



VP SA Storage Array User Guide

Release 23.09-SP3

Zadara

Apr 15, 2025

GETTING STARTED

1	Introduction	3
1.1	Intended Audience	3
1.2	Overview	3
1.3	VPSA Components	4
2	First steps	5
2.1	Register a Zadara Account	5
2.2	Creating a VPSA	5
2.3	Creating RAID Group, Pools, and Volumes	11
2.4	Managing your VPSA's resources	11
2.5	The VPSA Interface	19
2.6	System notifications	21
3	Drives	23
3.1	Replacing a Drive	23
3.2	Shredding a Drive	23
3.3	Viewing Drive properties	24
4	RAID Groups	27
4.1	Creating a RAID Group	27
4.2	Deleting a RAID Group	28
4.3	RAID Group Media Scan	28
4.4	Viewing RAID Group properties	30
4.5	Understanding Hot Spare Drives	32
4.6	Managing RAID Group Sync Speed	32
5	Pools	35
5.1	Understanding Storage Pools	35
5.2	Understanding Pool's Capacity	36
5.3	Viewing the List of Pools	40
5.4	Creating and Managing Pools	41
5.5	Viewing Pool properties	50
5.6	Managing Pool Capacity Alerts	57
5.7	Managing Pool Performance Alerts	61
5.8	Deleting a Pool	61
6	Volumes	63
6.1	Creating and Deleting a Volume	63
6.2	Filtering the List of Volumes	68
6.3	Attaching & detaching Volumes to Servers	68
6.4	Expanding a Volume	71

6.5	Managing SMB File History	71
6.6	Cloning a Volume	72
6.7	Online Volume Migration	73
6.8	Managing Data Reduction	74
6.9	Managing Encrypted Volumes	75
6.10	Audit Log Management	80
6.11	Volume File Lifecycle Management	83
6.12	Viewing Volume Properties	84
7	Servers	93
7.1	Filtering the List of Servers	93
7.2	Adding a Server	93
7.3	Viewing Servers Properties	109
7.4	Deleting a Server	112
8	Controllers	113
8.1	Failover	113
8.2	Restart File Services	113
8.3	Viewing Controller Properties	114
9	Remote VPSA	117
9.1	Connect to a remote VPSA	117
9.2	Remote VPSA Properties	118
10	Remote Object Storage	119
10.1	Connecting to Remote Object Storage	119
10.2	Viewing Remote Object Storage properties	120
11	Snapshots and Snapshot Policies	121
11.1	Managing Snapshots and Snapshot Policies	121
11.2	Filtering Snapshots	124
12	Mirroring	125
12.1	Creating a Local Mirror (on the same VPSA)	125
12.2	Creating a Remote Mirror	127
12.3	Replicate the same Volume to multiple destinations	128
12.4	Pause & Continue Remote Mirror	128
12.5	Managing Mirror Lifecycle	129
12.6	Viewing Remote Mirror Properties	133
13	Remote Clones	137
13.1	Remote cloning introduction	137
13.2	Remote cloning modes	137
13.3	Common remote cloning use cases	137
13.4	Connect source and destination VPSAs	139
13.5	Create a remote clone	139
13.6	Monitoring a remote clone	141
13.7	Attach a clone to a server	142
13.8	Data services	143
13.9	Delete a clone	143
14	Backup to Object Storage	145
14.1	Creating New Backups	145
14.2	Monitoring Backups	146
15	Restore from Object Storage	149

16 Images	153
16.1 Adding ZCS Engines	153
16.2 Creating an Image Repository	154
16.3 Creating a Container Image	154
17 Containers	157
17.1 Adding ZCS Engines	157
17.2 Creating a Container	158
17.3 Monitoring Containers	160
18 Container Memory Pools	163
18.1 Creating a Container Memory Pool	163
19 File Lifecycle	165
19.1 Understanding File Lifecycle Management Analytics	165
19.2 Reporting File Lifecycle Analytics	166
20 File Categories	169
20.1 Managing File Categories	169
21 Active Directory	171
21.1 Active Directory Authentication	171
22 LDAP	175
22.1 Enabling LDAP Authentication	175
23 Local NAS Users and Groups	177
23.1 Creating NAS Users	177
23.2 Creating SMB Users	178
23.3 Editing SMB Users Password	179
23.4 Enabling or Disabling User/Group/Project Quotas	180
23.5 Setting User/Group Quotas	183
23.6 Setting Project Quotas	184
24 Local users	187
24.1 User Roles	187
24.2 Adding and Deleting Users	187
24.3 Managing User Passwords	188
24.4 Managing Password Policy	189
24.5 Dual Factor Authentication	189
25 Performance	191
25.1 Understanding Performance Monitoring	191
25.2 The Performance Monitor	191
25.3 Customizing the Performance Monitor	193
26 Settings	195
26.1 General	195
26.2 Security	196
26.3 NAS	198
26.4 Metering	201
26.5 Container Service	201
26.6 Network	201
26.7 File Lifecycle Management	202
27 Diagnostics	203
27.1 Network Diagnostics	203

28 Logs	205
28.1 Access Log	205
28.2 Events Log	205
29 Support	207
30 CSI Driver	209
30.1 Zadara VPSA CSI for Kubernetes	209

Zadara provides the following VPSA offerings:

- VPSA Storage Array - a hybrid virtual array that supports both HDD and SSD drives
- VPSA Flash Array - an array optimized for efficient utilization of flash media

Explore our user guide to provision and administer your VPSA Storage Array.

Getting started Provision and start using Zadara's VPSA Storage Array

Volumes Configure your VPSA's Volumes

Servers Configure your VPSA's Servers

NAS Access Control Manage your NAS Users Access Control

INTRODUCTION

This documentation presents information specific to Zadara Storage SAN and NAS products.

✓ **Note:** You can find information specific to Zadara Storage's VPSA Object Storage product in the VPSA Object Storage [User Guide](#).

1.1 Intended Audience

This document is intended for end users and storage administrators subscribers of Zadara Storage's Enterprise Storage-as-a-Service product, called VPSA Storage Array and VPSA Flash Array, in both public and private clouds.

1.2 Overview

Zadara Storage Cloud was architected from the ground up to build the first "Enterprise-Storage-as-a-Service Data Storage System for the Cloud" with the following key targets:

- Enterprise quality, resilient, highly available, consistent performance storage for the most demanding data center application workloads
- Consumed as a Service - flexible, dynamic and billable
- Scale out - grow to hundreds of Storage Nodes, thousands of drives and multi-Petabyte Storage
- True Multi-tenancy - End-user controlled privacy and security. Separate workloads, resource allocation, and management per tenant, such that each tenant truly experiences "no noisy neighbors" secure storage.
- Universal Storage - Supports all data services on one common infrastructure: Block, File, Object

Starting with release 18.07 Zadara provides the following offerings:

- VPSA Storage Array - a hybrid virtual array that supports both HDD and SSD drives
- VPSA Flash Array - an array optimized for efficient utilization of flash media

Most of the content in this guide is for both the VPSA Storage Array and VPSA Flash Array. Sections and features that are specific to one offering only are marked accordingly.

Both VPSA Storage Array and VPSA Flash Array provide:


- Unified storage with both Block volumes and File shares exposure
- Data Protection (RAID-1, 10)

- Advanced Data management (Thin Volumes, Snapshots & Clones, Remote Mirroring, Built-in Backup to Object Storage, Built-in Anti-Virus service, SSD Flash Cache, etc....)
- Security features such as Dual factor authentication, Role-based access control, Data-at-Rest encryption and Data-in-Flight encryption
- Ability to run IO intensive applications as a Docker compatible container within the storage controller
- Management GUI and Rest API
- Flexibility to grow and shrink resources such as CPU, RAM, cache and drive allocations.

All of which is reliable, secured, private and consumed as a service.

VPSA Flash Array In addition Zadara VPSA Flash Array provides:

- Data reduction mechanism based on inline deduplication and compression to save on capacity of relatively expensive SSD media
- Tiered storage pools (SSD/HDD or SSD/S3 or azure API compatible object storage) with automatic data movement according to IO pattern.

 **Tip:** Virtual Private Storage Array (VPSA) is the first Software Defined, Enterprise Storage-as-a-Service. It is an elastic and private Block and File Storage System which provides Enterprise-grade data protection and data management storage services. As the VPSA Administrator, you will appreciate the level of control you have over the storage system while leveraging the benefits of consuming it as a service.

1.3 VPSA Components

1.3.1 Zadara Provisioning Portal

The Zadara Provisioning Portal is your gateway to the Zadara Storage ecosystem through which you can create, view and modify your VPSA configurations (engines, drives, Cache, etc....) on multiple Clouds that Zadara Storage offers.

1.3.2 Virtual Controller

A Virtual Controller (VC) is a Virtual Machine with dedicated CPUs & RAM which runs the VPSA IO stack and control stack. Two VCs are paired in an Active-Standby pair for high availability. The VC maintains a sophisticated and granular block-level mapping layer from virtual to physical address spaces, thus enabling enterprise-level data management capabilities like Thin Provisioning, Snapshots, Cloning and Remote Mirroring.

The VCs provide GUI and REST API end points for management and control.

1.3.3 Dedicated Drives

The Zadara Storage Cloud Orchestrator assigns dedicated drives for each VPSA. The drives are provisioned from different Storage Nodes (SN) for maximum redundancy and performance. Each drive is exposed as a separate iSCSI target from the SN and is LUN masked only to the VPSA's VCs. Your QoS is guaranteed because neighbors, with provisioned drives adjacent to yours, cannot access your drives, impact your performance or compromise your privacy and security.

FIRST STEPS

This chapter contains step-by-step instructions to create a VPSA and then to configure its storage properties from the Zadara's Provisioning Portal.

Important: Zadara's web applications allow only TLS 1.2 and higher, which is the recommended TLS level by industry standards. The TLS (Transport Layer Security) protocol secures transmission of data over the internet using standard encryption technology.

2.1 Register a Zadara Account

To register for a new Zadara account, go to <https://manage.zadarastorage.com/register/> and complete the registration form. If you wish to provision your new VPSA in a private location please use the URL provided by Zadara for the local Provisioning Portal instance.

2.2 Creating a VPSA

1. Log on to your **Zadara Provisioning Portal** at <https://manage.zadarastorage.com>, or at your private cloud, using your username/email and password.

Important: It is recommended to enable MFA (Multi-Factor Authentication) in order to add an additional layer of security to your account.

2. Click **Create New Service**.

The **Create Zadara Service** dialog opens.

Select:

- **Storage Array** to create a hybrid VPSA Storage Array.
- **Flash Array** to create a performance-optimized VPSA Flash Array with built-in inline data reduction support.

Click **Next** to progress to your selection.

✓ **Note:** VPSA Object Storage creation is described in the [VPSA Object Storage User Guide](#).

2.2.1 Creating a VPSA Storage Array

Scope: VPSA Storage Array

The **Create Storage Array** definitions dialog opens after selecting **Storage Array** in the Provisioning Portal's **Create Zadara Service** dialog.

1. **Select Provider and Define Name** dialog:

✓ **Note:** From the public cloud Provisioning Portal you can provision and manage all of your VPSAs, even if they are connected to different cloud provider groups and provider sites.

Enter the following fields:

- a. **Provider Group** - Select the cloud provider group for your deployment.
- b. **Provider** - Select the provider site for your deployment.
- c. **VPSA Name** - Give the VPSA a name. This is how it will appear in the Cloud Console and in the VPSA GUI. If you plan on having multiple VPSAs, you might want to give it as detailed a name as possible.
- d. **VPSA Description** (Optional) - Give the VPSA a description.

Click **Next** to continue to the **Select Engine Type and Drives** configuration.

2. **Select Engine Type and Drives** dialog:

Enter the following fields:

- a. **Engine Type** - The Zadara IO engine type defines the compute characteristics of your VPSA's Virtual Controllers (VCs). Each engine type defines the following characteristics:
 - Number of CPUs that are assigned to your VPSA's VCs.
 - Amount of RAM that is assigned to your VPSA's VCs.
 - Default size of protected SSD Cache.

When selecting the IO engine, take into account the capacity planned for this VPSA. Each engine has a limit to the number of drives it can support, and to the total raw capacity of the VPSA.

You can change the Zadara engine type (upgrade or downgrade) at any time throughout the lifetime of your VPSA according to your application's needs, on condition that you stay within the maximum limits of the engine type you are moving to.

✓ **Note:** The compute resources (CPU, RAM and Cache) are dedicated to your VPSA, which ensures consistent performance and isolation from other tenants' workloads and behavior.

- b. **Flash Cache** - (for VPSA Storage Array engines larger than 200) - From the dropdown, select the amount of flash cache to allocate to the VPSA. Note that each VPSA engine is provided with a minimum amount of cache memory. The extended cache is allocated in 200GB increments.
- c. **Drive Quantities** - Select the type and number of drives that you would like allocated to your VPSA.
 - The Zadara Cloud Orchestrator allocates dedicated drives.
 - Drives are allocated from as many different SNs as possible to provide maximum redundancy for your VPSA's RAID groups.

- There is a limit to the number of drives per Zadara IO engine type. The larger the engine is, the more drives you can add. There is also a limit to the total raw capacity of all drives. Make sure that the total capacity of all selected drives is within the limit.

The following table lists the maximum drives per VPSA Storage Array engine type:

IO Engine Type	Maximum # of Pools	Maximum # of Drives	Maximum Raw Capacity
200	3	5	24 TB
400	3	10	70 TB
600	3	20	140 TB
800	3	30	180 TB
1000	3	40	240 TB
1200	3	60	300 TB
1600	3	80	400 TB
2400	3	80	800 TB
3600	3	80	1000 TB

✓ **Note:** While it is technically possible to create a VPSA using the 200 engine to fulfill basic storage needs, it is important to note that such a setup may not be suitable for demanding production workloads. A smaller storage solution typically lacks the necessary scalability and performance capabilities required to handle the rigorous demands of a production environment. In scenarios where data growth, high-volume transactions, or complex data processing are involved, it is strongly advised to start with a larger engine that can ensure optimal performance and reliable storage management. Prioritizing a well-designed storage infrastructure can prevent potential bottlenecks and other issues that may hinder the smooth operation of critical business applications.

- d. **Create RAID-10** – By default, at VPSA creation time RAID-10 pools are automatically created, one pool per type of selected drive. The pool includes all the selected drives of each type. If you want to create different pools settings, uncheck this checkbox, and manually create your RAID groups and pools as described in [Creating RAID Group, Pools, and Volumes](#).

Click **Next** to continue to the **Advanced Services** configuration.

3. **Advanced Services** dialog:

Enter the following fields:

- a. **Container Services Engine** – The Zadara Container Services (ZCS) Engine defines the compute resources of the VPSA's Virtual Controllers that are allocated for Docker containers within this VPSA. Refer to [Containers](#) for details about Zadara Container Services.
- b. **File Lifecycle Management** - Enable file lifecycle management and analytics.

✓ **Note:** File lifecycle management indexing consumes some VPSA compute and memory resources.

- A dedicated indexing repository is created in the VPSA, for file lifecycle management.
- Additional SSD volumes are allocated from the VPSA's existing resources to support file lifecycle management.

- c. **Fibre Channel Support** – Check this checkbox if you will be connecting hosts to this VPSA over FC SAN.

Click **Next** to continue to review and confirm your selections.

4. **Review and Confirm** dialog:

- a. After selecting the VPSA characteristics, review the displayed summary.

You can click **Back** to return to previous dialog screens to modify your previous selections.

- b. Press the **Create** button to confirm the VPSA creation request.
 - The requested VPSA will appear in the Provisioning Portal's Service Inventory list, with the **Awaiting Approval** status.
 - Completing the VPSA creation requires the approval of a Zadara Storage Cloud admin. Once approved, the new VPSA takes only a few minutes to launch. During that time you'll see your VPSA with the **Launching** status and spinner, until launched.

5. VPSA First Access:

- a. When the VPSA is ready, you'll receive an email with a temporary passcode at your registered email address.
- b. To access the VPSA GUI, in the Provisioning Portal select the VPSA and under **Properties**, click the **Management Console** link.

✓ **Note:** By default, the VPSA interfaces are accessible to the storage front-end network only.

If you want to access it using a public IP, refer to the [Assigning Public IPs](#) section in this guide.

- c. Use your registered username or email address and the temporary passcode to enter the VPSA GUI. You will be immediately prompted to set a new password for your VPSA user account.

2.2.2 Creating a VPSA Flash Array

Scope:VPSA Flash Array

The **Create Flash Array** definitions dialog opens after selecting **Flash Array** in the Provisioning Portal's **Create Zadara Service** dialog.

1. Select Provider and Define Name dialog:

✓ **Note:** From the public cloud Provisioning Portal you can provision and manage all of your VPSAs, even if they are connected to different cloud provider groups and provider sites.

Enter the following fields:

- a. **Provider Group** - Select the cloud provider group for your deployment.
- b. **Provider** - Select the provider site for your deployment.
- c. **VPSA Name** - Give the VPSA a name. This is how it will appear in the Cloud Console and in the VPSA GUI. If you plan on having multiple VPSAs, you might want to give it as detailed a name as possible.
- d. **VPSA Description** (Optional) - Give the VPSA a description.

Click **Next** to continue to the **Select Engine Type and Drives** configuration.

2. Select Engine Type and Drives dialog:

Enter the following fields in the VPSA Flash Array's **Select Engine Type and Drives** dialog:

- a. **Engine Type** - The Zadara IO engine type defines the compute characteristics of your VPSA's Virtual Controllers (VCs). Each engine type defines the following characteristics:
 - Number of CPUs that are assigned to your VPSA's VCs.

- Amount of RAM that is assigned to your VPSPA's VCs.

When selecting the IO engine, take into account the capacity planned for this VPSPA. Each engine has a limit to the number of drives it can support, and to the total raw capacity of the VPSPA.

You can change the Zadara engine type (upgrade or downgrade) at any time throughout the lifetime of your VPSPA according to your application's needs, on condition that you stay within the maximum limits of the engine type you are moving to.

✓ **Note:** The compute resources (CPU, RAM and Cache) are dedicated to your VPSPA, which ensures consistent performance and isolation from other tenants' workloads and behavior.

b. **Drive Quantities** – Select the type and number of drives that you would like allocated to your VPSPA.

- The Zadara Cloud Orchestrator allocates dedicated drives.
- Drives are allocated from as many different SNs as possible to provide maximum redundancy for your VPSPA's RAID groups.
- There is a limit to the number of drives per Zadara IO engine type. The larger the engine is, the more drives you can add. There is also a limit to the total raw capacity of all drives. Make sure that the total capacity of all selected drives is within the limit.

The following table lists the maximum drives and capacity per VPSPA Flash Array Engine type:

✓ **Note:** Due to VPSPA Flash Array data reduction, the capacity limit per engine depends on both the physical capacity of the drives and the customer virtual capacity (as seen by the hosts), before any data reduction.

More about VPSPA Flash Array capacities: [Understanding Pool's Capacity](#)

IO Engine Type	Maximum # of Pools	Maximum # of Drives	Maximum Raw Capacity	Maximum Provisioned Capacity
H100	1	60	280 TB	140 TB
H200	2	80	440 TB	220 TB
H300	2	120	800 TB	400 TB
H400	2	140	1000 TB	500 TB

The high tier (also known as tier 0) is the more performant tier of the VPSPA, and comprises in-array SSD/NVME drives.

The low tier (also known as tier 1) is the capacity-oriented tier of the VPSPA.

It can be implemented via in-array SATA/NLSAS drives, or by connectivity to a remote object storage container.

- **SSD Storage Class:** The number of SSD type drives for Tier 0 (high tier).
- **HDD Storage Class:** The number of HDD type drives for Tier 1 (low tier).

For more information on tiers see **Tiers** in [Understanding Storage Pools](#).

✓ **Note:** The above capacities depend on the type of the pool(s) used.

The numbers shown are the limits of the aggregated size of all pools of type Throughput-Optimized.

See [Creating a Pool](#) for details

- c. **Create RAID-10** – By default, at VPSA creation time RAID-10 pools are automatically created, one pool per type of selected drive. The pool includes all the selected drives of each type. If you want to create different pools settings, uncheck this checkbox, and manually create your RAID groups and pools as described in [Creating RAID Group, Pools, and Volumes](#).

Click **Next** to continue to the Advanced Services configuration.

3. **Advanced Services** dialog:

Enter the following fields:

- a. **Container Services Engine** – The Zadara Container Services (ZCS) Engine defines the compute resources of the VPSA's Virtual Controllers that are allocated for Docker containers within this VPSA. Refer to [Containers](#) for details about Zadara Container Services.
- c. **File Lifecycle Management** - Enable file lifecycle management and analytics.

✓ **Note:** File lifecycle management indexing consumes some VPSA compute and memory resources.

- A dedicated indexing repository is created in the VPSA, for file lifecycle management.
 - Additional SSD volumes are allocated from the VPSA's existing resources to support file lifecycle management.
 - Usage of the File Lifecycle Management service incurs additional charges.
-

- d. **Data Reduction Bundle** - Enable support for:

- Inline data compression
- Inline data deduplication

✓ **Note:** Usage of the Data Reduction service incurs additional charges.

Click **Next** to continue to review and confirm your selections.

4. **Review and Confirm** dialog:

- a. After selecting the VPSA characteristics, review the displayed summary.

You can click **Back** to return to previous dialog screens to modify your previous selections.

- b. Press the **Create** button to confirm the VPSA creation request.
 - The requested VPSA will appear in the Provisioning Portal's Service Inventory list, with the **Awaiting Approval** status.
 - Completing the VPSA creation requires the approval of a Zadara Storage Cloud admin. Once approved, the new VPSA takes only a few minutes to launch. During that time you'll see your VPSA with the **Launching** status and spinner, until launched.

5. **VPSA First Access:**

- a. When the VPSA is ready, you'll receive an email with a temporary passcode at your registered email address.
- b. To access the VPSA GUI, in the Provisioning Portal select the VPSA and under **Properties**, click the **Management Console** link.

✓ **Note:** By default, the VP SA interfaces are accessible to the storage front-end network only.
If you want to access it using a public IP, refer to the [Assigning Public IPs](#) section in this guide.

- c. Use your registered username or email address and the temporary passcode to enter the VP SA GUI. You will be immediately prompted to set a new password for your VP SA user account.

2.3 Creating RAID Group, Pools, and Volumes

By default a new VP SA is created with all its drives configured in RAID Groups, and a Pool per each drives type. If the automatic pools satisfy your needs, go directly to the volumes creation below.

Otherwise follow the RAID Group and Pool creation instruction:

1. Create a RAID Group to define the level of data protection needed.
For more details, see [Creating a RAID Group](#).
2. Create a storage pool by using aggregated capacity from one or more RAID Groups.
For more details, see [Creating a Pool](#).
3. Create an iSCSI / FC / NFS / SMB Thin Provisioned Volume to be used by your servers.
For more details, see [Creating and Deleting a Volume](#).
4. Add a server. The server object represents the host using the storage volume.
Follow the instructions depending on the OS and connectivity of your server: [Adding a Server](#)
5. Attach the Volume to a Server.
For more details, see [Attaching & detaching Volumes to Servers](#)

Congratulations! You have a new VP SA provisioned and ready to use.

The following sections describe the various capabilities and services of your VP SA in detail.

2.4 Managing your VP SA's resources

You create, add, change, delete and manage the resources comprising your VP SAs via **the Zadara Provisioning Portal**.

This section describes the operations available in the Provisioning Portal (<https://manage.zadarastorage.com>).

2.4.1 Adding and removing Disk Drives

To add drives to your VPSA go to the Provisioning Portal, select the VPSA, and from the **Actions** dropdown select **Add Drives**.

- Select the number of drives from each available drive type you wish to add to your VPSA, and press **Add**. Keep in mind the RAID Groups you are going to build.
- This operation requires the approval of a Zadara Storage Cloud Admin. Once approved, you'll see the number of drives in the Provisioning Portal update accordingly. If you then refresh the **Drives** page in the VPSA GUI, the new drives will be displayed.

You can remove unused drives (indicated by the **Available** status) from within the VPSA.

In the VPSA GUI, go to the **Drives** page, select the drive you wish to remove and press **Remove**.

If you wish to remove a drive that is part of a RAID Group, you first need to replace it with another drive as described in [Replacing a Drive](#).

2.4.2 Managing Zadara Engines

The Zadara IO Engine type defines the following characteristics of your VPSA's Storage Controllers:

- **Dedicated CPU and memory resources** - These are dedicated solely to your VPSA. These resources are not shared with any other VPSA or tenant within the Zadara Storage Cloud.
- **Flash Cache Size** (VPSA Storage Array only) - Each VPSA is provisioned with a Flash Cache partition to be used for both metadata and read/write caching. The SSD cache partition is protected using RAID-1, where each mirror copy resides on a different SN, thus ensuring cache resilience to SN failure. Each Engine type is provided with a base SSD cache partition size. You can request additional flash capacity for caching. For more details see [Adjusting Flash Cache \(VPSA Storage Array\)](#).
- **Maximum number of drives** - The maximum number of drives that can be allocated to each VPSA engine type.

The following Zadara IO Engines are available for VPSA Storage Array:

IO Engine Type	Dedicated Compute Resources	Base Flash Cache	Max # of Drives	Max Raw Capacity
200	2 CPU, 6 GB RAM	20 GB	5	24 TB
400	4 CPU, 12 GB RAM	20 GB	10	70 TB
600	6 CPU, 20 GB RAM	40 GB	20	140 TB
800	8 CPU, 28 GB RAM	60 GB	30	180 TB
1000	10 CPU, 36 GB RAM	80 GB	40	240 TB
1200	12 CPU, 52 GB RAM	100 GB	60	300 TB
1600	16 CPU, 68 GB RAM	120 GB	80	400 TB
2400	24 CPU, 100 GB RAM	180 GB	80	800 TB
3600	36 CPU, 144 GB RAM	240 GB	80	1000 TB

The following Zadara IO Engines are available for VPSA Flash Array:

IO Engine Type	Dedicated Compute Resources	Maximum # of Drives	Maximum Raw Capacity	Maximum Provisioned Capacity
H100	12 CPU, 72 GB RAM	60	280	140 TB
H200	24 CPU, 116 GB RAM	80	440	220 TB
H300	36 CPU, 176 GB RAM	120	800	400 TB
H400	48 CPU, 236 GB RAM	140	1000	500 TB

✓ Note:

- Due to data reduction in a VPSA Flash Array, the capacity limit per engine depends on both the physical capacity of the drives and the provisioned capacity, before any data reduction.
- The above capacities depend on the type of the pool used. The numbers shown are the limits of the aggregated size of all pools of type Throughput-Optimized. See [Creating a Pool](#) for details.

More about VPSA Flash Array capacities: [Understanding Pool's Capacity](#)

The following Zadara Container Services Engines (see: [Containers](#)) are available:

Zadara ZCS Engine Type	Dedicated compute resources
01	2 CPU, 512 MB RAM
02	2 CPU, 1 GB RAM
04	4 CPU, 2 GB RAM
06	6 CPU, 4 GB RAM
08	8 CPU, 8 GB RAM

To change both types of Zadara Engines, in the Provisioning Portal select the VPSA, and from the **Actions** dropdown select **Change Engine Type**.

When selecting any engine larger than 200 you can also select the required flash cache size for that engine. For Flash Cache limits see [Adjusting Flash Cache \(VPSA Storage Array\)](#).

Completing this operation requires the approval of the Zadara Storage Cloud Admin.

The Zadara Engine upgrade/downgrade process may take a few minutes. During that time, the VPSA status will change to **Upgrade Pending**.


When the process completes, the VPSA status will change back to **Ready**.

2.4.3 Managing Virtual Networks

The Zadara cloud provides a flexible and dynamic virtual networking infrastructure that can be tailored to meet multiple storage architecture and use cases.

Each cloud tenant is allocated with one or more “Virtual Networks” which is a set of available IP addresses within a specific network segment. Virtual networks are allocated for a specific cloud tenant and within a specific available cloud VLAN.

The below diagram depicts the relationship between cloud tenants, virtual networks and VLANs:



In case a VPSA serves as storage for servers on different networks, the VPSA can be plugged on multiple “Virtual Networks”. Both block volumes and NAS shares can simultaneously be exposed through one or more Virtual Networks.

Each VPSA is created with a primary network for its front end (hosts connectivity). This network is routable and is mandatory.

You can manage your virtual networks from the Zadara Provisioning Portal. In the Provisioning Portal, click your account at the top right, and in the dropdown, click **Network Management**.

To create another virtual network press **Next**, and fill in the requested parameters such as: CIDR, Gateway, IP Range, and whether IPv4 or IPv6 should be used.

You can add and remove secondary networks to the VPSA. The VPSA internally maintains a “Virtual Networks Interface” (VNI) that connects into each virtual networks.

To add a Virtual Network Interface, in the Provisioning Portal select the VPSA, and from the **Actions** dropdown select **Add Virtual Network Interface**.

To remove a Virtual Network Interface, in the Provisioning Portal select the VPSA, and from the **Actions** dropdown select **Remove Virtual Network Interface**.

✓ **Note:**

- Number of VNIs per VPSA is limited to 5.
- VPSA REST API/GUI is accessible through any VNI.
- Only Primary VN IP is registered in DNSimple
- VPSA can't have two VNI with the same VLAN.

✓ Note:

- Only “Primary Virtual Network” is a routable network. Remaining virtual networks are not routable. There are some limitations on the remaining virtual networks:
 - Active Directory can be joined only through “primary virtual network”.
 - Backup (B2OS), Mirror, Remote Clone through FE network are only allowed via the “primary virtual network”.
 - ZCS container services exposed through FE network can be done only on “primary virtual network”.
 - “iSER” host connectivity is available only on the “Primary Virtual Network”.
-

2.4.4 Assigning Public IPs

By default you cannot access the VPSA from the public Internet for security and privacy reasons. The VPSA Front-End IP address which is used for VPSA management (via GUI and REST API) and for data IO workload (host connectivity via iSCSI/NFS/SMB protocols), is allocated on the Zadara Storage Cloud “Front-End” network 10GbE interface which is routable only from the Cloud Servers network. Servers outside of your Cloud Servers network cannot reach this IP address. This means you cannot access your VPSA GUI from the Internet.

A typical use case requiring Public IP addresses is when you’re running Asynchronous Remote Mirroring between two VPSAs in different regions, between on premise and cloud deployments or even between different Cloud Providers for Disaster Recovery (DR). Communication between the VPSAs is done via an authenticated and encrypted channel over the public Internet, thus requiring Public IPs.

To assign a Public IP address to your VPSA, go to the Provisioning Portal and press the **Assign Public IP** link. You can see the assigned IP address in your VPSA details in the Provisioning Portal and in the VPSA GUI, under **Settings > General > Public IP**. To remove it, click **Remove Public IP** in the Provisioning Portal.

✓ Note: Access to the VPSA GUI and API is blocked through the Public IP for security reasons.

✓ Note: NAT’d server IP connections are not supported for iSCSI, NFS, and SMB protocols over the Public IP.

2.4.5 Adjusting Flash Cache (VPSA Storage Array)

Scope: VPSA Storage Array

Each VPSA is provisioned with a base flash cache partition, which is utilized by the VPSA for both metadata and read/write caching. The initially assigned default SSD cache size is also the minimal cache size for a given Zadara Engine. The flash cache partition is protected using RAID-1, where each mirror copy resides on a different SN, thus ensuring cache resilience to multiple types of failure.

On top of the base flash cache described above, you can add an extended cache. The VPSA extended flash cache size is elastic, so you can increase or decrease the cache size according to the needs of your workload.

Each Engine type has a minimum (default) and maximum SSD Cache size, as shown in the table below:

Zadara Engine	Base Flash Cache	Default Extended Flash Cache Size	Max Extended Flash Cache Size
200	20 GB	0 GB	0 GB
400	20 GB	200 GB	400 GB
600	40 GB	400 GB	800 GB
800	60 GB	600 GB	1200 GB
1000	80 GB	800 GB	1600 GB
1200	100 GB	1200 GB	2400 GB
1600	120 GB	1600 GB	3200 GB
2400	180 GB	1600 GB	3200 GB
3600	240 GB	1600 GB	3200 GB

To change the Extended Flash Cache size for your VPSA, in the Provisioning Portal select the VPSA, and from the **Actions** dropdown select **Change Flash Cache**.

2.4.6 Enabling or Disabling Data Reduction Bundle

Scope: VPSA Flash Array

The Data Reduction Bundle provides the option to activate one, both or neither of inline compression and deduplication per Volume in a VPSA Flash Array.

A VPSA Flash Array's Data Reduction Bundle can be enabled or disabled at any time.

Enabling Data Reduction Bundle

To enable the ability to configure a VPSA Flash Array's Volumes' inline compression and deduplication options:

1. In Zadara's Provisioning Portal, go to **Console > Service Inventory** and locate the VPSA in the inventory list.
2. Click the VPSA to select it, and from its **Actions** dropdown select **Enable Data Reduction Bundle**.
3. To activate inline compression and deduplication on the VPSA's existing Volumes:
 1. In the VPSA's UI, go to the **Volumes** page.
 2. Select the Volume from the Volumes list, and in the south pane's **Properties** tab set the **Compress** and **Dedupe** properties to **Enable**.
 3. Repeat this procedure for all Volumes in the VPSA that should have inline compression and deduplication.

Disabling Data Reduction Bundle

To disable the ability to configure a VP SA Flash Array's Volumes' inline compression and deduplication options:

1. In Zadara's Provisioning Portal, go to **Console > Service Inventory** and locate the VP SA in the inventory list.
2. Click the VP SA to select it, and from its **Actions** dropdown select **Disable Data Reduction Bundle**.

2.4.7 Upgrading your VP SA

From zStorage 23.09, the self-service VP SA software update process is available in Zadara's Provisioning Portal. The self-service option streamlines and accelerates the VP SA update process. VP SA owners can now determine VP SA update schedules at their own convenience. The self-service software update can be run immediately, or scheduled to run in 30 minutes from the current time, up to 7 days from the current date and time.

Important:

- Per Zadara's Software Lifecycle Policy, only VP SAs that are running software versions that have not entered the **End of Support** phase are eligible for self-service software updates.
 - The update takes up to 20 minutes on the standby controller, followed by the VP SA initiating a single failover, during which IOs are paused for several seconds and resume automatically.
-

In the **Service Inventory** list of the Provisioning Portal's Console, the **Update Available!** badge displays next to the VP SA's **Service Type**.

To update a VP SA:

1. In the Provisioning Portal, in the Console's Service Inventory list, either click the eligible VP SA's **Update Available!** badge, or click the VP SA row to open the VP SA's details pane on the right, and in the **Actions** dropdown menu, select **VP SA Software Update**.

The update process runs checks to verify that the VP SA's state allows an upgrade. For example, if the VP SA's state is **Degraded** or the VP SA is running an older software version that cannot be upgraded directly, a relevant message is displayed directing users to contact Zadara Support.

On successful verification of the VP SA for the self-service update, the **VP SA Software Update** dialog opens. The VP SA's current running version is displayed, as well as the new zStorage software version that is available for update, together with the link to its release notes. This is followed by a brief description of the flow of the marked update process (**Now** or **Schedule an update for later**).

2. Select the update schedule:

- **Now**

Important: On confirmation, the VP SA software update process is executed immediately, and takes up to 20 minutes to complete.

The standby controller is updated with the latest software version. This is followed by the VP SA initiating a single failover to complete the software update, during which IOs are paused for several seconds and resume automatically.

1. Enter your Provisioning Portal login password, and click **UPDATE**.

The **Update software version started** message is displayed.

2. Click **OK** to close the dialog.

✓ **Note:** On the next refresh of the Console's Service Inventory list, the VPSA's **Status** display changes to **Upgrading Version**.

On completion of the update process, the VPSA's **Status** display changes briefly to **Reconfiguring**, and then back to **Created**.

- **Schedule an update for later**

1. Enter a date and time, in the range of 30 minutes from the current time, up to 7 days from the current date and time.

The default is one hour from the current time.

2. Enter your Provisioning Portal login password, and click **UPDATE**.

Important: When scheduling an update for your VPSA software at a later time:

1. The VPSA's standby controller is updated **immediately** with the new software version, in preparation for the scheduled failover to the new version. The update takes about 20 minutes.
 2. Prior to a scheduled update, the system sends email reminders to the VPSA owner:
 - A notification 8 hours before the scheduled update time.
 - A final notification 10 minutes before the scheduled update time.
 3. At the scheduled time, the VPSA initiates a single failover to complete the software update during which IOs are paused for several seconds, and resume automatically.
-

To **reschedule** a scheduled VPSA upgrade:

1. Click the eligible VPSA's **Update Available!** badge, or click the VPSA row to open the VPSA's details pane on the right, and in the **Actions** dropdown menu, select **VPSA Software Update**.
2. In the **Reschedule VPSA Software Update** dialog that opens, select one of:

- **Now**

The previously scheduled update had already updated the standby controller with the new software version, in preparation for its scheduled failover.

Select **Now** to trigger an immediate failover to the standby controller to complete the update, during which IOs are paused for several seconds, and resume automatically.

- **Schedule an update for later**

The previously scheduled update had already updated the standby controller with the new software version, in preparation for its scheduled failover.

Enter a new date and time to reschedule the failover to the updated standby controller, in the range of 30 minutes from the current time, up to 7 days from the current date and time.

3. Enter your Provisioning Portal login password.
4. Click **Reschedule** to apply the new schedule, or **Cancel** to remain with the original schedule.

The self-service update process:

- Logs the self-service VPSA software update request in the Provisioning Portal's access logs.
- Sends an automatic email notification to VPSA owners on completion of the VPSA software update, advising them of the successful VPSA software update.

2.4.8 Hibernating your VP SA


You can hibernate your VP SA when it is not in use for some period of time in order to reduce its associated service cost. While the VP SA is in a hibernated state you will only be billed for the drives, not the engine. Hibernating a VP SA involves the process of deleting its Virtual Controllers (the VP SA) while maintaining the data drives and all the necessary metadata to resume its operation at a later stage. **No data is lost!** The hibernated VP SA is not accessible to any GUI or REST API commands, nor will it present any iSCSI or NFS / SMB volumes. Resuming a hibernated VP SA only takes a few minutes.

To change the Extended Flash Cache size for your VP SA, in the Provisioning Portal select the VP SA, and from the **Actions** dropdown select **Hibernate**.

To resume access to the VP SA, in the Provisioning Portal select the VP SA, and from the **Actions** dropdown select **Restore**.

✓ **Note:** The **Hibernate** and **Restore** toggle depends on the current state of the VP SA

2.4.9 Deleting your VP SA

 **Warning:** Please note that deleting a VP SA is a significant action and should be done with caution. Ensure that you have backed up any critical data and that you are certain you want to proceed with the deletion

The VP SA owner can delete their VP SA using Zadara's Provisioning Portal in case it is no longer needed. The delete operation will delete all VP SA entities such as volumes,pools,raid-groups,drives and server records.

In order to delete a VP SA select the **Delete** action from the **Actions** dropdown menu in Zadara's Provisioning Portal. The delete operation will require an additional authentication confirmation.

Due to the sensitive nature of the delete operation, when the VP SA owner initiates a delete operation, it will generate a delete request that necessitates approval from a cloud administrator. Once approved, the VP SA will be permanently deleted.

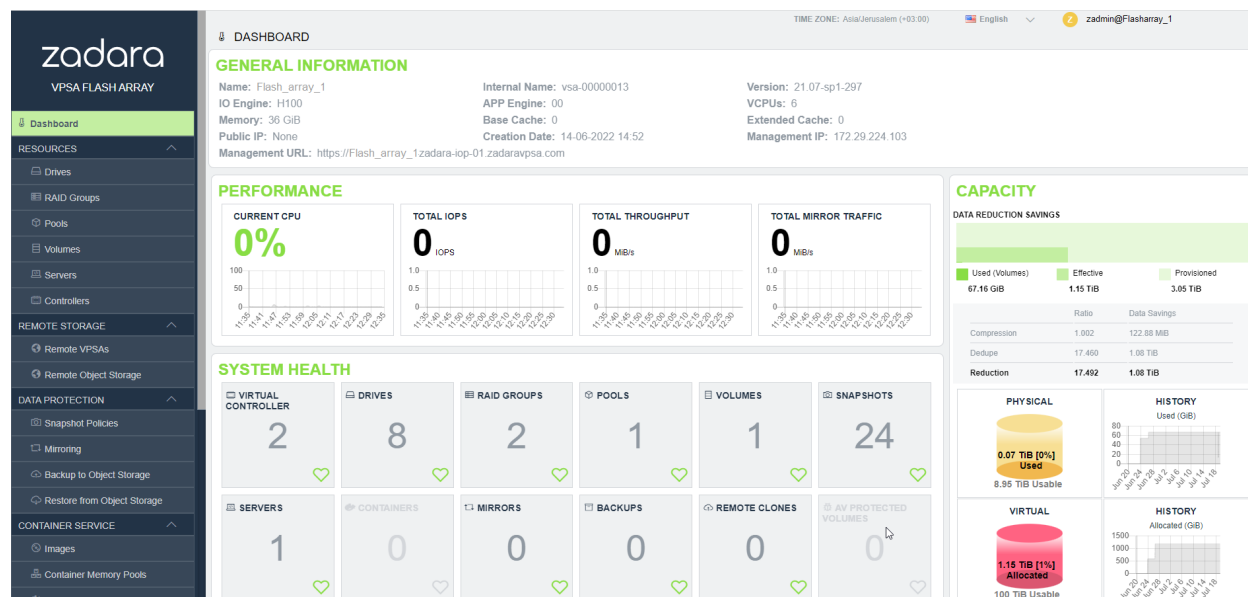
✓ **Note:** While the delete request is still pending, the VP SA will continue to operate normally and will continue to report its consumption to the cloud's billing services

2.5 The VP SA Interface

✓ **Note:** The VP SA management interface web application is supported in all modern browsers. We recommend using Google Chrome, Firefox or Microsoft Edge for an optimal user experience.

2.5.1 Understanding the VPSA Dashboard

VPSA Dashboard is the landing page, every time the GUI opens. It gives the overall state of the VPSA (Health, Capacity, Performance) at a glance. The Dashboard has the following components:



1. **VPSA Info:** General information of the VPSA such as its name, engine type and management IP address.
2. **System Health:** Shows the inventory of the objects managed by the system, such as Pools, Volumes, Mirrors, and so on. If all objects of a component are in the normal state, there is a green indicator on the component's tile. If there is situation that needs your attention a red indicator is shown on the component's tile, with the number of objects requiring attention.
3. **CPU:** Shows the CPU utilization of the active Controller of the VPSA, over time. This chart gives an indication of the load on the storage system.
4. **Capacity:** Shows the capacity state of the VPSA. The display is different between VPSA Storage Array and VPSA Flash Array. For the latter, it shows the capacity reduction saving. See [Understanding Pool's Capacity](#) for details.
 - a. Current capacity state
 - b. Capacity consumption over time during the last month
5. **Performance:** These charts show the aggregated performance of all Volumes.
 - a. Current IOPS (reads and writes) of all Volumes
 - b. IOPS activity during the last hour
 - c. Current throughput of all Volumes
 - d. Throughput of all volumes during the last hour
 - e. Current mirroring traffic of all mirrors (outbound and inbound)
 - f. Mirroring activity of all mirrors during the last hour

2.5.2 Understanding the VPSA GUI

The screenshot illustrates the VPSA GUI interface for managing drives. The left navigation panel (1) lists various system components. The center pane (2) shows a table of drives with the following data:

Name	Capacity	Storage Node	Type	Status	RAID Group	Zone
drive-000	3.49 TiB	sn-01	SSD	Normal	Pool1-r0	zone_0
drive-001	3.49 TiB	sn-02	SSD	Normal	Pool1-r0	zone_0
drive-020	3.49 TiB	sn-01	SSD	Available		zone_0
drive-021	3.49 TiB	sn-02	SSD	Available		zone_0
drive-002	5.46 TiB	sn-02	SAS	Normal	Pool1-r1	zone_0
drive-003	5.46 TiB	sn-01	SAS	Normal	Pool1-r1	zone_0
drive-018	5.46 TiB	sn-02	SAS	Available		zone_0
drive-019	5.46 TiB	sn-01	SAS	Available		zone_0

The south pane (3) provides detailed information for the selected drive (drive-000):

- Properties:** ID: volume-000000a2, Name: drive-000, Capacity: 3.49 TiB, Storage Node: sn-01, Type: SSD, UUID: 55CD2E4150D52CC2
- Metering:** (Empty)
- Logs:** (Empty)

The VPSA GUI provides full management and control of your VPSA. It contains the following main components (as numbered in the above screenshot):

- 1. Main Navigation Left Panel** – Traverse through the various VPSA entities. The selected entity is highlighted.
- 2. The Center Pane** – Displays a list of objects from the selected entity type (e.g. Drives in the above screenshot example) and for each object it displays the main properties.
- 3. The South Pane** – Displays detailed information regarding the selected object. All objects have at least 3 tabs:
 - a. Properties** – Detailed properties of the object.
 - b. Metering** – Typically IO workload metering info.
 - c. Logs** – List of event-log messages related to that object.
- 4. Logged-in username** – Displayed at the top right corner.
- 5. Selected Language** – Displayed on the top bar. You can use this drop down to change the displayed language.

2.6 System notifications

VPSA notifications, both informational and critical, necessitating user action, will be communicated via email to the service owner and other designated users set to receive notifications.

System notifications are categorized based on the following priorities:

- Urgent
- High
- Normal
- Low

2.6.1 Urgent priority notification

An alert that requires an immediate VP SA administrator action to ensure the storage service health or to restore the VP SA to normal operation, for example:

- Pool free capacity state
- System is pending for Master Encryption key from the administrator

2.6.2 High priority notification

An alert that requires awareness of the VP SA administrator to ensure the storage service health or other service issues that are currently being handled by Zadara's support team. In some cases VP SA administrator action is required.

Example for high priority notifications:

- Volume out of free space
- Duplicate Front-End IP discovery

2.6.3 Normal priority notification

An alert that requires the administrator's attention, yet does not necessarily have an immediate impact on the service. For example:

- Temporary Active Directory domain controller connectivity issue
- VP SA connectivity issue to a server that was set with connectivity test option

2.6.4 Low priority notification

A low-priority message with no impact on service health. While the message may require attention or action from administrators, its lower priority status indicates that it poses minimal or no threat to the current state of the service and can be addressed at the convenience of the administrators without causing service disruptions.

For example:

- Volume low capacity
- Volume auto-expand confirmation

DRIVES

3.1 Replacing a Drive

Press the **Replace** button on the Drives page to replace a drive. When selecting the replacement drive you must choose a drive that will not break the RAID Group redundancy (i.e. you cannot have two or more drives from the same Storage Node in a RAID Group). If you select a drive that has a different type or larger size than the other drives in the RAID Group, you will see a warning, but you can continue the operation.

You can replace a drive in any RAID Group whether the drive is healthy (Normal) or unhealthy (Failed).

You cannot replace a drive if the RAID Group is in a Resyncing state.

3.2 Shredding a Drive

Shredding is the process of erasing the data on a drive for security and privacy reasons by overwriting the entire drive with random data at least three times. Typically you will shred a drive before returning it to the Zadara Cloud or before deleting your VPSA.

You can only perform Shred on drives in Available status (i.e. not in a RAID group).

The Shredding progress appears in the drive status as “Shredding X%”.

You cannot remove a drive from a VPSA while it is being shredded. You need to either cancel the operation by pressing the Cancel Shred button, or wait until shredding is completed.

 **Caution: Shredding is irreversible!**

3.3 Viewing Drive properties

You can view the following properties and metering information in the Drives Details South Panel tabs:

Properties

Each drive displays the following Properties:

Property	Description
ID	An internally assigned unique ID.
Name	Drive name.
Capacity	Drive Capacity in MB.
Storage Node	The name of the Storage Node where the drive is physically located.
Type	SATA, SAS, or SSD
Status	The drive's status reflects the drive health as sensed by the Storage Node and by the VP SA RAID logic: <ul style="list-style-type: none"> • Available – The drive is healthy and free. • Normal – The drive is healthy and belongs to a RAID Group. • Absent – No access to the drive. • Failed – The Storage Node has reported failure accessing the drive. • Faulty – The VP SA RAID object has failed writing to or reading from this drive. • Recover Pending – The RAID Group has failed and the drive is awaiting recovery. • Shredding – The drive is being shredded.
RAID Group	Name of the RAID group that contains this drive.
Protection Zone	Displays the Protection Zone of the drive.
Usage	In-use or Available
Added	The date and time when the drive was added to the VP SA.
Modified	The date and time when the Drive object was last modified.

Metering

The Metering Charts provide live metering of the IO workload associated with the selected Drive.

The charts display the metering data as it was captured in the past 20 "Intervals." An interval length can be set to one of the following: 1 second, 10 Seconds, 1 Minute, 10 Minutes, or 1 Hour. The Auto button lets you see continuously-updating live metering info (refreshed every 3 sec).

The Following charts are displayed:

Chart	Description
IOPs	The number of read and write SCSI commands issued to the Drive, per second.
Bandwidth (MB\s)	Total throughput (in MB) of read and write SCSI commands issued to the Drive, per second.
IO Time (ms)	Average response time of all read and write SCSI commands issued to the Drive, per selected interval.

Logs

Displays all event logs associated with this Drive.

RAID GROUPS

4.1 Creating a RAID Group

VPSA RAID Groups define the level of protection against disk failure of the Pools and Volumes which contain the user's data. Careful consideration must be given when selecting the RAID level, along with the number and type of drives, in order to avoid potential impact on performance of your data. RAID groups always span across drives from different Storage Nodes, thus a RAID Group is resilient to both a single drive failure, as well as to a complete Storage Node failure.

To create your RAID Groups first select the Drives entity in the Main Navigation Panel (Left Panel) and then click the Create RAID Group button in the Center Panel.

Define the following attributes in the "Create RAID Group" dialog box:

- Enter the RAID Group name (you will later add it to a Pool so you may want to provide a meaningful name that describes the target usage of the Pool).

✓ **Note:** Objects names can be up to 128 chars long and can contain letters and digits, dashes "-" and underscores "_".

- Select Protection Type. Refer to the table below for a description of the various RAID levels.
- Select whether to allocate a drive as a Hot Spare for this RAID group. Adding Spare drive for RAID-1 groups is recommended. See [Understanding Hot Spare Drives](#) for more details about managing **Hot Spares**.
- Select the drives that will participate in the RAID Group. As noted in the table below, for RAID-1 a minimum of 2 drives is required.
- For maximum redundancy drives **MUST** be selected from different Storage Nodes so the VPSA will prevent you from doing otherwise.
- It is possible **but not recommended** to mix drives of different types in a single RAID Group.

✓ **Note:**

- RAID-5 is no longer supported.
- From zStorage version 23.09, the system does not allow RAID-6 RAID Group creation and RAID-60 Storage Pools, respectively. VPSAs that are currently using RAID-60 Storage Pools can be upgraded to version 23.09, and will also support storage expansions. For more information, contact <https://support.zadarastorage.com>.

4.1.1 Understanding RAID levels

The following RAID levels are supported:

RAID level	Description
RAID-1 (Mirroring)	RAID-1 mirrors the contents of one hard drive in the group onto another. If either hard drive fails, the other hard drive takes over and normal system operations are not interrupted. RAID-1, or Drive Mirroring, creates fault tolerance by storing duplicate sets of data on a minimum of two hard drives, and offers an excellent combination of data protection and performance. There must be 2 or 3 drives in a RAID-1 group. RAID-1 and RAID-10 are the most costly fault tolerance methods because they require 50 percent of the total combined drives capacity to store the redundant data.
RAID-10 (Mirroring and Striping)	RAID-10, or Drive Mirroring and Striping, is achieved in a VPSA by creating RAID-1 RAID Groups and striping them together at the Pool level. RAID-10 first mirrors each drive in the array to another, and then stripes the data across the mirrored pair. If a physical drive fails the mirror drive takes over and normal system operations are not interrupted. RAID 10 can withstand multiple simultaneous drive failures, as long as the failed drives are not mirrored to each other. RAID-10 creates fault tolerance by storing duplicate sets of data on a minimum of four hard drives and offers the best combination of data protection and performance. RAID-10 is the most costly fault tolerance method because it requires 50 percent of the total combined drives capacity to store the redundant data.

4.2 Deleting a RAID Group

The VPSA administrator can delete a specific RAID Group if it is not needed by clicking the **Delete** option under the RAID Group's top option menu.

✓ **Note:** A RAID Group can be deleted in case it is not allocated to a storage pool. The system will block deletion of allocated RAID groups.

4.3 RAID Group Media Scan

The VPSA performs regular media scans on the drives within each RAID group. A media scan is a proactive maintenance process that scans the underlying drives, identifies potential issues, and automatically tries to correct media errors when detected.

To minimize performance impact, the scans are performed on a RAID group basis with a rate limit, ensuring that the process is efficient and has minimal impact on the drives performance.

By default, the media scan is scheduled to run once a month, at 1:00am on a system-calculated different day of the month per VPSA in the same cloud.

To view or modify the automatically scheduled media scans, see [Media Scan Settings](#).

To bypass the automatically scheduled media scans, and trigger an on-demand immediate media scan, pause or resume a media scan on a specified RAID Group (whether a scheduled or on-demand scan), see [Starting, Pausing or Resuming a Media Scan](#).

4.3.1 Media Scan Settings

To view the VPSA's media scan settings:

1. In the navigation tree, select **Resources > RAID Groups**.

While a media scan is in progress on a RAID Group, its Status displays **Media Scan** and the scan's progress percentage.

On completion of the media scan, the RAID Group's Status changes to **Normal**.

2. In the **RAID Groups** window's top toolbar, select **Media Scan Settings**.

The **Media Scan Settings** dialog opens.

To modify the VPSA's media scan settings:

1. Adjust the relevant settings:

- **Periodic Media Scan:**

Toggle to activate or disable the scheduled periodic media scan.

Default: **ON**

- **Media Scan Schedule**

Frequency of automatically scheduled RAID Groups media scans.

One of:

- **Weekly**

Custom option to schedule media scans to automatically run on a user-selected **Day of week** (Sunday - Saturday) at 1:00am.

- **Monthly** (default)

Media scans run automatically on the **Day of month** (1 - 28) at 1:00am.

The default day of the month is system-calculated and differs between VPSAs in the same cloud, for the purpose of balancing the overhead. It can be overwritten by a user-selection.

The **Day of month** range is restricted to 1 - 28, to avoid discrepancies between longer and shorter months.

2. Optionally, click **Default** to reset the settings to their defaults.
3. Click **Submit** to save the changes, or **Cancel** to exit without applying changes to the settings.

4.3.2 Starting, Pausing or Resuming a Media Scan

The VPSA supports the option to bypass its automatically scheduled media scans:

- **Starting a Media Scan** an on-demand immediate media scan on a specified RAID Group, irrespective of whether the **Periodic Media Scan** toggle is activated or disabled.
- **Pausing a Media Scan** or **Resuming a Media Scan** on a specified RAID Group, irrespective of whether the media scan is a scheduled or an on-demand scan.

Starting a Media Scan

To start an on-demand immediate media scan on a RAID Group:

1. In the navigation tree, select **Resources > RAID Groups**.
2. In the **RAID Groups** window, select the RAID Group.
3. In the **RAID Groups** window's top toolbar, select **Start Media Scan**.

In the **Confirm** dialog, confirm starting an on-demand immediate media scan for the selected RAID Group.

The RAID Group's Status changes to **Media Scan**.

On completion of the media scan, the RAID Group's Status changes to **Normal**.

Pausing a Media Scan

To pause a media scan in progress on a RAID Group:

1. In the navigation tree, select **Resources > RAID Groups**.
2. In the **RAID Groups** window, select the RAID Group.

The RAID Group's Status displays **Media Scan**.

3. In the **RAID Groups** window's top toolbar, select **Pause Media Scan**.

In the **Confirm** dialog, confirm pausing the media scan that is in progress on the selected RAID Group.

The RAID Group's Status changes to **Media Scan paused**.

Resuming a Media Scan

To resume a paused media scan:

1. In the navigation tree, select **Resources > RAID Groups**.
2. In the **RAID Groups** window, select the RAID Group.

The RAID Group's Status displays **Media Scan paused**.

3. In the **RAID Groups** window's top toolbar, select **Resume Media Scan**.

In the **Confirm** dialog, confirm resuming the paused media scan for the selected RAID Group.

The RAID Group's Status changes to **Media Scan**.

On completion of the media scan, the RAID Group's Status changes to **Normal**.

4.4 Viewing RAID Group properties

The RAID-Group's details (properties and metering), are shown in the South Panel tabs:

Properties

Each RAID Group includes the following properties:

Property	Description
ID	An internally assigned unique ID.
Name	User assigned name. Can be modified anytime.
Comment	User free text comment. Can be used for labels, reminders or any other purpose
Protection	Selected RAID level - RAID-1.
Capacity	Total protected and usable capacity of the RAID Group.
Available Capacity	The RAID Group's usable capacity that is not allocated to any Pool.
Mirror Number	Number of mirror copies for RAID-1.
Status	<ul style="list-style-type: none"> • Normal – All drives are in sync • Resyncing X% – The RAID is in an initial rebuild process. • Degraded – One of the drives have failed. • Degraded Resyncing X% – The RAID is resyncing data following a drive recovery\replacement. • Repairing X% – Media Scan is in progress. • Repairing Paused – Media Scan is paused. • Failed – The array has lost too many drives and cannot serve Server IOs.
Created	Date & time when the object was created.
Modified	Date & time when the object was last modified.

Drives

Lists the disk Drives participating in the selected RAID Group. The following information is displayed per drive:

- Name
- Capacity (in GB)
- Location (Storage Node)
- Type (SAS/SATA/SSD/TBD)
- Status (Normal/Failed/TBD)
- Hot Spare (Yes/No)

Metering

The Metering Charts provide live metering of the IO workload associated with the selected RAID Group.

The charts display the metering data as it was captured in the past 20 intervals. An interval length can be one of the following: 1 second, 10 Seconds, 1 minute, 10 minutes, or 1 hour. The Auto button lets you see continuously-updating live metering info.

✓ **Note:** The metering info of the RAID Group doesn't include RAID-generated IOs, such as when doing a rebuild.

The following charts are displayed:

Chart	Description
IOPs	The number of read and write commands issued to the RAID Group per second.
Bandwidth (MB\s)	Total throughput (in MB) of read and write SCSI commands issued to the RAID Group per second.
IO Time (ms)	Average response time of all read and write SCSI commands issued to the RAID Group per selected interval.

Logs

Displays all event logs associated with this RAID Group.

4.5 Understanding Hot Spare Drives

When creating a RAID Group, you can decide whether you'd like to allocate hot spare drives to the RAID Group or not. You can change this selection at any time by clicking the **Add Spare** or **Remove Spare** buttons on a selected RAID Group in the VPSA GUI's **RAID Groups** page.

Allocating a hot spare drive for a RAID Group allows for immediate and automated drive replacement, with no human intervention, once the VPSA determines that the drive has failed.

If you choose not to allocate a hot spare drive to your RAID group, you can still replace a failed drive with any available drive that is not used in any other RAID Group within the VPSA. You can execute this process manually, or automate it via the VPSA REST APIs. Simply identify and select the failed drive, click the Replace button and select the available drive to use for the replacement. For more details see [Replacing a Drive](#).

4.6 Managing RAID Group Sync Speed

The RAID Group **Sync Speed** button allows you to control the rate with which data is synchronized during a RAID rebuild process on both a newly created RAID group and following a drive replacement.

Setting the Sync Speed is a tradeoff between the need to complete the RAID rebuild as quickly as possible in order to return to full redundancy level and the ability to supply good response time and throughput for application IOs. Therefore, the VPSA allows you to control two parameters impacting the sync Speed:

- **"Min Speed During Host I/Os"** – Controls the RAID sync speed when there are Server IOs. You will want to set it low if the Server's IOs are the priority. Set it high if you want to prioritize the RAID rebuild process.
 - Default value: 10 MiB/s
- **"Max Speed w/o Host I/Os"** – Controls the sync speed when there are no Server IOs. You would typically set it to max value (500 MB/s), unless it consumes too much of the VPSA's resources (depending on the Engine type) which impacts the performance of other RAID Groups (which do have active Server IOs).
 - Default value:
 - * HDD: 30 MiB/s
 - * SDD: 100 MiB/s

- Maximum value:
 - * HDD: 100 MiB/s
 - * SDD: 500 MiB/s

You can set and modify Sync Speed at any time, and it can vary between RAID groups. The Sync Speed also applies to Media Scan.

POOLS

5.1 Understanding Storage Pools

Storage Pools are virtual entities that manage storage provisioning from the aggregated capacity of one or more RAID Groups pooled into a single construct with some QoS attributes.

Volumes are thinly provisioned, allocating capacity from the Pool only when needed. The Pool has an underlying block virtualization layer which maps virtual address space to physically allocated Pool space and manages sharing of Pool physical chunks between Volumes, Snapshots and Clones.

Snapshots and Clones consume zero capacity when they are created because they share the same data chunks as the originating Volume. Anytime you actually modify the data in the Volume, or in one of the Clones, the data chunk is copied-on-write (COW) from the source in order to apply the new data write to a new Pool region without affecting the data set of any other objects that share the same data chunk.

The Pool's attributes define the way Volumes, Snapshots and Clones are provisioned.

A single VPSA can have multiple storage pools based on the VPSA type and engine type according to the following specifications:

- **VPSA Storage Array** - 3 Pools
- **VPSA Flash Array**
 - H100 - 1 Pool
 - H200 - 2 Pools
 - H300 - 2 Pools
 - H400 - 2 Pools

5.1.1 Tiers

Scope: VPSA Flash Array

✓ **Note:** From version 20.12, VPSA Flash Array supports a mixture of media types within one storage Pool, as tiers. Tiering optimizes costs, automating data storage placement by tracking segments, keeping high frequency activity on SSD storage and low frequency activity on HDD.

The high tier (also known as tier 0) is the more performant tier of the VPSA, and comprises in-array SSD/NVME drives.

The low tier (also known as tier 1) is the capacity-oriented tier of the VPSA. It can be implemented via in-array SATA/NLSAS drives, or by connectivity to a remote object storage container.

Supported configurations:

- Tier 0 is SSD and Tier 1 (low tier) is HDD.
- Tier 0 is SSD and Tier 1 (low tier) is Remote Object Storage.

Inline data reduction is supported irrespective of the actual data location. Data storage placement in the VPSA Pool is aligned dynamically, based on an internal heat index. Scanning, promotion and demotion of the storage location between tiers is embedded in the VPSA garbage collection cycle.

The VPSA tier manager keeps track of heat scores for the hottest LSA chunks in each tiered Pool.

The heat score is calculated according to:

- Frequency of chunk read-write operations
- Chunk deduplication references

The VPSA attempts to stabilize SSD utilization at a steady state around 80%.

SSD utilization	Promotion and demotion considerations
Below steady state	All data is retained in SSD and no demotions take place.
At stabilization target around steady state	Placement considerations, promotions and demotions.
Above steady state	Demotions are increased and promotions are blocked.

Chunk promotion: Chunks can be promoted to a higher tier:

- By the low tier defragger
- On host reads
- By the tier manager, based on periodic assessing of chunks with highest heat scores

Chunk tier demotion: Chunks can be demoted to a lower tier:

- By the SSD defragger
- By the tier manager, when SSD utilization is at the steady state and higher

5.2 Understanding Pool's Capacity

Scope: VPSA Flash Array

The introduction of data reduction makes the Pool capacity management more complex. Data reduction efficiency depends on the nature of the data, therefore it is harder to predict the drive capacity needed for each workload.

Capacity metrics to consider:

5.2.1 Physical View

Raw Capacity - Sum of all drives capacities in the Pool

Usable Capacity - Total capacity of all RAID Groups in the Pool

✓ **Note:** the system keeps about 0.5% of each RAID Group capacity as its internal spare

Used by Volumes - Capacity used to store the Volumes' data

Used by metadata - Capacity used to store the Pool's metadata

Used by data copies: - Capacity used to store the Snapshots and Clones

Used Capacity - The total size of all data written in the Pool

Used Capacity = Used by Volumes + Used by metadata + Used by data copies

Free Capacity - Available Capacity in the Pool that can be used for new Data and Metadata writes

Free Capacity = Usable Capacity - Used Capacity

%Full = "Used Capacity" / "Usable Capacity"

✓ **Note:** Capacity alerts are based on Free Capacity

5.2.2 Virtual View

Provisioned Capacity - Sum of Pool's Volumes and Clones capacities as seen by the hosts


Allocated Capacity - Pool's allocated address space of all Volumes, Snapshots and Clones

Allocation Limit - Max Capacity of the Pool's address space. Depends on the Pool type.

Free Address Space = Allocation Limit - Allocated Capacity

✓ **Note:** Address Space alerts are based on Free Address Space

Effective Capacity - Amount of data written in the Pool by all Volumes and can be accessed by hosts. Not including space taken by Snapshots



5.2.3 Data Reduction Saving

Thin Provision Ratio = Provisioned Capacity / Effective Capacity

Data Reduction Ratio = Effective Capacity / Used by Volumes

Data Reduction Saving = Effective Capacity - Used by Volumes

Data Reduction Percentage = $1 - (1 / \text{Data Reduction Ratio})$

e.g.

Data reduction ratio 2:1 , Data Reduction Percentage 50%

Data reduction ratio 5:1 , Data Reduction Percentage 80%


Data reduction ratio 20:1 , Data Reduction Percentage 95%

5.2.4 Pool Capacity Monitoring


The VP SA Flash Array **Dashboard** shows the capacity consumption and data reduction saving.

CAPACITY

DATA REDUCTION SAVINGS



	Ratio	Data Savings
Compression	1.002	122.88 MiB
Dedupe	17.460	1.08 TiB
Reduction	17.492	1.08 TiB



The upper bar shows the current capacity provisioned to the hosts by all Pools vs. the effective capacity written by the hosts vs. the physical space needed to store the data.

The lower chart shows trend of time of the physical capacity used and available.

The Pools table on the **Pools** page shows 2 bars per Pool:

Dashboard	POOLS						
RESOURCES	REFRESH	CREATE	DELETE	EXPAND	SHRINK	CREATE VOLUME	CREATE SHARE
Drives	Name	Status	RAID Group(s)	Physical Capacity	Virtual Capacity	Type	Data Reduction Ratio
RAID Groups	Pool1	Normal	Multiple(2)	8.88 TiB Free / 8.95 TiB	198.85 TiB Free / 200 TiB	Balanced	17.5 : 1
Pools							

- The **Physical capacity** bar shows the usable vs. used capacities.
- The **Virtual capacity** bar shows the allocated capacity vs. the allocation limit.

5.3 Viewing the List of Pools

The Pools page displays a table listing the VPSA's Pools in the upper part of the screen.

Both VPSA Storage Array Pools and VPSA Flash Array Pools have the following common columns:

Property	Description
Name	User assigned name. Can be modified anytime.
Status	<ul style="list-style-type: none"> • Normal • Creating • Deleting • Partial/Failed – At least one of the underlying RAID Groups has failed, or the Pool metadata cannot be initialized at Start Of the Day.
Raid Group(s)	RAID Group name, or “Multiple (X)” where X denotes the number of RAID Groups in the Pool.

- Columns specific to VPSA Storage Array Pools:

Property	Description
Capacity	Total available capacity for user data and system metadata.
Type	<ul style="list-style-type: none"> - Transactional Workloads - Repository Storage - Archival Storage
Cached	Yes/No – Indicates whether the Pool utilizes SSD for read/write caching.
Provisioned Capacity	Sum of Pool's Volumes and Clones capacities as seen by the hosts.

- Columns specific to VPSA Flash Array Pools:

Property	Description
Physical Capacity	The usable vs. used capacities in the Pool.
Virtual Capacity	The allocated capacity vs. the allocation limit in the Pool.
Type	<ul style="list-style-type: none"> - IOPS-Optimized - Balanced - Throughput-Optimized
Data Reduction Ratio	Capacity savings by all data reduction techniques. Data Reduction Ratio = Effective Capacity / Used by Volumes

5.3.1 Filtering the List of Pools

To filter the list of Pools displayed in the center pane, you can use a predefined custom tag.

The screenshot shows the 'SERVERS' section of the interface. Below the navigation bar (REFRESH, ADD, CONFIG, DELETE, VOLUMES), there is a 'Filter' section. It includes a 'Tag Name' input field, a 'Tag Value' input field with an 'Optional' dropdown and a trash icon, and an 'Add Filter' dropdown menu currently set to 'Tag'. Below these fields are 'SEARCH' and 'CLEAR' buttons.

1. Expand the **Filter** control.
2. On the right, click the **Add Filter** dropdown, and select **Tag**. The tag input fields appear on the left.
3. The **Tag** filter requires input of the predefined custom **Tag Name**, and optionally, a **Tag Value** to further refine filtering for tags that have specific values for assigned Pools. Wildcards are not accepted. See the [Tags tab](#) section for configuring predefined custom tags.
4. Click **Search** to apply the filter.
5. To remove the filter, click **Clear**, or click the trash icon to the right of the filter. Click **Search** again to refresh the Volumes list.

5.4 Creating and Managing Pools

5.4.1 Creating a Pool

✓ Note: By default when a new VPSA is created, a default Pool is automatically created for each type of drive selected for this VPSA.

If the default Pool does not meet the needs, you can delete it and follow the process described here to create your own Pools.

To create a new Storage Pool press either the Create button on the [Pools](#) page or the Create Pool button on the [RAID Groups](#) page. There are 2 methods to create a Pool:

- Create a Pool from RAID Groups
- Create a Pool from drives, and let the system automatically create the needed RAID Groups.

You can toggle between the two by clicking Use Drive Selection / Use RAID Group Selection at the lower left corner of the dialog.

Create a Pool from RAID Groups

To create a Pool from RAID Groups, click the **Create** button on the Pools page. The **Create Pool** dialog opens. At the lower left, select **Use RAID Group Selection** from the toggle:

Create Pool

Name: *

Pool1

RAID Group(s): *

<input type="checkbox"/>	Name	Protection	Status	Available
<input checked="" type="checkbox"/>	RAID-10-Pool-1-r0	RAID1	Normal	2.73 TB
<input checked="" type="checkbox"/>	RAID-10-Pool-1-r1	RAID1	Normal	2.73 TB
<input type="checkbox"/>	RAID-10-Pool-1-r2	RAID1	Normal	2.73 TB

Capacity (GiB):

5000

Calculate Max

Type:

☐ Transactional Workloads
 ☒ Repository Storage
 ☐ Archival Storage
 ☐ Depot Storage

Cached:

☒

Striped:

☒

Use Drive Selection

Create

Cancel

Select the Pool attributes:

- **Display Name** – You can modify this anytime later.
- **RAID Group(s) selection** – Check the box(es) of one or more RAID Groups from which protected storage capacity will be allocated for this Pool.
- **Capacity** – The Pool's physical capacity shown in GB. By default the capacity is the aggregated capacities of all the selected RAID Groups, but you do not have to allocate full RAID Groups. If you define a capacity smaller than is available in the selected RAID Groups the capacity will be evenly distributed between the RAID Groups.

✓ **Note:** The actual usable capacity of the Pools is a little less than the requested size, as the system reserves some space for the Pool's metadata (typically up to 100GB).

- **Type**
 - VPSA Storage Array supports Transactional, Repository, Archive Pool and Depot Pool types.
 - VPSA Flash Array supports IOPS-Optimized, Balanced Pool and Throughput-Optimized Pool types.

These Pool types use different chunk sizes for the mapping of virtual LBAs to Physical Drive addresses.

The following tables describe the tradeoffs for each type and the recommended use cases:

VPSA Storage Array **Pool types:**

	Transactional Pool	Repository Pool	Archive Pool	Depot Pool
Chunk size	256KB	1MB	2MB	4MB
Pros	<ul style="list-style-type: none"> - Faster COW operation - Space efficiency on random writes to Snapshots 	<ul style="list-style-type: none"> - Smaller metadata size - Sequential workload performance is similar to transactional 	<ul style="list-style-type: none"> - Allows large Pools - Sequential workload performance is the same 	<ul style="list-style-type: none"> - Optimal capacity consolidation for backup and media storage
Cons	<ul style="list-style-type: none"> - Increased metadata size 	<ul style="list-style-type: none"> - Slower COW operation - Less space efficient 	<ul style="list-style-type: none"> - Slower with frequent data modifications - Limited Snapshots frequency (1 hour min) 	<ul style="list-style-type: none"> - Slower for small block transactional workloads (optimal support for workloads with IO block size $\geq 256\text{KB}$) - Limited Snapshot frequency (1 per day with 30 days retention) - Less space-efficient for small files (recommended for file sizes $\geq 1\text{MB}$)
Use case	<ul style="list-style-type: none"> - Transactional workload with Snapshots 	<ul style="list-style-type: none"> - Repository type workload - Large Pools - Many Snapshots to keep 	<ul style="list-style-type: none"> - Relatively static data - Archive type workloads - Very large Pools/Volume ($> 100\text{TB}$) 	<ul style="list-style-type: none"> - Backup repositories - Media repositories - Large files archive
Maximum size	20TiB	100TiB	200TiB	500TiB

VPSA Flash Array **Pool types:**

	IOPS-Optimized Pool	Balanced Pool	Throughput-Optimized Pool
Thin Provision Chunk size	1MB	2MB	4MB
Deduplication Chunk size	16KB	32KB	64KB
Pros	<ul style="list-style-type: none"> - Smaller metadata size - Better deduplication - Lower COW overhead in cases of small block I/O 	<ul style="list-style-type: none"> - Allows large Pools - Better sequential workload throughput compared to IOPS-Optimized Pools - Better compression ratio compared to IOPS-Optimized Pools 	<ul style="list-style-type: none"> - Optimal capacity consolidation for backup and media storage - Better sequential workload throughput compared to other Pool types - Better compression ratio compared to other Pool types
Cons	<ul style="list-style-type: none"> - Increased metadata size 	<ul style="list-style-type: none"> - Higher COW overhead for IOPS < 16KB in comparison to IOPS-Optimized Pools - Lower deduplication efficiency compared to IOPS-Optimized Pools 	<ul style="list-style-type: none"> - Higher COW overhead for IOPS < 32KB in comparison to other Pool types - Higher latency for small block writes < 64KB due to RMW - Less space-efficient for small files (recommended for file sizes \geq 1MB) - Low deduplication efficiency (recommendation is to turn deduplication off)
Use Case	<ul style="list-style-type: none"> - Analytics - Small block IOPS workloads - High IOPS - Database (OLTP) - Deduplication-friendly data 	<ul style="list-style-type: none"> - Pools/Volumes (> 100TB) - Archive type workloads - Workloads with average IO block size of 32KB - General purpose - File system 	<ul style="list-style-type: none"> - Relatively static data - Backup repositories - Media repositories - Large file archives - Archive type workloads - Sequential workloads such as video streaming - Workloads with average IO block size > 128KB
Maximum size	100TiB	200TiB	500TiB

When there are a number of Pools in a given VP SA, there is a limit to the aggregated total size of all Pools.

The following table lists the maximum capacity per Pool type in TB, per VP SA Flash Array engine:

	Engine			
	H100	H200	H300	H400
IOPS-Optimized	60	100	100	100
Balanced	100	160	200	200
Throughput-Optimized	140	220	400	500

✓ **Note:** In most cases, the maximum provisioned capacity is the same as the maximum usable capacity for that engine and Pool type configuration.

An H400 engine with a Balanced Pool supports a maximum provisioned capacity of 250TB when the overall data reduction ratio for the array is 1:1.5 or higher.

- **Cached** – Check this box to use SSD to Cache Server’s reads and writes.
 - All Pools that are marked as “Cached” share the VP SA Cache.
 - Flash cache usually improves the performance of Volumes based on HDD’s Pools. However it depends on the specific workload and the size of the cache vs. the size of the active data set.
 - If the Pool consists of SSD drives this option will be disabled.
- **Striped** – This check box is enabled only when you select two or more RAID Groups. Striping over RAID-1 creates a RAID-10 configuration. Use striping to improve performance of random workloads, since the IOs will be distributed and all drives will share the workload.

VP SA Flash Array Pools are always striped, and the **Striped** check box is hidden.

VP SA Flash Array configurations

Scope: VP SA Flash Array

- **Additional Storage Class:** Adding a storage class defines a low tier for the Pool as HDD or remote Object Storage. Adding a storage class opens the **SSD Cool Off** configuration option.

Create Pool

Name: *

Pool2

Type:

☐ IOPs-Optimized
 ☒ Balanced
 ☐ Throughput-Optimized

SSD Cool Off (Hours):

12

SSD Storage Class

<input type="checkbox"/>	Name	Type	Capacity	Status	Storage Node	Zone
<input type="checkbox"/>	drive-020	SSD	3.49 TiB	Available	sn-01	zone_0
<input type="checkbox"/>	drive-021	SSD	3.49 TiB	Available	sn-02	zone_0

Use Raid Group Selection

Additional Storage Class: *

☐ None
 ☒ HDD
 ☐ OBS

HDD Storage Class

<input type="checkbox"/>	Name	Type	Capacity	Status	Storage Node	Zone
<input type="checkbox"/>	drive-018	SAS	5.46 TiB	Available	sn-02	zone_0
<input type="checkbox"/>	drive-019	SAS	5.46 TiB	Available	sn-01	zone_0

Use Raid Group Selection

Configured Pool Useable Capacity: 0 GiB

Create

Cancel

- **SSD Cool Off:** Set a goal for data retention in SSD with a value of 0 to 720 hours (30 days). The default is 0 (disabled). This hints to the system that within the cool-off period there will be a repeated access to a data chunk in the Pool. When SSD utilization is around the steady state, the tiering manager references the cool-off period definition in its decision to determine tier placement for the data chunk.

Create a Pool from Drives

To create a Pool from Drives, click the **Create** button on the Pools page. The **Create Pool** dialog opens. At the lower left, select **Use Drive Selection** from the toggle:

Create Pool

Name: *

Pool2

Drives: *

Select even number of drives of the same type

<input checked="" type="checkbox"/>	Name	Type	Capacity	Status	Storage Node	Zone
<input checked="" type="checkbox"/>	drive-000	SATA	2.73 TB	Available	qa9-sn2	zone_0
<input checked="" type="checkbox"/>	drive-001	SATA	2.73 TB	Available	qa9-sn1	zone_0
<input checked="" type="checkbox"/>	drive-002	SATA	2.73 TB	Available	qa9-sn2	zone_0
<input checked="" type="checkbox"/>	drive-003	SATA	2.73 TB	Available	qa9-sn1	zone_0
<input checked="" type="checkbox"/>	drive-004	SATA	2.73 TB	Available	qa9-sn2	zone_0
<input checked="" type="checkbox"/>	drive-005	SATA	2.73 TB	Available	qa9-sn1	zone_0

Capacity:

8.19 TB

Type:

☐ Transactional Workloads
 ☐ Repository Storage
 ☒ Archival Storage
 ☐ Depot Storage

Cached:

☐

Use RAID Group Selection

Create

Cancel

The parameters are the same as above. Check the boxes of drives that will be allocated for this Pool.

5.4.2 Expanding Pool Capacity

To Expand the Pool press the Expand button on the [Pools](#) page.

You can use capacity from any RAID Group to expand a Pool.



Warning: If the RAID Group from which the new capacity is added doesn't match the protection type or drive type of the existing capacity, a warning message displays, asking you to confirm the mismatch.

Keep in mind that continuing with the mismatched types may impact the Pool performance and protection QoS.

VPSA Flash Array configuration

Scope: VPSA Flash Array

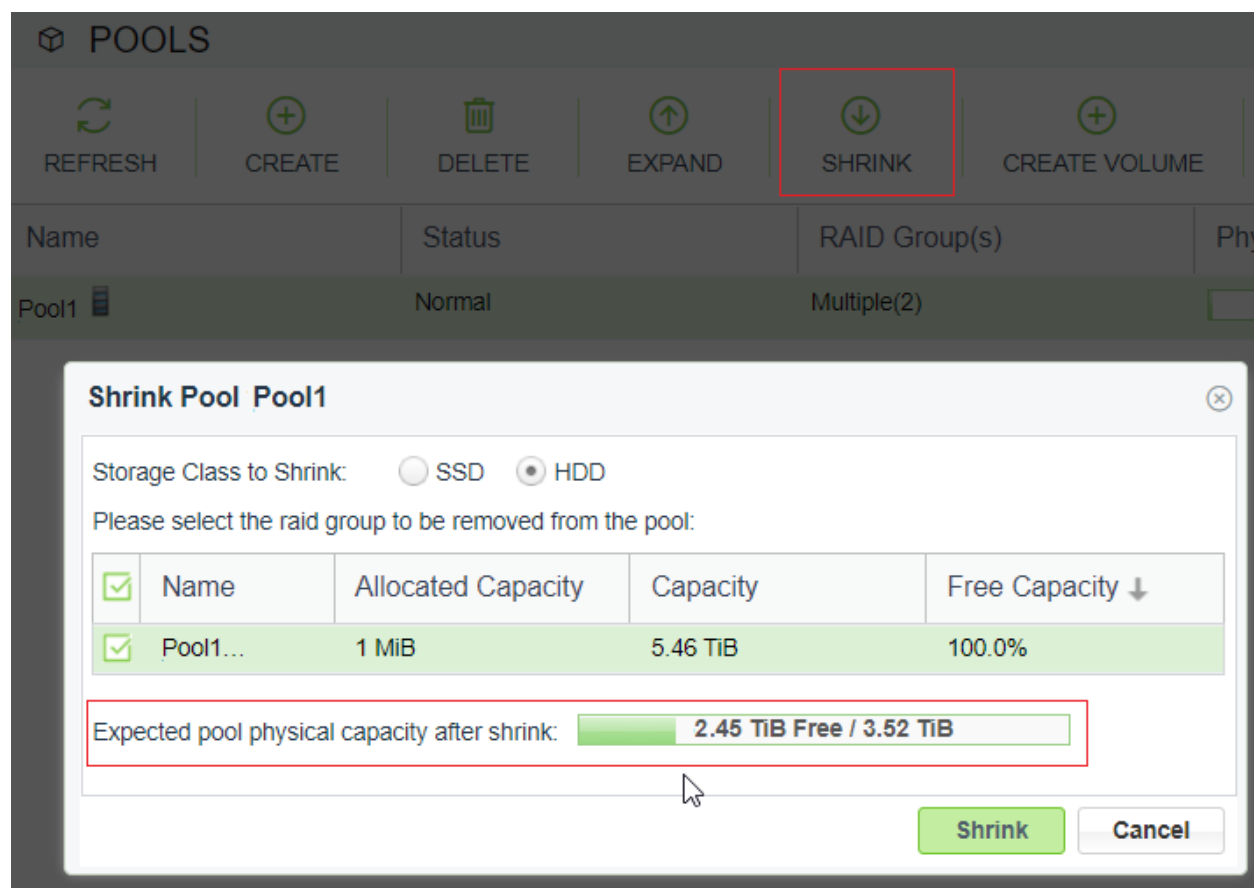
Expand in Storage Class: Choose SSD or HDD to list the storage class resource details and availability for expansion.

5.4.3 Shrinking Pool Capacity

Scope: VPSA Flash Array

✓ Note: Pool shrink is only supported in VPSA Flash Array.

If the Pool capacity is not fully used you can shrink it's size by removing one RAID Group at a time from the Pool. The VPSA will evacuate the selected RAID Group and will return the RAID Group to the VPSA for reuse, or for the RAID Group to be deleted and the drives removed from the VPSA. To Shrink the Pool press the Shrink button on the [Pools](#) page.

**Storage Class to Shrink:** Choose SSD or HDD to list the storage class RAID Group details and availability for shrinkage.

Select the RAID Group to remove from the Pool. Check the physical size expected after the shrinking operation is completed, and press Shrink. The operation might take a while, depending on the amount of data to be copied to other drives. The system will generate an Event once done.

5.4.4 Caching

Scope: VPSA Storage Array

It is possible to enable Caching on non-cached Pools.

One use case for leveraging this capability is to enable caching only after the initial copy of the data into the VPSA. The initial copy typically generates a sequential write IO workload, where non-cached Pools are most efficient. Once the initial copy is completed enable caching on the Pool if you expect a more random type of IO workload.

5.4.5 Disabling SSD cache on a Pool

Scope: VPSA Storage Array

By default every Pool is cached by the VPSA's SSD cache, but it is also possible to disable caching on cached Pools which will remove this feature. The Enable Cache/Disable Cache buttons toggle depending on the current caching state of the Pool.

5.4.6 Creating a Volume

A Volume can be created from the **Create Volume** button in the top menu in the Pools view or in the Volumes view. Both flows are identical.

See [Creating and Deleting a Volume](#) on the **Managing Volumes, Snapshots and Clones** page.

5.4.7 Creating a share

A NAS share can be created from the **Create Share** button in the top menu in the Pools view or in the Volumes view. Both flows are identical.

See [Creating a NAS share](#) on the **Managing Volumes, Snapshots and Clones** page.

5.4.8 Adding a tier

Scope: VPSA Flash Array

It is possible to add a low tier to an existing VPSA Flash Array Pool.

1. In the **Pools** view, search for the Pool in the Pool table, and mark it by clicking its row.
2. Click **Create** in the top menu. The **Create Pool** dialog opens.
 1. Select the high tier from the list in the **SSD Storage Class** table.
 2. Select the **Additional Storage Class**: Either **HDD** or **OBS** (remote Object Storage).
Depending on the selection, either the **HDD Storage Class** list or the **OBS Storage Class** list is displayed.
The **SSD Cool Off** field is also displayed.
 3. Select the low tier from the **HDD Storage Class** or **OBS Storage Class** list.
 4. Set the number of hours for **SSD Cool Off**, between 0 and 720 (30 days).

During the cool-off period there is repeated access to data chunks in the Pool. The tiering manager references the cool-off period definition to determine tier placement for data chunks.

5. Click **Create**.

The tier's properties will be viewable in the [Tiers tab](#).

5.5 Viewing Pool properties

The Pool's details are shown in the following South Panel tabs:

5.5.1 Properties tab

VP SA Storage Array Pool properties

Scope: Scope: VP SA Storage Array

Each VP SA Storage Array Pool has the following properties:

Property	Description
ID	An internally assigned unique ID.
Name	User assigned name. Can be modified anytime.
Status	<ul style="list-style-type: none"> • Normal • Creating • Deleting • Partial/Failed – At least one of the underlying RAID Groups has failed, or the Pool meta-data cannot be initialized at Start Of the Day.
Comment	User free text comment. Can be used for labels, reminders or any other purpose.
Type	<ul style="list-style-type: none"> • Transactional Workloads • Repository Storage • Archival Storage
Mode	<ul style="list-style-type: none"> • Simple – There are one or more concatenated RAID Groups. • Stripe – There are two or more striped RAID Groups. • Mixed – There are two or more concatenated and striped RAID Groups.
Cached	Yes/No – Indicates whether the Pool utilizes SSD for read/write caching.
Cache COW Writes	Yes/No – Indicates whether flash cache is used for internal Snapshots Copy-On-Write Operations. Enabled by default. Disabled only on rare cases where frequent Snapshots cause extreme load of metadata operations. Consult Zadara support.
Raid Group(s)	RAID Group name, or “Multiple (X)” where X denotes the number of RAID Groups in the Pool.
Stripe Size	Applicable only for Pools of Striped mode (i.e. when data is striped between 2 or more RAID Groups). The Stripe size is always 64KB.
Created	Date & time when the object was created.
Modified	Date & time when the object was last modified.
Capacity	Total available capacity for user data & system metadata.
Available Capacity	Available (free) capacity to be used for User data. VPSA reserves 2% of the total Pool capacity for system metadata. If the VPSA needs more capacity for the metadata (very rare scenario), it will be consumed from the available capacity.
Metadata Capacity	Capacity used by the metadata that is required for managing pool allocation space.
Capacity State	<ul style="list-style-type: none"> • Normal • Alert • Protected • Emergency <p>See Managing Pool Capacity Alerts for more details.</p>

VPSA Flash Array Pool properties

Scope: VPSA Flash Array

Each VPSA Flash Array Pool has the following properties:

Property	Description
	General
ID	An internally assigned unique ID.
Name	User assigned name. Can be modified anytime.

continues on next page

Table 1 – continued from previous page

Property	Description
Status	<ul style="list-style-type: none"> • Normal • Creating • Deleting • Partial/Failed – At least one of the underlying RAID Groups has failed, or the Pool meta-data cannot be initialized at Start Of the Day.
Comment	User free text comment. Can be used for labels, reminders or any other purpose.
Type	<ul style="list-style-type: none"> • IOPS-Optimized • Balanced • Throughput-Optimized
Raid Group(s)	RAID Group name, or “Multiple (X)” where X denotes the number of RAID Groups in the Pool.
SSD Cool Off	Period (hours) of repeated access to a data chunk in the Pool. Used for automatically determining tier placement for data chunks.
Created	Date & time when the object was created.
Modified	Date & time when the object was last modified.
	Physical Capacity
Usable Capacity	Total capacity of all RAID Groups in the Pool.
Used Capacity	The total size of all data written in the Pool Used Capacity = Used by Volumes + Used by metadata + Used by data copies
Used by Volumes	Capacity used to store the Volumes data.
Used by Data Copies	Capacity used to store Snapshots and Clones.
Used by Metadata	Capacity used to store the Pool’s metadata.
Currently Inactive	Capacity of the Pool’s data chunks that are currently not being accessed.
Free Capacity	Available Capacity in the Pool that can be used for new Data and Metadata writes.
Physical Capacity State	<ul style="list-style-type: none"> • Normal • Alert • Protected • Emergency <p>See Managing Pool Capacity Alerts for more details.</p>
	Virtual Capacity
Provisioned Capacity	Sum of Pool’s Volumes and Clones capacities as seen by the hosts.
Allocated Capacity	Pool’s allocated address space of all Volumes, Snapshots and Clones.
Pattern Capacity	Capacity savings by data blocks with predefined patterns, such as all-zeroes. Blocks that match the patterns are deduped.
Effective Capacity	Amount of data written in the Pool by all Volumes and can be accessed by hosts. Not including capacity taken by Snapshots.
Block-Virt Metadata Capacity	Capacity used by the metadata that is required for managing pool allocation space.
Free Virtual Capacity State	<ul style="list-style-type: none"> • Normal • Alert • Protected • Emergency <p>See Managing Pool Capacity Alerts for more details.</p>
	Capacity Savings

continues on next page

Table 1 – continued from previous page

Property	Description
Data Reduction Ratio	Capacity savings by all data reduction techniques. Data Reduction Ratio = Effective Capacity / Capacity used by Volumes
Dedupe Ratio	Capacity savings by deduplication.
Compression Ratio	Capacity savings by compression.
Thin Provision Ratio	Capacity savings by thin provisioning technique. Thin Provision Ratio = Provisioned Capacity / Effective Capacity

5.5.2 RAID Groups tab

Scope: VPSA Storage Array

The **RAID Groups** tab has 2 display styles to present the RAID Groups allocated to the selected Pool:

- **RAID Groups View**

Click the **RAID Groups View** button to list the RAID Groups allocated to the selected Pool.

A table lists the following information for each RAID Group:

- Name
- Protection (RAID-1)
- Status
- Contributed Capacity

Details for RAID-10-Pool-1			
Properties	RAID Groups	Volumes	Dest. Volumes
		Recycle Bin	Logs
		Metering	Capacity Alerts
		Performance Alerts	Tags
Name	Protection	Status	Contributed Capacity
RAID-10-Pool-1-r0	RAID1	Normal	5.46 TiB
RAID-10-Pool-1-r1	RAID1	Normal	5.46 TiB

Raid Groups View Segments View

- **Segments View**

Click the **Segments View** button to display the structure of a Pool made of concatenated or striped segments.

Click the **Expand** or **Collapse** controls to display the detail or summary levels.

Details for RAID-10-Pool-1			
Properties	RAID Groups	Volumes	Dest. Volumes
		Recycle Bin	Logs
		Metering	Capacity Alerts
		Performance Alerts	Tags
EXPAND ALL COLLAPSE ALL			
11172GB(concatenation)			
<ul style="list-style-type: none"> RaidGroup-5 (RAID-10-Pool-1-r0) 5586GiB RaidGroup-5 (RAID-10-Pool-1-r1) 5586GiB 			
Raid Groups View Segments View			

5.5.3 Tiers tab

Scope: VPSA Flash Array

The **Tiers** tab displays details of the tier types allocated to the selected Pool.

POOLS

REFRESH
 CREATE
 DELETE
 EXPAND
 SHRINK
 CREATE VOLUME
 CREATE SHARE

Name	Status	RAID Group(s)	Physical Capacity	Virtual Capacity	Type	Data Reduction Ratio
Pool1	Normal	Multiple(2)	8.88 TiB Free / 8.95 TiB	198.85 TiB Free / 200 TiB	Balanced	17.5 : 1

<< < PAGE 1 OF 1 > >>
Displaying 1 - 1 of 1

Details for Pool1

Properties **Tiers** Volumes Dest. Volumes Recycle Bin Logs Metering Capacity Alerts Performance Alerts

Type	Members	Capacity	Status	Utilization
SSD	RaidGroup-6	3.47 TiB 3.49 TiB Usable 68.15 GiB Used By Volumes 2.23 GiB Currently Inactive 3.42 TiB Free	Normal	 1.91% Used By Volumes / 0.06% Currently Inactive / 98.03% Free
Capacity	RaidGroup-7	5.43 TiB 5.46 TiB Usable 0 B Used By Volumes 0 B Currently Inactive 5.46 TiB Free	Normal	 0.00% Used By Volumes / 0.00% Currently Inactive / 100.00% Free

Each tier type includes the following information:

- **Type** of tier.
- **Members:** RAID Groups or Drives and their capacities, that are members of each tier.
- **Capacity:** Measurements in units for Usable capacity, Used By Volumes, Currently Inactive, and Free.
- **Status:** Operational status of the tier.
- **Utilization:** Progress bar and measurements in percentages: Used By Volumes, Currently Inactive, and Free.

5.5.4 Volumes and Dest Volumes tabs

These two tabs display the provisioned Volumes and the Provisioned Remote Mirroring Destination Volumes. Please note that the Dest Volumes are not displayed in the main **Volumes** page, since most operations are not applicable to them. Displaying the list of the Dest Volumes in the Pools South Panel provides a complete picture of the Objects that consume capacity from the Pool.

The **Volumes** tab displays the following information:

- Name
- Capacity (virtual, not provisioned)
- Status
- Data Type (Block or File-System)

The **Dest Volumes** tab displays the following information:

- Name
- Capacity (virtual, not provisioned)
- Data Type (Block or File-System)
- Mapped Capacity
- Data Copies Capacity

5.5.5 Recycle Bin tab

By default when you delete a Volume it moves to a Pool's Recycle Bin for 7 days until it is permanently deleted. From the Recycle Bin, an administrator can purge (permanently delete) or restore a Volume.

A Volume in the **Recycle Bin** tab has the same details that were displayed for it in the **Volumes** tab, before deletion:

- Name
- Capacity (virtual, not provisioned)
- Status
- Data Type (Block or File-System)

5.5.6 Logs tab

The **Logs** tab displays all event logs associated with this Pool:

- Type of log entry, for example: Information, Warning, or Error.
- Title - details of the logged event.
- Time - the date and time of the logged event.

5.5.7 Metering tab

The Metering Charts provide live metering of the IO workload associated with the selected Pool.

The charts display the metering data as it was captured in the past 20 "intervals". An interval length can be set to one of the following: 1 Second, 10 seconds, 1 Minute, 10 Minutes, or 1 Hour. The Auto button lets you see continuously-updating live metering info (refreshed every 3 seconds).

Pool Metering includes the following charts:

Chart	Description
IOPS	The number of read and write SCSI commands issued to the Pool, per second.
Bandwidth (MB/s)	Total throughput (in MB) of read and write SCSI commands issued to the Pool, per second.
IO Time (ms)	Average response time of all read and write SCSI commands issued to the Pool, per selected interval .

5.5.8 Capacity Alerts tab

The **Capacity Alerts** tab lists the configurable attributes of the Pool Protection Mechanism.

- VPSA Storage Array capacity alerts:

VPSA STORAGE ARRAY

RESOURCES

Drives

RAID Groups

Pools

Volumes

POOLS

REFRESH

CREATE

DELETE

EXPAND

ENABLE CACHE

CREATE VOLUME

CREATE SHARE

Filter

Name	Status	RAID Group(s)	Capacity	Type
pool1	Normal	Multiple(3)	38.09 TiB Free / 38.09 TiB	Repository Storage

<< < PAGE 1 OF 1 > >> ↺

Details for pool1

PropertiesRAID GroupsVolumesDest. VolumesRecycle BinLogsMeteringCapacity AlertsPerformance AlertsTags

Your VPSA™ will alert you through a support ticket when your Pool reaches specified capacity thresholds.

Alert Threshold
Alert me when it is estimated that the Pool will be at full capacity in:
720 minutes

Protection Threshold
Do not allow new Volumes, Shares, or Snapshots to be created when it is estimated that the Pool will be at full capacity in:
360 minutes

Emergency Threshold
Delete snapshots, starting from the oldest, when there is less than the following capacity left in the Pool:
50 GiB

Alert Interval
Time estimations above are based on capacity usage during the last:
60 minutes

- VPSA Flash Array capacity alerts, and where they differ:

VPSA FLASH ARRAY

RESOURCES

Drives

RAID Groups

Pools

Volumes

POOLS

REFRESH

CREATE

DELETE

EXPAND

SHRINK

CREATE VOLUME

CREATE SHARE

Filter

Name	Status	RAID Group(s)	Physical Capacity	Virtual Capac
Pool-11	Normal	Multiple(5)	14.97 TiB Free / 44.88 TiB	127.54 TiB

<< < PAGE 1 OF 1 > >> ↺

Details for Pool-11

PropertiesTiersVolumesDest. VolumesRecycle BinLogsMeteringCapacity AlertsPerformance AlertsTags

Your VPSA™ will alert you through a support ticket when your Pool reaches specified capacity thresholds.

Physical

Physical Capacity Alert Threshold
Alert me when:
It is estimated that the Pool will be at full capacity in: 600 minutes
or
The pool is: 90 % full

Physical Capacity Protection Threshold
Do not allow new Volumes, Shares, or Snapshots to be created when:
It is estimated that the Pool will be at full capacity in: 180 minutes
or
The pool is: 95 % full

Physical Capacity Emergency Threshold
Delete snapshots, starting from the oldest, when:
There is less than the following capacity left in the Pool: 50 GiB
or
The pool is: 100 % full

Physical Capacity Alert Interval
Time estimations above are based on capacity usage during the last:
60 minutes

Virtual

Allocated Capacity Alert Threshold
Alert me when it is estimated that the Pool Allocated Capacity will be at full capacity in:
360 minutes

Allocated Capacity Protection Threshold
Do not allow new Volumes, Shares, or Snapshots to be created when it is estimated that the Pool Allocated Capacity will be at full capacity in:
60 minutes

Allocated Capacity Emergency Threshold
Delete snapshots, starting from the oldest, when there is less than the following Allocated capacity left in the Pool:
5 GiB

Allocated Capacity Alert Interval
Time estimations above are based on capacity usage during the last:
60 minutes

✓ **Note:** Note the differences between the capacity alert options available for VPSA Storage Array and VPSA Flash Array

See [Managing Pool Capacity Alerts](#) for detailed information on capacity alert comparisons and configuration options.

5.5.9 Performance Alerts tab

The **Performance Alerts** tab lists the Pool's configurations for sending alerts when performance drops below expectations. See [Managing Pool Performance Alerts](#) for more details.

5.5.10 Tags tab

Predefined custom tags can be configured in the **Tags** tab. An example use case for tags is [Filtering the List of Pools](#) in the center pane.

A tag is identified by its **Tag Name** and has a **Tag Value** associated with it. A tag can be defined only once for a Pool. However, the same **Tag Name** can be defined with a different **Tag Value** for other Pools.

- **Create:** To create a new tag for a Pool, in the Pool's **Tags** tab click **Create**, and enter the **Tag Name** and **Tag Value**. The tag is added to the list of tags in the **Tags** tab.
- **Edit:** To change the **Tag Value** of an existing tag, click on that tag in the tags list to mark it, and then click **Edit**. The **Edit Tag** dialog box opens, allowing overwriting of the **Tag Value**.

✓ **Note:** Only the **Tag Value** can be edited. A tag cannot be renamed. It must be deleted, and a tag with the new name configured in its place.

- **Delete:** To delete a tag, click on that tag row in the tags list to mark it, and then click **Delete**. A confirmation dialog box opens.
- **Refresh:** Displays the updated tags list.

5.6 Managing Pool Capacity Alerts

The VPSA's efficient and sophisticated storage provisioning infrastructure maximizes storage utilization, while providing key enterprise-grade data management functions. As a result, you can quite easily over-provision a Pool with Volumes, Snapshots and Clones, hence requiring a Pool Protection Mechanism to alert and protect when free Pool space is low.

The VPSA Pool Protection Mechanism is either time-based or capacity consumption based. The goal is to provide you sufficient time to fix the low free space situation by either deleting unused Volumes/Snapshots/Clones or by expanding the Pool's available capacity (a very simple and quick process due to the elasticity of the VPSA and the Zadara Storage Cloud).

The VPSA measures the rate at which the Pool's free space is consumed and calculates the estimated time left before running out of free space.

User-configurable parameters impact alerts and operations that are performed as part of the Pool Protection mechanism. Capacity alerts can be viewed and configured in the Pool's south pane [Capacity Alerts tab](#).

The following table lists the default values of the configurable capacity alert parameters for the VPSA Storage Array compared with those of the VPSA Flash Array:

Type	Alert Parameter Name	Alert Measurement	VPSA Storage Array	VPSA Flash Array
Physical	Capacity Alert Threshold	Estimated remaining time to full capacity	360 minutes	600 minutes
		Percentage of used capacity	▢	90% full
Physical	Capacity Protection Threshold	Estimated remaining time to full capacity	60 minutes	180 minutes
		Percentage of used capacity	▢	95% full
Physical	Capacity Emergency Threshold	Free capacity less than	50 GiB	50 GiB
		Percentage of used capacity	▢	99% full
Physical	Capacity Alert Interval	Timeframe to calculate consumption rate	60 minutes	60 minutes
Virtual	Allocated Capacity Alert Threshold	Estimated remaining time to full address space	▢	360 minutes
Virtual	Allocated Capacity Protection Threshold	Estimated remaining time to full address space	▢	60 minutes
Virtual	Allocated Capacity Emergency Threshold	Free address space less than	▢	5 GiB
Virtual	Allocated Capacity Alert Interval	Timeframe to calculate consumption rate	▢	60 minutes

5.6.1 VPSA Storage Array and VPSA Flash Array physical alerts

Scope: VPSA Storage Array and VPSA Flash Array

Labels indicate where there are differences between the default values for VPSA Storage Array and VPSA Flash Array configurations of the capacity alert parameters, and for parameters available only in VPSA Flash Array implementations.

- **Physical Capacity Alert Threshold**

“Alert me when it is estimated that the Pool will be at full physical capacity in X Minutes.”

The estimated time (in minutes) before running out of free space. When triggered, an online support ticket is submitted and an email is sent to the VPSA user. When crossing this threshold the Free Capacity State changes to “Alert” and the available capacity will be shown in Yellow. A secondary “reminder” ticket and an email will be generated when only half of this threshold’s estimated time is left.

- Estimated time to reach full capacity.
 - * Default:
 - 360 minutes (6 hours) VPSA Storage Array
 - 600 minutes (10 hours) VPSA Flash Array
 - * Minimum: 1 minute (0 means disable this time-based alert)
- Used capacity percentage (for VPSA Flash Array only):

In addition to the time calculation, a VPSA Flash Array can also track the Pool’s used capacity percentage, and issues an alert on reaching either the time or percentage threshold, whichever is first.

- * Default Percentage: 90% full

- * Minimum: 1 % (0 means disable this percentage-based alert)

• Physical Capacity Protection Threshold

“Do not allow new Volumes, Shares, or Snapshots to be created when it is estimated that the Pool will be at full physical capacity in X Minutes.”

The estimated time (in minutes) before running out of free space. When triggered, the VPSA starts blocking the creation of new Volumes, Snapshots and Clones in that Pool. A support ticket and email are also generated. When crossing this threshold, the Free Capacity State changes to “Protect” and the available capacity will be shown in Red.

- Estimated time to reach full capacity.

- * Default:

- 60 minutes (1 hour) VPSA Storage Array
- 180 minutes (3 hours) VPSA Flash Array

- * Minimum: 1 minute (0 means disable this time-based alert)

- Used capacity percentage (for VPSA Flash Array only):

In addition to the time calculation, a VPSA Flash Array can also track the Pool’s used capacity percentage, and issues an alert on reaching either the time or percentage threshold, whichever is first.

- * Default Percentage: 95% full

- * Minimum: 1 % (0 means disable this percentage-based alert)

• Physical Capacity Emergency Threshold

“Delete Snapshots, starting from the oldest, when there is less than the following capacity left in the Pool.”

When the Pool’s free capacity drops below this fixed threshold (in GiB), the VPSA starts freeing Pool capacity by deleting older Snapshots. The VPSA will delete one Snapshot at a time, starting with the oldest Snapshot, until it exceeds the Emergency threshold (i.e. when free capacity is greater than the threshold). A support ticket and email are also generated. When this threshold is crossed the Free Capacity State changes to “Emergency” and the available capacity will be shown in Red.

- Free capacity less than:

- * Default: 50 GiB

- * Minimum: 1 GiB

- Used capacity percentage (for VPSA Flash Array only):

In addition to the free capacity amount, a VPSA Flash Array can also track the Pool’s used capacity percentage. On reaching either the capacity amount or percentage threshold, whichever is first, the VPSA issues an alert and proceeds with snapshot deletion until free capacity drops below the Emergency threshold.

- * Default Percentage: 99% full

- * Minimum: 1 % (0 means disable this percentage-based alert)

• Physical Capacity Alert Interval

“Time estimations above are based on capacity usage during the last X minutes.”

The size of the window (in minutes) that is used to calculate the rate at which free space is consumed. The smaller the window is, the more this rate is impacted by intermediate changes in capacity allocations, which can result from changes in workload characteristics and/or the creation/deletion of new Snapshots and Clones.

- Default: 60 minutes (1 hour)

- Minimum: 1 minute

5.6.2 VPSA Flash Array allocated capacity alerts

Scope: VPSA Flash Array

In addition to the physical capacity alerts, the VPSA Flash Array also provides configurable thresholds to alert in cases where the Pool allocation (virtual address space) is near capacity.

Free Address Space = Allocation Limit – Allocated Capacity

The following user-configurable parameters impact alerts and operations that are performed as part of the VPSA Flash Array Pool Protection mechanism:

- **Allocated Capacity Alert Threshold**

“Alert me when it is estimated that the Pool’s address space will be at full capacity in X Minutes.”

The estimated time (in minutes) before running out of free address space. When triggered an online support ticket is submitted and an email is sent to the VPSA user. When crossing this threshold the Allocated Capacity Alert Mode changes to “Alert” and the available address space will be shown in Yellow. A secondary “reminder” ticket and an email will be generated when only half of this threshold’s estimated time is left.

- Default: 360 minutes (6 hours)
- Minimum: 1 minute (0 means disable this alert)

- **Allocated Capacity Protection Threshold**

“Do not allow new Volumes, Shares, or Snapshots to be created when it is estimated that the Pool’s address space will be at full capacity in X Minutes.”

The estimated time (in minutes) before running out of free address space. When triggered the VPSA starts blocking the creation of new Volumes, Snapshots and Clones in that Pool. A support ticket and email are also generated. When crossing this threshold, the Allocated Capacity Alert Mode changes to “Protect” and the available address space will be shown in Red.

- Default: 60 minutes (1 hour)
- Minimum: 1 minute (0 means disable this alert)

- **Allocated Capacity Emergency Threshold**

“Delete Snapshots, starting from the oldest, when there is less than the following free address space left in the Pool.”

When the Pool’s free address space drops below this fixed threshold (in GiB), the VPSA starts freeing Pool capacity by deleting older Snapshots. The VPSA will delete one Snapshot at a time, starting with the oldest Snapshot, until it exceeds the Emergency threshold (i.e. when free address space is greater than the threshold). A support ticket and email are also generated. When this threshold is crossed the Free Capacity State changes to “Emergency” and the available address space will be shown in Red.

- Default: 5 GiB
- Minimum: 1 GiB

- **Allocated Capacity Alert Interval**

“Calculate the estimated time until the Pool’s address space is full based on new capacity usage in the previous X minutes.”

The size of the window (in minutes) that is used to calculate the rate at which free address space is consumed. The smaller the window is the more this rate is impacted by intermediate changes in capacity allocations, which can result from changes in workload characteristics and/or the creation/deletion of new Snapshots and Clones.

- Default: 60 minutes (1 hours)
- Minimum: 1 minute

5.7 Managing Pool Performance Alerts

A VPSA administrator has the option to set Pool Performance Alerts in addition to the default Pool Capacity Alerts. Performance Alerts are available for:

Read IOPS Limit – Creates an alert when the average read IOPS, during the past minute, for a Pool exceeds a user-specified threshold.

Read Throughput Limit - Creates an alert when, during the past minute, the average read MB/s for a Pool exceeds a user-specified threshold.

Read Latency Limit – Creates an alert when, during the past minute, the average read latency for a Pool exceeds a user-specified threshold.

Write IOPS Limit – Creates an alert when, during the past minute, the average write IOPS for a Pool exceeds a user-specified threshold.

Write Throughput Limit - Creates an alert when, during the past minute, the average write MB/s for a Pool exceeds a user-specified threshold.

Write Latency Limit – Creates an alert when, during the past minute, the average write latency for a Pool exceeds a user-specified threshold.

5.8 Deleting a Pool

The VPSA administrator can delete a specific Pool if it not needed by clicking the **Delete** option under the Pool top option menu.

Due to the sensitivity of the operation the system will block a delete request where the underlying entities (Volumes and Snapshots) still exist.

The VPSA has a built-in Recycle-Bin mechanism that protects against human error (enabled by default). The Pool cannot be deleted where Volumes are stored in the Recycle Bin (Pool level). If you wish to delete the Pool, ensure that all Volumes are Purged from its Recycle Bin.

VOLUMES

VPSA virtual Volumes are thinly provisioned utilizing an efficient and sophisticated block-level mapping layer. The Volume's virtual address space is carved into virtual contiguous blocks (a.k.a. "Chunks"). When you create a Volume it consumes zero Pool capacity. Pool capacity is provisioned to volumes on demand. Only at the first write to each chunk the physical space is allocated from the Pool capacity to the Volume, and mapping update of the virtual-to-physical addresses.

The Volume's virtual Capacity is not limited to the available Pool capacity.

Snapshots are read-only representations of the Volume's data at a given point-in-time. They are thinly provisioned and share the same data chunks with their Volume as much as possible until you actually modify the chunk's data. This triggers a Redirect On Write (ROW) operation where a new chunk is provisioned and the modified data is written there.

Cloned Volumes are Volumes created by cloning another Volume's data set at a specified point-in-time Snapshot. Volumes and their Clones share unmodified Pool Chunks. A COW is triggered whenever you modify a chunk in the Volume or in the Clone.

Volumes can be Block Volumes (exposed via an iSCSI or Fibre Channel protocols) or NAS Shares (exposed via NFS or SMB protocols).

The following protocols are supported in the VPSA Storage Array and VPSA Flash Array:

Table 1: Supported volume protocols

Volume Type	Version
Block	iSCSI, iSER, FC
NAS (SMB/CIFS)	2.x, 3.x
NAS (NFS)	3, 4.0, 4.1, 4.2

Important: As of VPSA version 20.12, SMB version 1 is no longer supported

6.1 Creating and Deleting a Volume

To Create a Volume go to the [Volumes](#) Page and press the Create button. Select whether you wish to create a Block Volume or a NAS Share.

6.1.1 Creating a Block Volume

Define the following Volume attributes in the Create Block Volume dialog:

- **Name** – the Volume’s display name. This must be unique, and can be modified throughout the Volume’s lifetime.

✓ **Note:** Objects names can be up to 128 chars long and can contain letters and digits, dashes “-” and underscores “_”

- **Capacity** – Virtual Capacity of the Volume in GB. All Volumes are thinly provisioned. No actual capacity is allocated when the Volume is created, so the aggregated Virtual capacity of the volumes is not bounded by the Pool capacity. It is possible to over-provision a Pool, but you need to manage and monitor this it carefully, using a Pool Protection Mechanism (see [Managing Pool Capacity Alerts](#) for more details).

Warning: ReFS doesn’t support trim/unmap operation for non-Microsoft Storage Spaces storage array. ReFS doesn’t support thin-provisioned volumes. It is highly recommended to avoid VP SA’s storage pool over-provisioning in case ReFS filesystem is planned to be used as deleted data will not be reflected in the VP SA’s pool capacity

- **Pool** – Select the Pool that is most appropriate for your Volume’s QoS requirements (based on available capacity, caching, RAID protection, drive types, etc.).
- **Encrypted** – Select this checkbox if you wish to encrypt the volume’s data on the drives. Please note that you must first define an encryption password via the [Controllers](#) Page. For more details about Volume encryption see [Managing Encrypted Volumes](#).
- VP SA Flash Array **Compress** and **Dedupe** options:
 - **Compress** – Check the checkbox if you want the new volume to be compressed.
 - **Dedupe** – Check the checkbox if you want the new volume to be deduped.

✓ **Note:** The ability to apply the Compress and Dedupe options is dependent on the Data Reduction feature being enabled at the VP SA level in the Provisioning portal. See [Enabling or Disabling Data Reduction Bundle](#).

- **Attach Default Snapshot Policies** – Refer to [Managing Snapshot Policies](#) for a detailed explanation regarding snapshot policies. You can apply and remove snapshot policies from a Volume at any time.
- **Performance Capping** - To limit the volume’s maximum input or output operations or throughput per second, enter values for:

Performance Capping Parameter	Description
Read IOPS	The maximum number of read operations per second.
Write IOPS	The maximum number of write operations per second.
Read MBPS	The maximum throughput of data in Megabytes per second for read operations.
Write MBPS	The maximum throughput of data in Megabytes per second for write operations.

6.1.2 Creating a NAS share

Define the following Volume attributes in the Create Share dialog:

- **Name** – The share’s display name. It must be unique, and can be modified throughout the share’s lifetime

✓ **Note:** Objects names can be up to 128 chars long and can contain letters and digits, dashes “-” and under-scores “_”

- **Capacity** – Virtual Capacity of the Volume in GB. All Volumes are thinly provisioned. No actual capacity is allocated when the Volume is created, so the aggregated Virtual capacity of the volumes is not bounded by the Pool capacity. It is possible to over-provision a Pool, but you need to manage and monitor this it carefully, using a Pool Protection Mechanism (see [Managing Pool Capacity Alerts](#) for more details).
- **Export Name** – The name of the NFS/SMB mount point as seen by the Server. This must be unique. By default it is identical to the Share name.

✓ **Note:** In addition to the primary Export Name defined here, there is an option to add secondary Export Names to the same share. This can be done in the Volume properties page. See [Viewing Volume Properties](#).

✓ **Note:** Changing Export Name requires an unmount/remount of all NFS clients for the changed name to take effect.

- **Pool** – Select the Pool that is most appropriate for your Share’s QoS requirements (based on available capacity, caching, RAID protection etc.).
- **Attach Default snapshot Policy** – See [Managing Snapshot Policies](#) for a detailed explanation regarding snapshot policies. You can apply and remove snapshot policies from a Share at any time. If you select this checkbox you need to select one of the existing snapshot policies.
- **Filesystem Write Policy** - refer to different ways in which data is written to the underlying VPSA volume filesystem during filesystem operations.
 - **Asynchronous Writing (default)** When the filesystem is mounted with the “Asynchronous Writing” option, data modifications are not immediately synchronized with the volume file-system. Instead, the system buffers these changes in memory and may perform the actual write to the filesystem at a later time. This can lead to faster write performance as the system doesn’t need to wait for each individual write to complete before proceeding with other tasks.
 - **Synchronous Writing** When a filesystem is mounted with the “Synchronous Writing” option, all data modifications (writes) are immediately synchronized with the storage device. This means that before a write operation is considered complete, the data is physically written to the VPSA filesystem, ensuring that changes are safely stored on stable storage. This can ensure data integrity but can also lead to slower write performance, as the system waits for the filesystem to confirm the write operation before proceeding.
- **Encrypted** – Select this checkbox if you wish to encrypt the Share’s data on the drives. Please note that you must first define an encryption password via the [Controllers](#) Page. For more details about Volume encryption see [Managing Encrypted Volumes](#).
- **atime Update** – Set this checkbox to indicate whether you want to enable updating the access time of files and directories on every access, including read-access. (default: enabled).
- **User Quotas** – Select On or Off, to enable/disable the User Quotas mechanism for this Volume. for more information about quotas see [Setting User/Group Quotas](#).

- **Group Quotas** – Select On or Off to enable/disable the Group Quotas mechanism for this Volume.

✓ **Note:** If both User and Group quotas are “On” the first limit to be met takes effect.

- **Project Quotas** – Select On or Off to enable/disable the Project Quotas mechanism for this Volume. Project is defined as a set of folders (one or more) regardless of their User/group ownership. See [Setting Project Quotas](#).

✓ **Note:** Project and Group Quotas are mutually exclusive. One cannot define both on the same volume

- **File Access Audit** – Select On or Off to enable or disable the File Access Audit mechanism for this volume.

✓ **Note:** To be able to apply file access auditing on a volume, **File Access Audit** must be enabled globally in the [Security](#) tab on the **Settings** page.

SMB Options

- **SMB Only** – Set this checkbox if you know that this NAS share will only be attached to Servers via the SMB protocol. When this is the case the VP SA is able to do some locking optimization that enhances performance.
- **Allow Guest Access** – Set this checkbox if you want to enable connection and access to the NAS share by anonymous users without requiring a password.
- **Encryption Mode** – Select this to use SMB Encryption Secure protocol. Connected Windows hosts should support SMB encryption. See Microsoft MSDN for details: <https://learn.microsoft.com/en-us/archive/blogs/openspecification/encryption-in-smb-3-0-a-protocol-perspective>
Select “Off” to disable SMB Encryption, “Required” to enforce SMB Encryption (Windows host must enable encryption to connect) or “Desired” to let the client side decide if encryption is used or not.
- **Enhanced Windows ACLs** – Set this checkbox to enable the Enhanced Windows ACLs. These include support for Windows NT format ACLs, permission inheritance and additional extended attributes specific to Windows.
- **File Creation Mask** – Use this field to set the default bitmask used for file creation at the UNIX level.
- **Directory Creation Mask** – Use this field to set the default bitmask used for directory creation at the UNIX level.
- **Map Archive** – Set this checkbox to enable mapping of an archive bit. The DOS archive bit is used to flag a file that has been changed since it was last archived. Many programs do not work properly if the archive bit is not stored correctly for DOS and Windows files.
- **Browseable** – Select this checkbox for this share to be shown in the list of available shares in a network view and in the browse list.
- **Hidden Files** – Use this field to enter a list of files or directories that will not be visible, but will still be accessible. The DOS ‘hidden’ attribute is applied to any files or directories that match. Each entry in the list must be separated by a ‘/’, which allows spaces to be included in the entry. ‘*’ and ‘?’ can be used to specify multiple files or directories as in DOS wild cards. Each entry must be a UNIX path, not a DOS path, and must not include the Unix directory separator ‘/’. Note that this list is case sensitive.
- **Hide Unreadable** – Set this checkbox to prevent clients from seeing the existence of files that cannot be read.
- **Hide Unwritable** – Set this checkbox to prevent clients from seeing the existence of files that cannot be written to.
- **Store DOS Attributes** – Set this checkbox to preserve DOS file attributes Specifically , Hidden, Archive, Read-Only and System in the when creating/copying files into an SMB share. Turn on for compatibility with file system created on early NTFS versions.

- **SMB Serial small IO workload optimized** – Select this checkbox if your workload is serial small IOs from a single client (non concurrent)

NFS Options

- **NFS Root Squash** – Select this checkbox to block external root access to this share. If this box is checked, the system maps requests from uid/gid 0 (root) to the anonymous uid/gid.
- **NFS All Squash** – Select this checkbox to consolidate permission set for all users accessing this export (can be used to coordinate permissions between multiple server/applications or for setting up public file shares). If this box is checked, the system maps all external user requests to the anonymous uid/gid.

✓ **Note:**

- All Squash also applied for uid/gid 0 (root) making all squash and root Squash mutually exclusive
 - VMWare NFS V3 Mounts require NFS Root Squash & All Squash to be disabled (not checked)
-

- **NFS anonymous GID** – explicitly sets a specific group id for the anonymous account. this option is useful when set in conjunction with NFS Root/All Squash.
- **NFS anonymous UID** – explicitly sets a specific user id for the anonymous account. this option is useful when set in conjunction with NFS Root/All Squash.

File Lifecycle Management

- **Enable File Lifecycle Management Indexing** - The VPSA supports file lifecycle management and analytics. If the VPSA is configured for file lifecycle management and analytics, mark the checkbox to activate the file lifecycle management and analytics feature for this volume.

Click **Submit** to create a NAS Share with the supplied parameters.

✓ **Note:** Share creation involves the process of initializing a file system which may take a few minutes depending on the Virtual capacity of the Share. During this time the share is shown in a “Creating” state, but will be available for immediate use. When initialization is completed, the Share’s status changes to “Available” and an event-log message is saved.

6.1.3 Deleting a Volume/Share

You can delete a Volume only if it is not attached to a server.

On the **Volumes** page select the Volume and press the Delete button. After confirming that you want to delete, it will immediately move the Volume to “Deleting” status. The deletion process may take some time depending on the Volume size and the number of Snapshots and Clones which share the data Chunks. The VPSA then updates chunk mapping and references accordingly. When the deletion process completes, the Volume will disappear from the **Volumes** page, and an event-log message will be saved.

If the Volume has snapshots associated with it the VPSA will delete them together with the Volume. You will be prompted to confirm the deletion of the Snapshots as well.

Clones of the deleted Volume are **not** affected by the deletion of the Volume.

✓ **Note:** By default when you delete a volume it isn’t destroyed immediately, but it moves to the Pool’s Recycle Bin for 7 days until it is permanently deleted. From the Recycle Bin an administrator can purge (permanently delete) or restore the volume.

6.2 Filtering the List of Volumes

To filter the list of volumes displayed in the center pane, you can use one or more of the predefined filters, or a predefined custom tag.

1. Expand the **Filter** control.
2. On the right, click the **Add Filter** dropdown to select a filter. The selected filter appears on the left.
3. Repeat the **Add Filter** actions to select additional filters.

✓ **Note:** Selected filters are italicized and highlighted with a gray background in the **Add Filter** dropdown.

4. Refining the filtering:
 - For most filters, select one of its dropdown options.
 - The **Name** filter accepts input of a case-sensitive string. If the input string matches part or all of a volume's name, the volume is listed. Wildcards are not accepted.
 - The **Tag** filter requires input of a predefined custom **Tag Name**, and optionally, a **Tag Value** to further refine filtering for tags that have specific values for assigned volumes. Wildcards are not accepted.

✓ **Note:** Unlike the **Name** filter, **Tag Name** and **Tag Value** require full case-sensitive strings, and do not return matches on partial strings.

See the [Tags](#) section for configuring predefined custom tags.

5. Click **Search** to apply the filter.
6. To remove a single filter, click the trash icon to the right of the filter. Click **Search** again to refresh the volumes list.
7. To remove all the filters, click **Clear**.

6.3 Attaching & detaching Volumes to Servers

Volumes can be attached to many Servers. Block Volumes are attached via the iSCSI protocol. NAS Shares are attached via the NFS/SMB protocol.

To attach a Volume

Go to the [Volumes](#) page, select the Volume and press the Servers > Attach to Server(s) button:

- Select the Server(s) that you'd like to provide with access to the Volume.
- For NAS Shares, select the access type: NFS or SMB.

- For Block Volumes over Fibre Channel, select FC
- Press Submit to confirm.

Mounting an NFS Share on a Linux machine

1. Install the NFS client:

On Ubuntu Servers do:

```
apt-get install nfs-common
```

On Redhat/CenOS Servers do:

```
yum install nfs-utils
```

2. Create a mount point:

```
$ mkdir /mnt/nfs_share
```

3. Run the following command as the superuser (or with sudo):

```
$mount -t nfs4 <NFS_Export_Path>/<mount point>
```

You can find the NFS_Export_Path in the Volumes > Properties tab.

4. Follow the step in [Creating NAS Users](#) to setup basic NFS authentication.

Mounting an SMB Share on a Windows Server

1. On the Windows Server, go to **Computer > Map Network Drive**.
2. In the **Map Network Drive** dialog:
 1. **Drive:** Select a drive letter.
 2. **Folder:** Enter the SMB Export Path of the SMB share in the format \\<VPSA_IP>\<volume_export_name>. You can find the SMB Export Path parameter in the VPSA GUI **Volumes > Properties** tab.

✓ **Note:** On the first time that you connect from a Windows Server to a VPSA share, you are requested to enter an SMB User name and Password. See [Creating SMB Users](#) for more details, or use SMB guest access.

Format a Volume

Once the Volume is attached to the Server and identified by the Operating System as a drive, use the specific OS tools to format the drive to the needs of the OS or file-system used. Allocation units of 512B to 64KB are supported.

Important: Recommended best practice for Windows NTFS and ReFS performance

- Allocation unit size should match the LSA chunk size
 - Log structured array (LSA)

A zStorage pool uses a log structured array (LSA) to allocate capacity. A log structured array (LSA) has a directory that defines the physical placement of data blocks.

Each logical block device has a range of Logical Block Addresses (LBAs), starting from 0 and ending with the block address that fills the block device capacity.

An LSA enables allocating data sequentially (log structure), and its directory provides a lookup for matching the LBA with the physical address within the array.

An LSA does not overwrite existing data in its original location. Every write operation appends data at the end of the allocated space, and old data is marked for garbage collection. This reduces Read-Modify-Write (RMW) IO operations, and improves performance.

Since a pool uses an LSA to allocate capacity, a volume created from a pool consists of a directory that stores the allocation of blocks within the capacity of the pool.

- Formatting Windows NTFS drive

When formatting a Windows guest VM's NTFS drive, by default the allocation unit size is 4KB. For a large LSA chunk size, for example 32KB, a small allocation unit size of 4KB causes almost all IOs to be Read-Modify-Write (RMW) operations, that can impact performance.

Zadara recommends aligning the allocation unit size on Windows VM NTFS drives to match the LSA chunk size, for significantly reduced RMW on IOs.

✓ **Note:** For LSA chunk sizes, refer to the VPSA Storage Array **pool types** and VPSA Flash Array **pool types** tables under [Creating a Pool](#) in the **Configuring Storage Pools** page.

• Quick Removal policy for Windows virtual disk devices

The Windows **Better performance** disk device policy that enables write caching gains in performance at the expense of non-redundancy. In cases where customer applications already handle relaxed consistency by triggering flushes, or don't need strong consistency, they can use the default **Better performance** policy option.

For other use cases, Zadara advises disabling the write caching policy for Windows virtual disk devices, by configuring the Windows **Device Manager** > <guest VM> > **Disk drives** > **Properties** > **Policies** > **Removal Policy** to **Quick Removal**.

To detach a Volume

When you detach a Volume from a Server, the Server will lose access to the Volume's data. The recommended practice is to unmount the Volume on the Server side before detaching it on the VPSA.

To detach a Volume from a Server, go to the Volumes page, select a Volume and click **Servers** > **Detach from Server(s)**. You will be requested to select the Servers from which to detach this Volume.

Alternatively, you can view the attached Servers list in the Volume's South Panel, select the Server from which to detach the Volume, and click **Detach Server(s)** in the top-left corner of the South Panel.

6.4 Expanding a Volume

You can expand a Volume anytime, regardless of whether the Volume has Snapshots, Clones or is being remotely mirrored.

To expand a Volume go to the Volumes page, select the Volume and click **Expand**. Enter the amount of additional virtual capacity with which you'd like to expand the Volume, and press **Submit**.

6.4.1 Volume Automatic Expansion

To avoid out-of-space situations for File shares, the VPSA provides an Automatic Expansion mechanism.

It allows the customer to define an automatic NAS volume expansion policy.

Automatic Expansion is controlled by 3 parameters:

- **Emergency Threshold** - Volume will be expanded once the free capacity of the NAS share is below the given threshold. Default: 10% of the volume provisioned capacity.
- **Expand By** - The additional provisioned capacity to be added. Default: 50GiB
- **Maximum Volume Capacity** - The maximum allowed volume provisioned capacity (up to MAX Pool capacity) Default: 0GiB (Unlimited)

By default all volumes are created with Automatic Expansion disabled. To enable it, check the **Automatic Expansion** checkbox on the **Create > Create NAS share** dialog, or enable it from the **Properties** tab's **Capacity** column.

6.5 Managing SMB File History

SMB File History is a mechanism that allows restoration of previous versions of any given file or folder on a NAS volume, attached to Windows. SMB File History is similar to the VPSA snapshots mechanism, and driven by the same Snapshots Policies.

✓ **Note:** The deletion of these snapshots might affect snapshot policy routines, such as creating new snapshots or deleting expired ones.

If a folder from the SMB File History Snapshot remains open on a client, the snapshot deletion might fail.

In such cases, a "**Failed to unmount busy file history snapshot**" notification is sent to the VPSA owner, and a support ticket is opened.

To resolve this scenario, try one or both of the following:

- Ensure that all SMB File History Snapshot files are closed by the client.
 - On the active controller, [Restart File Services](#).
-

6.5.1 To Apply a SMB File History Policy on a Volume

1. Go to the Volumes page, and select the NAS Volume.
2. Select **Data Services > Attach File History Policy**.
3. In the **Apply Snapshot Policy to Volume** dialog, select the Snapshot Policy to apply to the Volume and press **Submit**.

6.5.2 To detach a SMB File History Policy from a Volume

1. Go to the Volumes page, and select the Volume and press the **Snapshot Policies** south tab to view the Volume's applied Snapshot Policies.
2. Select the Snapshot Policy to delete, and press the **Detach Policy** button at the top left corner of the South Panel.
You will be prompted to select whether or not to delete all of the Volume's Snapshots associated with this Policy.

6.5.3 To restore files from SMB File History

1. On a Windows Server, open Windows Explorer and navigate to the file/folder you want to restore.
2. Right-click on the file, and select **Restore previous versions**.
3. In the dialog that opens, go to the **Previous Versions** tab, select the version to restore, and click **Restore**.

✓ **Note:**

- Each share can keep up to 64 snapshots for File History recovery purposes, (e.g. once a day for a month) and maximum of 512 snapshots for a VPSA Storage Array.
 - When a Volume with SMB File History Snapshots is migrated to another Pool, the SMB File History snapshots will not be migrated to the new Pool.
-

6.6 Cloning a Volume

Cloning a Volume is the process of creating a Read/Write zero-capacity replica of a Volume, with a data set identical to that of the Volume, from a selected point-in-time (which can be the time the Clone is created, or one of the existing Snapshots' point-in-time).

The result of the Cloning operation is a new Volume. The two Volumes now share all of the non-modified chunks. Only upon a first-write to a chunk, a Copy-On-Write occurs which allocates a new chunk and breaks the chunk sharing.

You can create an unlimited number of Clones of a given Volume, either from the same data set (from the same Snapshot) or from different data sets.

Clones are completely independent from each other, from the source Volume and from the Snapshot from which they were created. For example, you can delete the original Volume and/or Snapshot and it will leave the Cloned Volume unaffected. You can also modify Volume attributes of each Clone independently.

You can only create Clones within the Pool where the original Volume resides.

To create a new Clone

1. Go to the Volumes page, select the Volume to be cloned and press **Data Services > Clone**.
2. In the **Clone Volume** dialog:
 - **Clone Name** – Enter a name for the Cloned Volume.
 - **Clone from** – Select the point-in-time Snapshot whose data set you wish to replicate.

If you wish to clone the current data set of the Volume, don't select any Snapshot.

3. Press **Submit**.

Alternative method: **Clone from Snapshot**

1. On the Volumes page, select the Volume to be cloned.
2. Press the **Snapshots** tab in the South Panel, and select the desired point-in-time Snapshot.
3. Press **Clone** at the top left corner of the South Panel.
4. In the **Clone from Snapshot** dialog, enter a name for the new cloned Volume.

The newly created Clone will appear as a regular Volume in the Volume list.

The NFS/SMB Export name of a cloned Volume will be identical to the Cloned Volume display name. |

6.7 Online Volume Migration

Volumes created in a VPSA pool can be easily migrated to a different pool in the same VPSA. All entities bounded to the volume (snapshot policies, servers attachments etc.) will be migrated as well. Existing snapshots migration is configurable by the user.

The online migration process is completely seamless to the end user and will not cause any service disruption to the hosts connected to the volume.

A common use case for using the Online Volume Migration feature is migrating performance demanding volume to a more performant storage pool(e.g. SATA pool to an SSD pool) on-the-fly.

Online Volume Migration can be initiated from the VPSA GUI or via VPSA REST API. For the REST API usage and examples please refer to the Volumes section of the [VPSA REST API Guide](#).

6.7.1 Migrating a Volume

In the left pane menu, navigate to the Volumes section under the Resources section.

1. Select the Volume that will be migrated to another VPSA Pool.
2. From the upper options menu select **Data Services > Migrate**.
3. In the **Migrate Volume** dialog:
 - **Destination pool** – Select the destination Pool to migrate to, from the list of available pools. Make sure to select a Pool with sufficient free capacity.
 - **Migrate Existing Snapshots** – Check the checkbox if the migration of the volume should include the existing snapshots of the volume. If **Migrate Existing Snapshots** is checked, all snapshots will be migrated to the destination Pool.



Note: If **Migrate Existing Snapshots** is not checked, the Volume snapshots will be deleted.

- VPSA Flash Array options:
 - **Compress** – Check the checkbox if you want the new volume to be compressed.
 - **Dedupe** – Check the checkbox if you want the new volume to be deduped.
- Press **Submit** to start the migration.
- In the **Create Migration** confirmation dialog, review the details and confirm the Online Volume Migration operation.

6.7.2 Monitoring the migration

Once started, the online migration task can be monitored from the VPSA GUI.

1. In the left pane menu navigate to the Volumes section under Resources.
2. Select the Volume that is currently being migrated.
3. On the south panel, the new **Migration Status** tab is available. The **Migration Status** tab will provide real-time migration information while the migration is still running.
4. The user has complete control on the migration task as it can be Paused or Aborted in the **Migration Status** tab.
5. Upon completion, the **Migration Status** tab will be removed from the Volume south panel. A log entry will be added in the **Logs** tab as an indication of a successful migration.

6.8 Managing Data Reduction

Scope: VPSA Flash Array

The ability to apply the Compress and Dedupe options on a Volume is dependent on the Data Reduction feature being enabled at the VPSA level in the Provisioning portal. See [Enabling or Disabling Data Reduction Bundle](#).

- Data Reduction on a VPSA can be enabled or disabled at any time.
- Compress and Dedupe can be enabled or disabled on a Volume at any time, on a VPSA that has Data Reduction enabled. Subsequent block write operations on the Volume follow the latest Compress and Dedupe Enabled/Disabled settings.
- To apply the Compress and Dedupe feature to data at rest on an existing Volume, the Volume should undergo a live migration to a new Storage Pool. The new Volume on the destination Storage Pool should be created with the Compress and Dedupe features enabled, so that inline Data Reduction is applied as the live migration occurs. See [Online Volume Migration](#).

6.9 Managing Encrypted Volumes

Encrypting data at rest is a highly recommended security measure to protect sensitive information from unauthorized access. This provides an extra layer of security, even if the data drive is stolen or compromised. It is especially important for organizations that handle sensitive information such as personal data, financial information, and confidential business information.

Encryption management of Data-at-Rest (data on the Disk Drives) is applied by the VPSA on a per-Volume basis. Encrypted and unencrypted Volumes can coexist in the same VPSA Pool.

Volume encryption has a negligible impact (if any) on volume performance. It's important to note that the benefits of encryption far outweigh any potential performance impact, as it ensures the protection of sensitive data.

A VPSA generates a random 256-bit unique Volume Encryption Key per encrypted Volume and uses the Advanced Encryption Standard (AES) to encrypt and decrypt the Volume data.

✓ **Note:** In previous versions of the VPSA software, AES 128 was used. Volumes that were created on those versions are encrypted with 128 bit keys.

The Volume Encryption Keys are stored on disk as ciphertext, using AES with a 256-bit Master Encryption Key, which is generated from a user-supplied **Master Encryption Password**.

The User owns the Master Encryption Password. It is never stored on any persistent media. Instead, only its SHA3 hash-sum is saved on disk for password validation.



Caution: Since the system does not keep the Master Encryption Password, you are **fully responsible to retain and protect the Master Encryption Password**.

During VPSA operation, the Master Encryption Password itself is held in kernel memory of the VPSA. Core-dumping any user-mode process within the VPSA will not reveal the Master Encryption Key.

This method ensures that encrypted Data-at-Rest cannot be accessed without explicitly knowing the user-supplied Master Encryption Password, thus providing you full protection if you opt for Data-at-Rest Volume encryption.

Important: The encryption attribute of a volume cannot be changed! If you'd like to encrypt the data of a non-encrypted volume, or vice versa, you will need to create a new volume and copy the data.

Supported volume encryption options:

- [Encryption using an Encryption Password](#)
- [Encryption using AWS KMS Store](#)
- [Encryption using KMIP supporting KMS](#)

6.9.1 Encryption using an Encryption Password

To create a Master Encryption Password:

1. Go to the Settings page.
2. In the **Security** tab, in the **Encryption** section press **Edit**.
3. Read the instructions and warning.
4. Enter your Password and **Save**.

Once the Master Encryption Password is set, you can change or reset it at any time. Master Encryption Password does not affect the encrypted data.

Store your Master Encryption Password in a secure place.

To create an encrypted volume, follow the steps in the [Creating and Deleting a Volume](#) section.

Encrypted volumes will present the Encrypted property in the VPSA Volumes view and in the Volume's properties.

6.9.2 Encryption using AWS KMS Store

Amazon's AWS Key Management Service (KMS) centrally manages keys and policies across integrated services and applications from a single point. AWS KMS generates a data key, encrypts it under the KMS key, and sends both the plaintext data key and the encrypted data key to Amazon S3. Amazon S3 encrypts the data using the data key and removes the plaintext key from memory as soon as possible after use.

Zadara VPSA supports use of AWS KMS for VPSA volume encryption.

Configuring AWS KMS encryption for a volume involves configurations on both the AWS Management Console and the VPSA admin GUI.

AWS KMS Configuration

- Create a new customer managed key:
 - Symmetric (default)
 - Encrypt & Decrypt (default)
 - Single region key (default)
- The user allowed to use the key should have access to the KMS Service along with the following action:
 - "kms:Encrypt"
 - "kms:Decrypt"
 - "kms:ReEncrypt*"
 - "kms:GenerateDataKey*"
 - "kms:DescribeKey"



Warning: KMS Key rotation is currently not supported!

VP SA Configuration

After setting up AWS KMS keys in the AWS Management Console, in the VP SA GUI, under **System**, open the **Settings** page, and configure the following parameters.

1. Click the **Security** tab.
2. Click **Edit** on the right of **Encryption**.
3. Click **Encryption using KMIP supporting KMS**.
4. Select **Encryption using AWS KMS Store**.
5. **Region**: Select the region of your AWS KMS Store.
6. **KMS Key ID**: Enter the UUID of your AWS Key ID.
7. **AWS Access Key**: Enter your AWS access key ID.
8. **AWS Secret Key**: Enter your AWS secret access key.
9. Read the KMS disclaimer. To acknowledge your acceptance of responsibility to maintain KMS access, mark the checkbox.
10. Click **Save**.

Once the system is set to encrypt volumes using the AWS KMS Service, the status for Encryption will be set to “Activated, AWS KMS”.

Encrypted volumes will present the Encrypted property in the VP SA Volumes view and in the Volume’s properties.

To create an encrypted volume, follow the steps in the [Creating and Deleting a Volume](#) section.

Disable AWS KMS Encryption

If volume encryption with AWS KMS is no longer required, click **Edit** and select **Remove AWS KMS Encryption**. The system will only allow the removal of the encryption setting if the VP SA does not have any encrypted volumes.

Encryption Key recovery

Similar to the other Encryption modes (KMIP or Encryption Password) during VP SA operation, the Master Encryption Password itself is held in kernel memory of the VP SA. Since it is not stored persistently in the VP SA, the user will be prompted to validate their AWS KMS configuration in case of a failure.

In such case:

1. The VP SA Administrator will be notified
2. Volumes will not be accessible
3. User action is required from the VP SA Settings view

In case of AWS KMS, the VP SA administrator will be required to enter the following information:

- AWS Region
- AWS Access Key
- AWS Secret Key

The same KMS key can be used by multiple IAM users, therefore in case of a rotation of the user’s credentials or in case the original credentials are no longer valid (access key & secret) a new pair can be used to restore VP SA to KMS integration given it is allowed to use the original KMS key.

6.9.3 Encryption using KMIP supporting KMS

The Key Management Interoperability Protocol (KMIP) enables the secure creation and storage of keys and other security objects on a key management server (KMS).

Zadara VPSA supports Fortanix Data Security Manager (DMS) for VPSA Storage Array volume encryption.

Configuring Fortanix DSM for a volume involves configurations on both the Fortanix DSM UI and the VPSA admin GUI:

- [Fortanix DSM configuration](#)
- [VPSA KMIP configuration](#)

Fortanix DSM configuration

1. In the **Groups** panel, create a new group with a meaningful name and click **Save**.
2. In the **Apps** panel, create a new app with the following parameters:
 1. **Name:** A meaningful name for the app.
 2. **Interface:** Select **KMIP**.
 3. **Authentication:** Select **API key** or **Certificate**. If you are uncertain, choose **API key**.
 4. Assign the group that you created in the previous step.
 5. Click **Save**.
3. In the **Security objects** panel, create a new object with the following parameters:
 1. **Name:** A meaningful name for the security object.
 2. **Group:** Select the group you created earlier.
 3. Select **GENERATE** as the key creation method.
 4. **Key type:** Select **AES** as the key type.

✓ **Note:** Zadara VPSA currently supports AES encryption for Fortanix DSM.

5. **Key size:** Select **256** for the AES key size.
6. **Key operations permitted:** Configure **both**:
 - ENCRYPT
 - DECRYPT
7. Accept the default settings for the rest of the configuration.
8. Click **GENERATE** at the bottom of the form.

VPSA KMIP configuration

In the VPSA GUI, under **System**, open the **Settings** page, and configure the following parameters.

1. Click the **Security** tab.
2. Click **Edit** on the right of **Encryption**.
3. Click **Encryption using KMIP supporting KMS**.
4. **KMS Type**: Select **Fortanix DSM**.
5. **KMS Host**: Select the region where the SmartKey KMS is registered. Supported regions:
 - North America
 - European Union
 - United Kingdom
 - Asia Pacific
 - Australia
6. **Connect Via**: Select the interface used by the VPSA to connect to the KMS. Currently the VPSA frontend and public IP are supported.
7. **KMS Key ID**: Enter the ID for the KMS key used for the VPSA volume keys encryption. For Fortanix DSM, this is the UUID of the key object.
8. **KMS Username**: Enter the username used by the VPSA for KMS authentication. It can be retrieved from the Fortanix DSM app by clicking on **View credentials > Username/Password**.
9. **KMS Password**: Enter the password used by the VPSA for KMS authentication. It can be retrieved from the Fortanix DSM app by clicking on **View credentials > Username/Password**.
10. **Use Proxy**: If you need a proxy server:
 1. Mark **Use Proxy**.
 2. Enter the proxy **Host** and **Port**.
 3. If the proxy requires authentication credentials:
 1. Mark **Use Authentication**.
 2. Enter the proxy's **User** and **Password**.
11. **Login with credentials + certificate**: VPSA KMS integration supports an enhanced security login mode using login credentials together with a certificate. To use the enhanced security login, mark **Login with credentials + certificate**.
 1. **Keyfile Content**: Copy your keyfile content, and paste it here.
 2. **Certfile Content**: Copy your certfile content, and paste it here.
12. Read the KMS disclaimer. To acknowledge your acceptance of responsibility to maintain KMS access, mark the checkbox.
13. Click **Save**.

To create an encrypted volume, follow the steps in the [Creating and Deleting a Volume](#) section.



Encrypted volumes are displayed with the icon.

KMS and VP SA Key rotation



Caution: Do not discard your old KMS key until **after** the key rotation is complete, as the VP SA still uses it to protect its master encryption password.

1. In the Fortanix DSM UI, go to your SmartKey account and rotate the key. After key rotation, copy the new KMS key's UUID.
2. In the VP SA GUI, under **System**, open the **Settings** page.
3. Click the **Security** tab.
4. Click **Edit** on the right of **Encryption**.
5. Paste the new KMS key's UUID into **KMS Key ID**.
6. Click **Save**.



Note: After applying the updated settings, the VP SA will re-encrypt its master key using the new UUID. On successful completion of this phase, the old KMS key can be discarded.

6.10 Audit Log Management

The VP SA supports audit logging of specific file system events.

6.10.1 Configuring file access auditing

The auditing policy must first be configured globally in the [Security](#) tab on the **Settings** page, before it can be applied to volumes.

File access auditing can be enabled when creating a new volume (see **Creating a NAS Share** under [Creating and Deleting a Volume](#)) and also on existing volumes.

6.10.2 Enabling file access auditing

To **Enable** file access auditing on an existing volume:

1. In the center pane, click on the volume to mark it.
2. Click the **Audit Log** dropdown and click **Enable**.

6.10.3 Downloading the volume audit log

To **Download** the audit log:

1. In the center pane, click on the volume to mark it.
2. Click the **Audit Log** dropdown and click **Download Audit Log**.
3. In the **Download Audit Log** dialog box that displays, enter the date and time range to download.

✓ **Note:**

- Audit logs are downloaded as a zip file comprising logs in CSV format.
 - The audit log file download is limited to a maximum size of 1GB. If the audit log data exceeds 1GB, extract it in multiple downloads of shorter date and time ranges.
-

Volume's audit log output

The audit log is exported to a field-delimited file, containing rows with 11 columns. In some cases, a 12th column is included, which contains extended parameters of the logged command.

Column	Parameter
1	Timestamp
2	VC name
3	Share name
4	Protocol
5	Client host IP address
6	Client hostname (if available)
7	User ID or username
8	Group ID or groupname
9	Command
10	Return status
11	Command-specific parameters
12	Extended command-specific parameters

✓ **Note:**

The timestamp value is according to the timezone displayed in the VP SA's UI header, usually UTC.

It is not possible to download audit logs with a different timezone.

The following table lists and describes possible logged commands:

Category	Command description	SMB command	NFS-3 command	NFS-4 command
CONNECT	User connects to the share	connect		
CONNECT	User disconnects from the share	disconnect		
CREATE	Create new regular file	create(file)	create(file)	create(file)
	Create new directory	mkdirat	mkdir	mkdir
	Create new symbolic link	symlinkat	symlink	symlink
	Create new hard link	linkat	link	link
	Create new special file		mknod(blk)	mknod(blk)
			mknod(chr)	mknod(chr)
			mknod(socket)	mknod(socket)
OPEN FILE	Open existing file	open(file)		open(file)
RENAME	Rename file or directory	renameat	rename	rename
REMOVE	Delete file	unlinkat	unlink	remove
	Delete directory	unlinkat	rmdir	remove
AT-TRIBUTES	Change permissions	fchmod fset_dos_attributes	setattr	setattr
		chmod (deprecated since v22.06) set_dos_attributes (deprecated since v22.06)		
	Change owner	fchown lchown	setattr	setattr
	Change ACLs	fset_nt_acl		

The following snippet depicts several lines of an audit log.

✓ **Note:** The header row is added in this example for informational purposes only, and is not part of an audit log output.

```

Timestamp          |VC name          |Share|Protocol|IP address |host |User ID or name |Group ID
↳or name          |Command|St|Command parameters

2024-09-11 08:21:43.435816|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↳domain users|connect|ok|smb1
2024-09-11 08:21:57.133796|vsa-000001ab-vc-0| smb2|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↳domain users|connect|ok|smb2
2024-09-11 08:22:56.135063|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↳domain users|mkdirat|ok|/export/smb1/Test
2024-09-11 08:22:56.138699|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↳domain users|fset_dos_attributes|ok|/export/smb1/Test
2024-09-11 08:22:56.141804|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↳domain users|fset_nt_acl|ok|/export/smb1/Test|<extended parameters>
2024-09-11 08:22:56.150438|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↳domain users|mkdirat|ok|/export/smb1/Test/proc0
2024-09-11 08:22:56.151271|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↳domain users|fset_dos_attributes|ok|/export/smb1/Test/proc0
2024-09-11 08:22:56.154084|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↳domain users|fset_nt_acl|ok|/export/smb1/Test/proc0|<extended parameters>
2024-09-11 08:22:56.277061|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↳domain users|mkdirat|ok|/export/smb1/Test/proc1
2024-09-11 08:22:56.277842|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\

```

(continues on next page)

(continued from previous page)

```

↪ domain users|fset_dos_attributes|ok|/export/smb1/Test/proc1
2024-09-11 08:22:56.281236|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↪ domain users|fset_nt_acl|ok|/export/smb1/Test/proc1|<extended parameters>
2024-09-11 08:22:56.296915|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↪ domain users|create(file)|ok|/export/smb1/Test/proc0/data_0.din|w
2024-09-11 08:22:56.298040|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↪ domain users|fset_dos_attributes|ok|/export/smb1/Test/proc0/data_0.din
2024-09-11 08:22:56.300355|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↪ domain users|fset_nt_acl|ok|/export/smb1/Test/proc0/data_0.din|<extended parameters>
2024-09-11 08:22:56.305856|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↪ domain users|create(file)|ok|/export/smb1/Test/proc1/data_0.din|w
2024-09-11 08:22:56.306814|vsa-000001ab-vc-0| smb1|SMB3_11|<IP address>|vm123|ZSTORAGE2\user123|ZSTORAGE2\
↪ domain users|fset_dos_attributes|ok|/export/smb1/Test/proc1/data_0.din

```

6.10.4 Disabling file access auditing


To **Disable** the audit log for a volume:

1. In the center pane, click on the volume to mark it.
2. Click the **Audit Log** dropdown and click **Disable**. A confirmation dialog box opens.

✓ **Note:** Disabling a volume's audit logs does not remove audited entries.

6.11 Volume File Lifecycle Management

The VPSA supports file lifecycle management and analytics. When file lifecycle management is enabled for a VPSA, the following options are available for configuring each volume:



- **Indexing Enable** - Selecting this option activates file lifecycle management and analytics on the selected volume.
- **Indexing Pause** - Suspend analytics collection for the selected volume. This option is available for volumes that are enabled for file lifecycle and analytics.
- **Indexing Resume** - Resume analytics collection for the selected volume. This option is available for volumes that are enabled for file lifecycle and analytics, and their indexing is paused.
- **Indexing Disable** - Selecting this option deactivates file lifecycle management and analytics on the selected volume.

✓ **Note:** Disabling file lifecycle indexing for a volume removes all existing data collected for that volume.

6.12 Viewing Volume Properties

Filtering Volumes

In a VPSA with many volumes it might be difficult to locate a specific volume in the Volumes page. [Filtering the List of Volumes](#) can be useful.

The Volumes page displays the list of Volumes (Block and NAS) in the VPSA. Select a Volume to see its detailed information in the following South Panel tabs:

6.12.1 Properties

Each Volume includes the following properties:

Property	Description
	General
ID	An internally assigned unique ID.
Name	User assigned name. Can be modified anytime.
Comment	User free text comment. Can be used for labels, reminders or any other purpose
Status	<ul style="list-style-type: none"> - Creating - Initializing Volume's metadata. - Deleting - In process of deleting the Volume and updating data chunks references. - Partial/Failed - The Volume is inaccessible due to lower construct failure (on Pool or RAID Group level). - Available - The Volume is healthy but is not attached to any Server. - In-use - The Volume is healthy and is attached to one or more Servers.
Data Type	<ul style="list-style-type: none"> - "Block" for Block Volume. - "File-system" for NAS Shares.
Pool	The Pool name where this Volume is provisioned.
Server(s)	Server Name attached to the Volume. Multiple(X) will be displayed when X servers are attached.
NFS Export Path	(NAS Shares Only) The NFS Share export path to be used when mounting it. All defined paths are listed here. Additional path can be defined.
SMB Export Path	(NAS Shares Only) The SMB Share export path(s) to be used when connecting to it from a Windows Server. All defines paths listed.

continues on next page

Table 2 – continued from previous page

Property	Description
Access Type	(NAS Shares Only) Access protocols which are used by the Servers which are attached to a NAS Share: NFS, SMB, or Multiple.
atime Update	(NAS Shares Only) Yes/No – Indicates whether to update access time of NAS Share files and directories on every access, including read-access.
SMB Only	(NAS Shares Only) Yes/No – enable/disable locking optimizations
SMB Guest Access	(SMB Only) Yes/No – Allow/Block anonymous user access
SMB Encryption Mode	(SMB Only) Off/Desired/Required – Sets SMB encrypt secured protocol behaviour
Enhanced Windows ACLs	(SMB Only) Yes/No
Directory Creation Mask	(NAS Shares Only) Default directory umask value
File Creation Mask	(NAS Shares Only) Default file umask value
Map archive	(NAS Shares Only) Yes/No – Maps the windows archive bit to the unix execute bit.
SMB Browsable	(SMB Only) Yes/No – seen in the list of available shares
SMB Hidden Files	(SMB Only) This is a list of files or directories that are not visible but are accessible (delimited by /)
SMB Hide Unreadable	(SMB Only) Yes/No – Prevents clients from seeing the existence of files that cannot be read.
SMB Hide Unwritable	(SMB Only) Yes/No – Prevents clients from seeing the existence of files that cannot be written.
SMB Hide Dot Files	(SMB Only) Yes/No – Prevents clients from seeing the existence of “.” files.
SMB serial small IO workload Optimized	(SMB Only) Yes/No
SMB Store DOS Attributes	(SMB Only) Yes/No – Preserve DOS attributes (hidden, archive, read-only, system)
User Quotas	(NAS Shares Only) On/Off – user quotas on volume.
Group Quotas	(NAS Shares Only) On/Off – group quotas on volume.
Project Quotas	(NAS Shares Only) On/Off – Project quotas on volume.
NFS Root Squash	(NFS Only) Yes/No – map requests from uid/gid 0 (root) to the anonymous uid/gid. Note: Set to “Yes” to block external root access to the volume.
NFS All Squash	(NFS Only) Yes/No – map requests from and uid/gid to the anonymous uid/gid. Note: Useful for inter server/application correlation or Public File shares

continues on next page

Table 2 – continued from previous page

Property	Description
NFS anonymous GID	(NFS Only) explicitly sets a specific group id for the anonymous account
NFS anonymous UID	(NFS Only) explicitly sets a specific user id for anonymous account
File Lifecycle Index Management	Indicates the state of indexing for analytics data collection. Possible values: Enabled/Disabled for analytics data collection. Paused when analytics data collection is enabled, but pending resume.
File Lifecycle Index Management Full Scan State	When File Lifecycle Index Management is enabled, the VPSA performs a single full scan of NAS share files. Subsequent detected filesystem changes are updated in the indexing. Possible statuses of the volume's full scan: In Progress: indicates the full scan progress as a percentage. Finished: indicates that the full scan is complete. Paused: indicates that File Lifecycle Index Management is paused. Disabled: indicates that File Lifecycle Index Management is disabled.
Extended Metering	Yes/No – Enabling extended metering. When “Extended Metering” is disabled, the VPSA records the volume's performance statistics of reads and writes operations. When “Extended Metering” is enabled, the VPSA also records performance statistics of other file operations, including create, delete, etc.. Note: “Extended Metering” enabled puts extra load on the VPSA, and the metering DB might grow rapidly. It is recommended to use it for only limited period of time, for planning or troubleshooting purposes.
WWID	(Block Only) SCSI unique World-wide ID. Use this value on Linux Servers to identify the Volume device when multi-pathing is configured.

continues on next page

Table 2 – continued from previous page

Property	Description
Compress VPSA Flash Array	Indicates the state of data compression. Possible values: Enabled/Disabled Compression can be enabled or disabled at any time. See Enabling or Disabling Data Reduction Bundle . Subsequent block write operations on the volume follow the latest Enabled/Disabled setting.
Dedupe VPSA Flash Array	Indicates the state of data deduplication. Possible values: Enabled/Disabled Deduplication can be enabled or disabled at any time. See Enabling or Disabling Data Reduction Bundle . Subsequent block write operations on the volume follow the latest Enabled/Disabled setting.
Encrypted	Yes/No
Encryption Size Key	The number of bits in a key used by the encryption algorithm.
Created	Date & time when the Volume was created.
Modified	Date & time when the Volume was last modified.
	Capacity
Virtual Capacity	Capacity of the Volume as seen by the attached Servers.
Available Capacity	(NAS Shares Only) Free capacity of the NAS Share.
Mapped Capacity	The used capacity (allocated from the Pool) of the Volume excluding its Snapshots and Clones.
Data Copies Capacity	The used capacity (allocated from the Pool) of the Volume's Snapshots and Clones. Note: the total capacity allocated for a Volume and all its Clones and Snapshots is the sum of Mapped Capacity + Data Copies Capacity
Read IOPS	The maximum number of read operations per second.
Write IOPS	The maximum number of write operations per second.
Read MBPS	The maximum throughput of data in Megabytes per second for read operations.
Write MBPS	The maximum throughput of data in Megabytes per second for write operations.

6.12.2 Snapshots

Lists the point-in-time Snapshots of this Volume. If you retain many Snapshots per Volume, you may want to use the Snapshot Filtering tool to find a specific Snapshot. For more details see [Filtering Snapshots](#).

The following Properties are provided per Snapshot:

Attribute	Description
ID	Snapshot ID
Name	Display Name.
TimeStamp	Snapshot creation time stamp
Status	Normal\Pending Deletion\Deletion

6.12.3 Object Storage Snapshots

Lists the point-in-time Snapshots of this Volume which are stored in an Object Storage (e.g S3). These Snapshots are created by the Backup to Object Storage feature, as defined in [Backup to Object Storage](#)

The following Properties are provided per Object Storage Snapshot:

Attribute	Description
ID	Snapshot ID
Name	Display Name.
Region	Object storage region
Bucket	Object storage bucket
TimeStamp	Snapshot creation time stamp
Status	Normal\Pending Deletion\Deletion

6.12.4 SMB File History (SMB Only)

Lists the point-in-time Snapshots of this Volume which are kept for SMB File History recovery purposes. These Snapshots are created by the SMB File History mechanism. For details see [Managing SMB File History](#).

The following Properties are provided per File History Snapshot:

Attribute	Description
ID	Snapshot ID
Name	Display Name.
TimeStamp	Snapshot creation time stamp
Status	Normal\Pending Deletion\Deletion
Pool	Pool where the file history is kept

6.12.5 Snapshot Policies

The Snapshot Policies tab lists the policies that are attached to the selected Volume. The following Properties are provided per Snapshot Policy:

Attribute	Description
Name	Display Name.
Status	Active or Paused.
Type	The VPSA application controlling the Policy: <ul style="list-style-type: none"> • Snapshot Manager • Remote Mirroring • Backup to Object Storage • SMB File History
Create Policy	Frequency of Snapshot creation.
Delete Policy	Number of Snapshots to retain.
Dest. Delete Policy	Number of Snapshots to retain on Remote Mirror destination Volume.

For more details on Snapshot Policies management, see [Managing Snapshots and Snapshot Policies](#).

6.12.6 Servers

The Servers tab lists the Servers to which the Volume is attached. For Block Volumes the LUN Number associated with each Server is displayed. It also indicates if the server accesses the volume via iSCSI or FC.

6.12.7 Containers

Lists the Docker Containers that are able to access the selected Volume, along with their statuses. For details about attaching Volumes to Containers see [Containers](#).

6.12.8 Metering

The Metering Charts provide live metering of the IO workload associated with the selected Volume.

The charts display the usage data as it was captured in the past 20 “intervals”. An interval length can be set to one of the following: 1 Second, 10 Seconds, 1 Minute, 10 Minutes, or 1 Hour. The Auto button lets you see continuously-updating live metering info (refreshed every 3 seconds).

The following charts are displayed:

Chart	Description
IOPs	The number of read and write SCSI commands issued to the selected Volume from all attached Servers.
Bandwidth (MB\s)	Total throughput (in MB) of read and write SCSI command issued to the selected Volume from all attached Servers.
IO Time (ms)	Average response time of all read and write SCSI command issued to the selected Volume from all attached Servers.

6.12.9 Logs

Displays all event logs associated with this Volume.

6.12.10 Performance Alerts

Displays Performance Alerts for the selected Volume.

- **Read IOPS Limit** – Creates an alert when, during the past minute, the average read IOPS for the selected Volume exceeds a user-specified threshold.
- **Read Throughput Limit** – Creates an alert when, during the past minute, the average read MB/s for the selected Volume exceeds a user-specified threshold.
- **Read Latency Limit** – Creates an alert when, during the past minute, the average read latency for the selected Volume exceeds a user-specified threshold.
- **Write IOPS Limit** – Creates an alert when, during the past minute, the average write IOPS for the selected Volume exceeds a user-specified threshold.
- **Write Throughput Limit** – Creates an alert when, during the past minute, the average write MB/s for the selected Volume exceeds a user-specified threshold.
- **Write Latency Limit** – Creates an alert when, during the past minute, the average write latency for the selected Volume exceeds a user-specified threshold.

6.12.11 Capacity Alerts

Displays capacity Alerts for the selected NAS Volume. The Capacity Alerts tab lists the configurable attributes of the NAS Volume capacity Protection Mechanism, similar to the pool capacity alerts. See [Managing Pool Capacity Alerts](#) for more details.

- **Alert Threshold** – Creates an alert when it is estimated that the Volume will be at full capacity in X Minutes.
 - Default Value: 360 minutes
- **Alert Interval** – Calculates the estimated time until the Volume is full based on the capacity usage in the previous X minutes.
 - Default Value: 60 minutes
- **Emergency Threshold** – Creates an alert when the volume is running out of free space and reaching the given threshold.”
 - Default Value: 1 GB

6.12.12 File Lifecycle

The File Lifecycle tab provides a shortcut button to navigate directly to the file lifecycle analytics page for the selected volume.

6.12.13 Tags

Predefined custom tags can be configured in the **Tags** tab. An example use case for tags is [Filtering the List of Volumes](#) in the center pane.

A tag is identified by its **Tag Name** and has a **Tag Value** associated with it. A tag can be defined only once for a volume. However, the same **Tag Name** can be defined with a different **Tag Value** for other volumes.

- **Create:** To create a new tag for a volume, in the volume's **Tags** tab click **Create**, and enter the **Tag Name** and **Tag Value**. The tag is added to the list of tags in the **Tags** tab.
- **Edit:** To change the **Tag Value** of an existing tag, click on that tab in the tags list to mark it, and then click **Edit**. The **Edit Tag** dialog box opens, allowing overwriting of the **Tag Value**.

✓ **Note:** Only the **Tag Value** can be edited. A tag cannot be renamed. It must be deleted, and a tag with the new name configured in its place.

- **Delete:** To delete a tag, click on that tag row in the tags list to mark it, and then click **Delete**. A confirmation dialog box opens.
- **Refresh:** Displays the updated tags list.

SERVERS

Servers Objects in the VPSA represent Cloud Servers that consume VPSA Volumes. A Server needs to be properly defined and connected in order to access the VPSA Volumes via iSCSI, FC, NFS or SMB protocols.

7.1 Filtering the List of Servers

To filter the list of servers displayed in the center pane, you can use a predefined custom tag.

The screenshot shows the 'SERVERS' management interface. At the top, there is a toolbar with icons for REFRESH, ADD, CONFIG, DELETE, and VOLUMES. Below the toolbar is a 'Filter' section. It contains a 'Tag Name' input field, a 'Tag Value' input field with a dropdown menu set to 'Optional', and a trash icon. To the right of these fields is an 'Add Filter:' dropdown menu. Below the input fields are 'SEARCH' and 'CLEAR' buttons. The 'Add Filter:' dropdown menu is open, showing a list of tags, with 'Tag' selected.

1. Expand the **Filter** control.
2. On the right, click the **Add Filter** dropdown, and select **Tag**. The tag input fields appear on the left.
3. The **Tag** filter requires input of the predefined custom **Tag Name**, and optionally, a **Tag Value** to further refine filtering for tags that have specific values for assigned servers. Wildcards are not accepted. See the [Tags](#) section for configuring predefined custom tags.
4. Click **Search** to apply the filter.
5. To remove the filter, click **Clear**, or click the trash icon to the right of the filter. Click **Search** again to refresh the volumes list.

7.2 Adding a Server

Establishing a connection between a Server and the VPSA involves the following steps:

- Creating a Server Object in the VPSA database.
- Setting the Server IQN for iSCSI connectivity and/or Servers FC Connectivity and/or the server IP address for NFS/SMB connectivity.
- Establishing CHAP authentication handshake between the Server and the VPSA for iSCSI.
- Registering Server OS information (optional).

Important: The VPSA can provide server connectivity monitoring alerting upon connectivity issue with a specific server. Connectivity monitoring can be enabled on demand for any server with IP connectivity, for additional information refer to [Connectivity Monitoring](#) in this user guide section.

7.2.1 Adding a Server for NAS access

Adding servers to the VPSA in order to access files over NFS/SMB requires introducing the server's IP address to the VPSA.

- Go to Servers > Add and select Manual:
- On the Create Server dialog give the server a name

✓ **Note:** Objects names can be up to 128 chars long and can contain letters and digits, dashes “-” and underscores “_”.

- Select the server's Operating System (optional)
- Turn on File Access
- Provide the IP address

✓ **Note:** You can add a single server object to the VPSA representing an IP Network Range rather than adding each Server in the range separately. This is especially useful when attaching SMB/NFS shares to large number of servers in a subnet. Use the manual procedure shown below to add this type of Server while specifying the IP range in CIDR notation (e.g. 192.168.11/24)

- If you want to secure the IP connectivity with an IPsec tunneling check select the Enable IPsec checkbox. Please note that your Server must be properly configured to utilize IPsec, and that performance is impacted.
- An alternative to IPsec tunneling is SMB Encrypt that works for Windows servers that support it. SMB Encryption provides end-to-end encryption of SMB data and protects data from eavesdropping occurrences on untrusted networks. It has no requirements for Internet Protocol security (IPsec tunneling) and is much easier to configure. SMB Encryption can be configured on a per share basis. No setting is required on the VPSA. Just enable SMB Encryption on the Windows Server.

✓ **Note:** iSCSI or Fiber Channel are not required for a NFS/SMB connection so these settings can be left OFF.

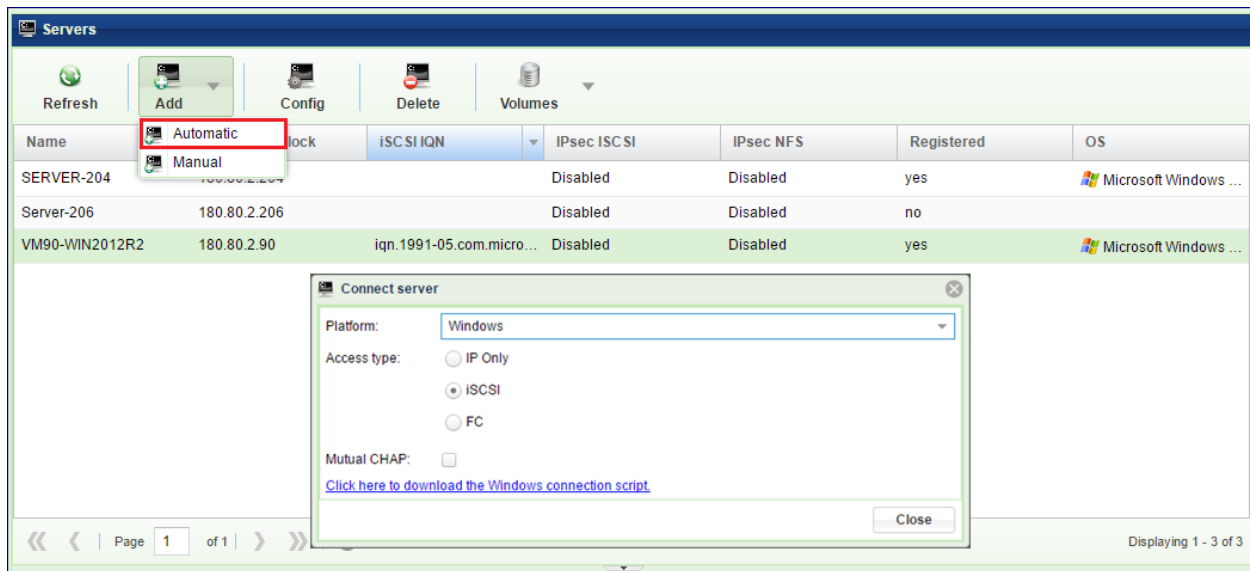
✓ **Note:** The VP SA NAS services will require the following ports and protocols to be accessible from the servers to the VP SA:

- NFS - 111(UDP/TCP), 2049(UDP/TCP), 3000(UDP/TCP), 4000(UDP/TCP), 4001 (UDP/TCP), 4045(UDP/TCP).
- SMB - 137(UDP), 138(UDP), 139(TCP), 445(TCP).

To add a server for block storage access follow the procedure shown below:

7.2.2 Adding a Server automatically


The VP SA automates the above steps for you via the “Connect Server” script. Go to **Servers > Add** and select Automatic:



Adding a Server automatically over iSCSI

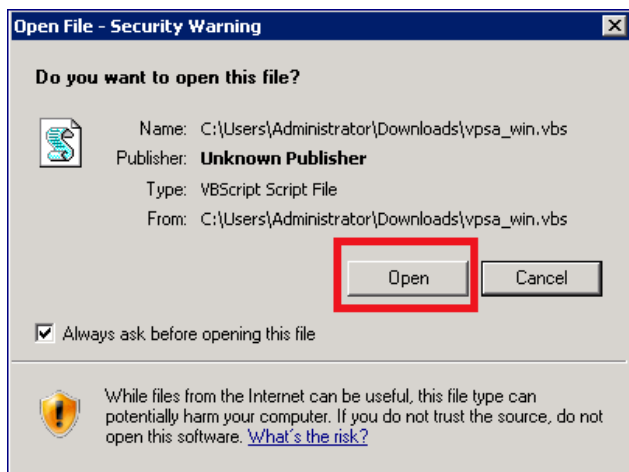

To Add a Windows Server:

- The first time you connect an iSCSI Volume to a Windows Server, you need to start the iSCSI service on the Windows Server **before** running the VP SA connect script.
 - In Windows Start->Run dialog, type iSCSI and select the “iSCSI Initiator” program. You will be prompted to start the service. Press Yes to confirm:



- Open the VP SA GUI on the Windows Server

- On the **VPSA GUI > Connect Server** dialog, select platform: Windows.
- Select the iSCSI checkbox if you wish to expose VPSA Block Volumes to this Server via iSCSI.
- Click the download link to download the connect script from the VPSA to your Server.
- Depending on your browser, locate the downloaded script, open and run it. The below screenshots are using the Chrome browser.



- Once the connect script successfully completes, the new connected Server will be listed in the VPSA Servers page with status = "Active" Registered = "Yes" and the correct OS details.

To Add a Linux Server:

- Verify that open-iscsi is installed on the Server:

- On RedHat Servers do:

```
$ yum install iscsi-initiator-utils
```

- On Ubuntu Servers do:

```
$ sudo apt-get update
$ sudo apt-get install open-iscsi open-iscsi-utils
```

- On the [VPSA GUI > Connect Server](#) dialog, select platform: Linux.
- Select the iSCSI checkbox if you wish to expose VP SA Block Volumes to this Server.
- Run the three steps as detailed in the connect server dialog to execute the vpsa_linux.sh script.

Connect Server

Platform:

iSCSI: ☒

Step 1

```
sudo wget -O ~/vpsa_linux.sh -
"https://172.31.240.108/api/servers/vpsa_linux.sh?
access_key=AL2GMYQT5ZH7Q92GD71C&iSCSI=yes" --no-check-certificate
```

Step 2

```
sudo chmod u+x ~/vpsa_linux.sh
```

Step 3

```
sudo ~/vpsa_linux.sh
```

Close

- Once the connect script completes successfully, the new connected Server will be listed in the [VPSA GUI > Servers](#) page with status = "Active" Registered = "Yes" and the correct OS details.

Multiple Virtual Networks:

If the VP SA has more than one Virtual Network assigned, the server can be assigned to it over any of them. Servers can be added using any VNI, automatic server registration in the presence of VNIs shows drop down list of available VNIs for you to select:

Connect server

Platform:

Source:

iSCSI:

Mutual CHAP:

[Click here to download](#)

Virtual Network Interface (VLAN 81) - 14.14.10.100

Virtual Network Interface (VLAN 50) - 10.2.10.231

Virtual Network Interface (VLAN 84) - 20.20.10.100


Virtual Network Interface (VLAN 52) - 21.21.10.100

Adding a Server automatically over Fibre Channel


Before defining a server using FC connectivity, make sure to install and configure multipathing software on the server. Also make sure to setup the zoning on the FC switch to allow connectivity between the connecting server and the VP SA FC ports.

To Add a Windows Server:

The first step is to configure MPIO. Open the Windows MPIO dialog:

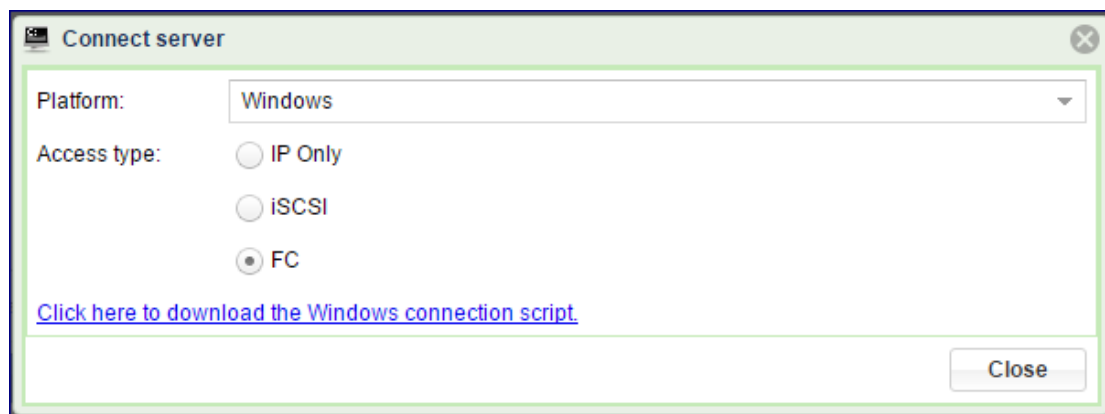


Enter the device HW ID as follows: “Zadara VP SA ” (“Zadara” followed by 2 blank spaces, and then “VP SA” followed by 4 blank spaces. 16 characters total) and press OK.



✓ **Note:** Adding MPIO requires Windows server to reboot before you can proceed.

- Open the VP SA GUI on the Windows Server
- On the [VP SA GUI > Connect Server](#) dialog, select FC access type and download the connection script



- Run the script as described above in [Adding a Server automatically over iSCSI](#).

To Add a Linux Server:

To add Linux Server over FC you need to set up the multipathing on the server. For detailed instructions follow this KB article:

<https://support.zadara.com/hc/en-us/articles/115003851406-How-To-Setup-Fiber-Channel-Multipath-in-Linux>

✓ **Note:** This change requires a restart of the MPIO service

- On the VP SA GUI goto the [Connect Server](#) dialog, select FC access type and download the connection script.
- Run the three steps as detailed in the [Connect Server](#) dialog to execute the vpsa_linux.sh script.

Connect server

Platform: Linux

Access type:

☐ IP Only

☐ iSCSI

☒ FC

Step 1

```
sudo wget -O ~/vpsa_linux.sh "https://180.80.2.103/api/servers/vpsa_linux.sh?iscsi=no&fc=yes&access_key=4I7K7ESILP07W2BJ56NI" --no-check-certificate
```

Step 2

```
sudo chmod u+x ~/vpsa_linux.sh
```

Step 3

```
sudo ~/vpsa_linux.sh
```

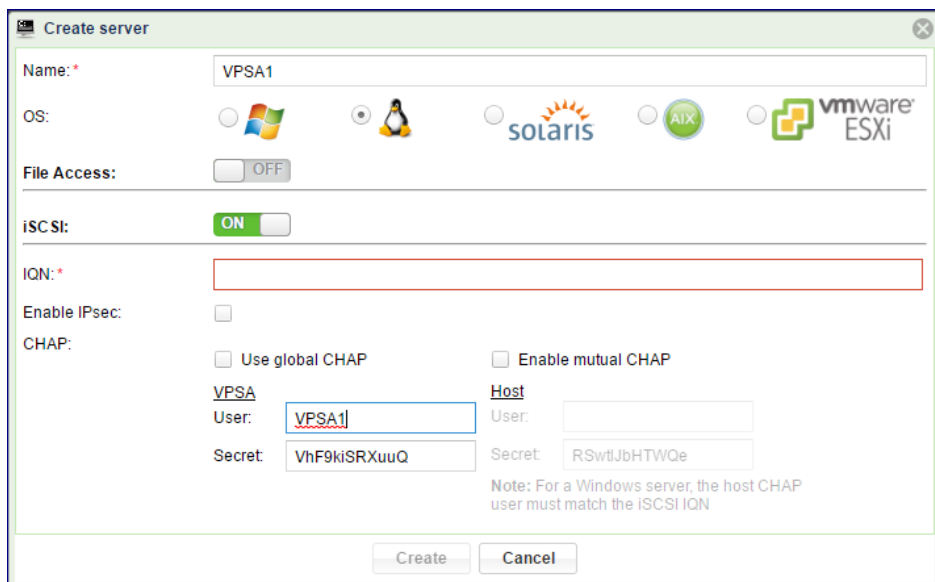
Close

- Once the connect script completes successfully, the new connected Server will be listed in the VP SA Servers page with status = “Active” Registered = “Yes” and the correct OS details.

7.2.3 Adding a Server manually

Establishing an iSCSI connection

If for some reason adding a server automatically doesn't work, follow these steps to add the server manually. Go to [Servers > Add](#) and select Manual:



The "Create server" dialog box contains the following fields and options:

- Name:** A text field containing "VPSA1".
- OS:** Radio buttons for Windows, Linux (Tux), Solaris, AIX, and VMware ESXi.
- File Access:** A toggle switch set to "OFF".
- iSCSI:** A toggle switch set to "ON".
- IQN:** An empty text field with a red border.
- Enable IPsec:** An unchecked checkbox.
- CHAP:**
 - ☐ Use global CHAP
 - ☐ Enable mutual CHAP
- VPSA User:** A text field containing "VPSA1".
- VPSA Secret:** A text field containing "VhF9kiSRXuuQ".
- Host User:** An empty text field.
- Host Secret:** A text field containing "RSwtIjbHTWQe".
- Note:** "For a Windows server, the host CHAP user must match the iSCSI IQN".
- Buttons:** "Create" and "Cancel".

- Enter the Server Name.


✓ **Note:** Objects names can be up to 128 chars long and can contain letters and digits, dashes "-" and underscores "_".

- Select the server OS
- Enter the server iSCSI IQN
- Check the "Enable IPsec" checkbox if you wish to secure iSCSI traffic between the Server and the VPSA. Please note that your Server must be properly configured to utilize IPsec and that performance is impacted.
- To enable CHAP, select between global CHAP (for the VPSA) or per host.
- Provide the CHAP user name and password (secret). Global CHAP parameters can be copied from here [Viewing Controller Properties](#).

After manually adding a Server you need to establish an iSCSI connection between the Server and the VPSA. Please note that you can skip this step if the Server was added automatically or if the Server is only consuming NFS/SMB type Volumes.


On Windows Servers:

- Open iSCSI Initiator: In Windows Start->Run dialog, type iSCSI and select the "iSCSI Initiator" program. If this is the first time you have run iSCSI initiator on this Server you will be prompted to start the service. Press Yes to confirm.




- The Windows iSCSI Initiator Properties dialog box will open, and the Targets tab will be displayed.


- On the Targets tab, type the iSCSI IP address of the VPSPA (which is displayed in the [VPSPA GUI > Controllers](#) page) in the Quick Connect target text box and then click the Quick Connect... button.
- The Quick Connect dialog box will be displayed, with the VPSPA discovered iSCSI target in an “Inactive” status. Press Done.



- To activate the connection, select the VPSPA target and press the Connect button. Please note that if you have multiple targets listed you can identify the VPSPA target by its IQN name which is in the form of “iqn.2011-04.com.zadarastorage:vsa-xxxx” and is displayed in the Controller properties page in the VPSPA GUI.
- You may check the Enable multi-path check-box if you wish to use MPIO multi-pathing. Then, click Advanced...



- Check the Enable CHAP log-on check-box and enter the CHAP Username: and Target Secret. You can retrieve those values from the VP SA GUI, under the [Controllers](#) page, in the properties tab. Press OK to confirm the operation.



- In the Targets tab you'll see that the VPSA iSCSI target has moved from "Inactive" to "Connected" status. A new Server is created automatically in the VPSA and is displayed in the Servers GUI page. The name of the server is its iSCSI initiator IQN. You may change the Server Display Name.

✓ Note: To achieve best performance it is recommended to use multiple sessions & MPIO. To enable MPIO please follow the instructions at <https://support.zadarastorage.com/hc/en-us/articles/360030099351-How-To-enable-iSCSI-MPIO-and-set-up-multiple-iSCSI-sessions-on-Windows-Server-2016>.

On Linux Servers:

Locate the VPSA iSCSI IP address and the CHAP Username and Password in the VPSA GUI Controller Properties Page:

Run the following commands to issue an iSCSI login using CHAP credentials:

```
$ iscsiadm -m node -T <VPSA-Target-IQN> -p <VPSA-Management-IP> --op new
$ iscsiadm -m node -T <VPSA-Target-IQN> -p <VPSA-Management-IP> --op update -n node.session.auth.
authmethod -v CHAP
$ iscsiadm -m node -T <VPSA-Target-IQN> -p <VPSA-Management-IP> --op update -n node.session.auth.username
-v <CHAP-username>
$ iscsiadm -m node -T <VPSA-Target-IQN> -p <VPSA-Management-IP> --op update -n node.session.auth.password
-v <CHAP-secret>
$ iscsiadm -m node -T <VPSA-Target-IQN> -p <VPSA-Management-IP> --login
```

Where:

- VPSA-Target-IQN – Target IQN of the VPA. Can be found in the **VPSA GUI > Controllers** page, Properties South Panel, Target parameter. It is of this format:

```
iqn.2011-04.com.zadarastorage:vsa-000009e5:1
```

- VPSA-Management-IP - The iSCSI IP of your VPSA. Can be found in the [VPSA GUI > 'Controllers'](#) page, under the iSCSI IP column.

To ensure automatic login of your Server to the VPSA after each reboot (or iscsid restart), run the following command on your Linux Server:

```
$ iscsiadm -m node -T <VPSA-Target-IQN> -p <VPSA-Management-IP> --op update -n node.startup -v automatic
```

✓ **Note:** To achieve best performance, it is recommended to use multiple sessions & MPIO. To enable multi-sessions and MPIO, please follow the instructions at: <https://support.zadarastorage.com/hc/en-us/articles/213024386-How-To-setup-Multiple-iSCSI-sessions-and-MultiPath-on-your-Linux-Cloud-Server>

On VMware ESX Servers:

On VMware ESX use the native multipathing. No special configuration is required.

Establishing FC connection

To add Servers (of any OS) connecting to the VPSA over Fibre Channel: Go to [Servers > Add](#) and select Manual:

- Enter the Server Name.

✓ **Note:** Objects names can be up to 128 chars long and can contain letters and digits, dashes “-” and underscores “_”

- Select the server OS
- Turn on the FC connectivity:

Click Edit and select the WWPN of your chosen Server



Before defining a Server using FC connectivity, make sure to install and configure multipathing software on the Server.

For Windows, Linux and ESX servers:

follow the instructions listed above [Adding a Server automatically over Fibre Channel](#).

For Solaris server:

While on Solaris x86 multipath is a default, on SPARC servers it must be configured using:

```
bash-3.2# stmsboot -e
```

✓ **Note:** A reboot is needed after issuing the command.

Multipathing parameters should be set in the following configuration file: `/kernel/drv/scsi_vhci.conf`:

```
#load-balance="round-robin";
load-balance="none";
#
auto-failback="disable";
#
# For enabling MPxIO support for 3rd party symmetric device need an
# entry similar to following in this file. Just replace the "SUN      SENA"
# part with the Vendor ID/Product ID for the device, exactly as reported by
# Inquiry cmd.
#
device-type-scsi-options-list =
"Zadara  VPSA      ", "f_tpgs";
# Tunable for updating path states after a UNIT ATTENTION reset.
# There are arrays which do not queue UAs during resets
# after an implicit failover. For such arrays, we need to
# update the path states after any type of UA resets, since
# UA resets take higher precedence among other UNIT ATTENTION
```

(continues on next page)

(continued from previous page)

```
# conditions. By default, scsi_vhci does not update path states
# on UA resets. To make scsi_vhci do that for such arrays, you need
# to set the tunable scsi-vhci-update-pathstate-on-reset to "yes"
# for the VID/PID combination as described below.
#
#      "012345670123456789012345",      "yes" or "no"
#      "|-VID--||-----PID-----|",
#
scsi-vhci-update-pathstate-on-reset =
    "Zadara VPSA      ",      "yes";
```

For AIX server:

For AIX server to connect to VPSA volumes using multipathing **Veritas Dynamic Multi-Pathing (DMP)** is required. Install DMP on the AIX server.

ODM Package (Zadara.aix.fcp.nonmpio.rte.1.0.0.0.bff) should be installed (to set the storage parameters to the OS) by running the following command:

```
# installp -ad <package_folder> -e <log_folder>/Zadara.aix.fcp.nonmpio.rte.
```

✓ **Note:** After ODM installation you must reboot the AIX server.

Also make sure to setup the zoning on the FC switch to allow connectivity between the connecting Server and the VPSA FC ports.

Connectivity Monitoring

The VPSA can be configured to allow a server connectivity monitoring for specific attached server(s). Connectivity monitoring will ping a specific server and will notify the VPSA administrators (or any other VPSA user with the "Notify on alerts" permissions), notifications are optional (but enabled by default).

Connectivity monitoring can be enabled for a specific server record during server record creation time or any time later by using the "Config" option in the VPSA Servers view. Toggle the option on for "Connectivity monitoring"

Important: Please note that in order to monitor existing records, ICMP traffic from the VPSA to the selected server should be allowed.

The VPSA will attempt to reach to the remote server and in case of a failure that doesn't meet the success threshold (Default 60%) an alert will be sent to the VPSA administrator.

Once enabled, the south pane for a server record will have an additional section presenting the connectivity monitoring status (last state, last success rate, reachability status change and latest connectivity test)

The default configuration for the connectivity monitoring is set globally on the VPSA level and applicable for all server records. For additional information on how to customize the alerting/connectivity success threshold, refer to [Server Connectivity Monitoring](#) in Settings.

7.2.4 Configure Server Attributes

For iSCSI Servers you can change the following Server Attributes using the Config Server dialog:

- Server IQN
- Server IP address
- Enable/Disable IPsec
- CHAP settings

The 'Config server' dialog for an iSCSI server shows the following fields and options:

- Name:** server-215
- File Access:** ON (toggle)
- IP or CIDR Block:** 130.30.2.215
- Enable IPsec:** ☐
- iSCSI:** ON (toggle)
- IQN:** iqn.1993-08.org.debian:01:51843c6c3a9a
- Enable IPsec:** ☐
- CHAP:**
 - VP SA:** User: Liran_Large_Cache1, secret: CmH07LyZg1mg
 - Host:** Mutual CHAP Disabled
- Edit** (link)
- Update** (button)
- Cancel** (button)

Both the server IQN and IP address must be unique. Therefore, the VPSA will block you from changing those attributes to conflicting values used by other Servers.

For FC servers you can change the following Server Attributes using the Config Server dialog:


- WWPN's


The 'Config server' dialog for an FC server shows the following fields and options:


- Name:** SERVER-204
- File Access:** ON (toggle)
- IP or CIDR Block:** 180.80.2.204
- Enable IPsec:** ☐
- iSCSI:** (disabled)
- Fiber Channel:** ON (toggle)
- WWPNs:**
 - 21000024ff7bc175
 - 21000024ff7bc174
- Edit** (link)
- Update** (button)
- Cancel** (button)


7.3 Viewing Servers Properties


The Servers Page displays a list of the available Server objects. You can view the following detailed information in the Servers details South Panel tabs:



Details for SERVER-204



Properties



Volumes


Paths


Metering


Logs


Performance Alerts

ID:	srv-00000001
Name:	SERVER-204 
VPSA CHAP User:	Dima_VPSA5
VPSA CHAP Secret:	gWmfDEXt3u2q
Host CHAP User:	
Host CHAP Secret:	
IP or CIDR Block:	180.80.2.204
iSCSI IQN:	
IPsec iSCSI:	Disabled
IPsec NFS:	Disabled
WWPN 1:	21000024ff7bc174
WWPN 2:	21000024ff7bc175
Registered:	yes
OS:	Microsoft Windows Server 2012 R2 Datacenter 6.3.9600
Added:	2016-03-30 16:48:54
Modified:	2016-03-30 16:48:54

7.3.1 Properties

Each server displays the following properties:

Property	Description
ID	An internally assigned unique ID.
Name	User assigned name. If the Server was created as a result of an iSCSI login, the VPSA will assign it a name similar to its IQN. Name can be modified anytime
Comment	User free text comment. Can be used for labels, reminders or any other purpose
VPSA CHAP User	VPSA CHAP User
VPSA CHAP Secret	VPSA CHAP Secret
Host CHAP User	Host CHAP User
Host CHAP Secret	Host CHAP Secret
IP or CIDR Block	IP Address or CIDR block of the Server(s).
iSCSI IQN	Unique “iSCSI Qualified Name” of the Server.
IPSec iSCSI	Enabled\Disabled
IPSec NFS	Enabled\Disabled
WWPN1	WorldWide Port Name for FC connectivity
WWPN2	WorldWide Port Name for FC connectivity
Registered	Yes – The Connect script was used to create the Server. No – The Server was created manually or via iSCSI login.
OS	OS version detailed string, such as: “Microsoft Windows Server 2008 R2 Datacenter 6.1.7601” Available only for registered Servers.
Added	Date & time when the Server object was added.
Modified	Date & time when the Server object was last modified.

7.3.2 Volumes

A list of all the Volumes attached to this Server.

7.3.3 Paths

This tab lists all the paths between this Server and each controller of the VPSA. If multipathing is set it shows all paths, along with the number of active sessions.

For iSCSI connections the initiator and target IQNs are listed.

For Fibre Channel connections initiator and target WWPN are listed.

7.3.4 Metering

The Metering Charts provide live metering of the IO workload associated with the selected Server.

The charts display the usage data as it was captured in the past 20 “intervals”. An interval length can be set to one of the following: 1 Second, 10 Seconds, 1 Minute, 10 Minutes, or 1 Hour. The Auto button lets you see continuously-updating live metering info (refreshed every 3 seconds).

The following charts are displayed:

Chart	Description
IOPs	The Number of read and write SCSI commands issued from this Server to all its attached Volumes.
Bandwidth (MB/s)	Total Throughput (in MB) of read and write SCSI issued from this Server to all its attached Volumes.
IO Time (ms)	Average response time of all read and write SCSI issued from this Server to all its attached Volumes.

7.3.5 Logs

Displays all event logs associated with this Server.

7.3.6 Performance Alerts

A VP SA administrator has the option to set the following Server Performance Alerts:

- **Read IOPS Limit** – Creates an alert when, during the past minutes, the average read IOPS for this Server exceeds a user-specified threshold.
- **Read Throughput Limit** - Creates an alert when the average read MB/s during the past minute for this server exceeds a user-specified threshold.
- **Read Latency Limit** – Creates an alert when, during the past minute, the average read latency for this server exceeds a user-specified threshold.
- **Write IOPS Limit** – Creates an alert when, during the past hour, the average write IOPS for this server exceeds a user-specified threshold.
- **Write Throughput Limit** - Creates an alert when, during the past minute, the average write MB/s for this server exceeds a user-specified threshold.
- **Write Latency Limit** – Creates an alert when, during the past minute, the average write latency for this server exceeds a user-specified threshold.

7.3.7 Tags

Predefined custom tags can be configured in the **Tags** tab. An example use case for tags is [Filtering the List of Servers](#) in the center pane.

A tag is identified by its **Tag Name** and has a **Tag Value** associated with it. A tag can be defined only once for a server. However, the same **Tag Name** can be defined with a different **Tag Value** for other servers.

- **Create:** To create a new tag for a server, in the pool's **Tags** tab click **Create**, and enter the **Tag Name** and **Tag Value**. The tag is added to the list of tags in the **Tags** tab.
- **Edit:** To change the **Tag Value** of an existing tag, click on that tag in the tags list to mark it, and then click **Edit**. The **Edit Tag** dialog box opens, allowing overwriting of the **Tag Value**.

✓ **Note:** Only the **Tag Value** can be edited. A tag cannot be renamed. It must be deleted, and a tag with the new name configured in its place.

- **Delete:** To delete a tag, click on that tag row in the tags list to mark it, and then click **Delete**. A confirmation dialog box opens.
- **Refresh:** Displays the updated tags list.

7.4 Deleting a Server

In case a server record is no longer needed, it can be Deleted from the VP SA. The server record marked for deletion must not have any volumes attached to it. Once volumes are detached, the server will lose access to the volumes.

In order to delete a server record select the **Delete** option from the top actions menu in the **Servers** section.

CONTROLLERS

Controller Objects in the VPSA represent the clustered virtual controller pair of a VPSA.

8.1 Failover

Failover halts the active controller and activates the standby controller with minimal I/O interruption.

In the Controllers screen, click **Failover** to manually trigger failover.

8.2 Restart File Services

Restart File Services terminates and reactivates SMB or NFS file service processes on the selected controller.

The restart action is almost instantaneous.

Restarting file services is recommended in cases where a folder from the SMB File History Snapshot has been left open on a client, preventing deletion upon expiry, and generating a failure error. In this case, restarting file services releases the client file handle, allowing the deletion to proceed.

To restart file services:

1. In the navigation tree, select **Resources > Controllers**.
2. In the Controllers list, select the controller whose file services will be restarted.
3. In the top menu bar, select Restart File Services.

1. Select the file services to restart.

- **SMB Service**
- **NFS Service**

A confirmation dialog opens, alerting that **Restarting the file share service will disconnect all connected clients** for the selected service (SMB or NFS).

2. Confirm proceeding with the file services restart.

8.3 Viewing Controller Properties

8.3.1 Properties

This tab display the following Properties for the selected Controller:

Property	Description
ID	Controller name (which is automatically assigned)
Target	The iSCSI qualified name (IQN)
WWPN1	Worldwide name of the FC virtual HBA
WWPN2	Worldwide name of the FC virtual HBA
IPSec Key	IPsec Key for secured iSCSI connectivity over IPSec
Encryption Set	
VPSA CHAP User	iSCSI CHAP authentication user name. For use when setting global CHAP on servers
VPSA CHAP Secret	iSCSI CHAP authentication password. For use when setting global CHAP on servers
Cache Size	Size of Flash Cache of this VPSA
Heartbeat1	Heartbeat status between 2 virtual controllers
Heartbeat2	Heartbeat status between 2 virtual controllers
Software Version	Virtual Controller SW build version

8.3.2 Paths

This tab lists all the paths between this virtual controller and the attached servers. If multipathing is set, each server will show all paths along with the number of active sessions.

For iSCSI connections, the initiator and target IQNs are listed.

For Fibre Channel connections, the initiator and target WWPN are listed.



Warning: Fibre Channel - VPSA implements the implicit ALUA format. In case all ports of the underlying Storage Node's HBA (of the active VC) will be disconnected, the VPSA will not initiate a failover.

8.3.3 Virtual Networks

If the VPSA was assigned multiple Virtual Networks (See: [Managing Virtual Networks](#)) they are listed in the Controller's south panel.

Each network is displayed with its IP address and the associated VLAN ID.

8.3.4 System Usage

The System Usage Charts provide live metering of the consumption of compute resources on the selected Controller.

The charts display the usage data as it was captured in the past 20 "intervals". An interval length can be set to one of the following: 1 Second, 10 Seconds, 1 Minute, 10 Minutes, or 1 Hour. The Auto button lets you see continuously-updating live metering info (refreshed every 3 seconds).

System Usage includes the following charts, detailed in the table below:

Chart	Description
CPU Usage (%)	Average usage of all CPU cores
Memory Usage (%)	Used memory out of available controller memory
Bandwidth (KB\s)	Total networking traffic (in KB) of read and write SCSI commands issued to the Controller, per second.
SSD Cache Usage (%)	Amount of Flash Cache used

✓ **Note:** Bandwidth graph in the system usage tab shows cumulative virtual network metering stats across all virtual networks. Traffic per virtual network can be seen in VPSSA performance dashboard.

8.3.5 Cache Metering

The Metering Charts provide live metering of the IO workload associated with Flash Cache of the selected Controller.

The charts display the usage data as it was captured in the past 20 “intervals”. An interval length can be set to one of the following: 1 Second, 10 Seconds, 1 Minute, 10 Minutes, or 1 Hour. The Auto button lets you see continuously-updating live metering info (refreshed every 3 seconds).

Controller metering includes the following charts:

Chart	Description
IOPs	The number of read and write SCSI commands issued to the Flash Cache of the Controller, per second.
Bandwidth (MB\s)	Total throughput (in MB) of read and write SCSI commands issued to the Flash Cache of the Controller, per second.
IO Time (ms)	Average response time of all read and write SCSI commands issued to the Flash Cache of the Controller, per selected interval.
Hit Rate (%)	Read and write cache hit-rate during the selected interval.

8.3.6 Logs

Displays all event logs associated with the selected Controller.

8.3.7 Performance Alerts

The **Performance Alerts** tab lists the configurable alerts of the selected Controller:

- Average CPU Usage in the last minute is above the given threshold
- Average memory consumption in the last minute is above the given threshold

REMOTE VPSA

VPSA Asynchronous Remote Mirroring provides the ability to replicate your VPSA's data asynchronously to a different Pool within the same VPSA, to a different VPSA (either locally within the same Zadara Cloud, or remotely to a VPSA located in a remote region), or even to a different cloud provider. You can replicate a single source Volume to any number of remote (or local) Mirrors.

Asynchronous Mirroring has minimal impact on IO throughput and response time from the Server perspective since the Server IO returns immediately after being written to the local VPSA storage (without waiting for acknowledgment from the remote VPSA, like is required with Synchronous Mirroring). Later, the data is synchronized to the Remote VPSA in the background.

Remote VPSA communication is strongly authenticated and secured using [cryptographic protocols](#) designed to provide secure communication over the Internet. Mirrored data is encrypted before being shipped to the remote VPSA.

Mirrored data is also compressed before being shipped to a remote VPSA in a different region, for efficient bandwidth utilization.

You can establish a “many-to-many” remote mirroring relationship for different Volumes between different VPSAs. This means that a VPSA can mirror Volumes to many remote VPSAs, while at the same time also be the Destination VPSA for other Volumes in any other VPSA.

✓ **Note:** Remote VPSA communication is done over ports 1339/1340(TCP).


9.1 Connect to a remote VPSA

The first step to building a DR plan (i.e. setting up a Mirrored Volume on a remote VPSA) is establishing a trusted relationship between your VPSAs.

If the VPSAs are located in different Zadara Storage Clouds you will need to first assign a Public IP to each VPSA (See [Assigning Public IPs](#) for more details) alternatively.

If the remote VPSA is located in the same Zadara cloud replication can take place using the following:

1. Frontend network - assuming both VPSAs are sharing the same network or on different network where proper route is in place.
2. Using a VPSA VNI (Virtual Network Interface - additional local network interface allocated to the VPSA)

 **Warning:** Limitation - Replication using VNI (both source and target) over different network is allowed only if a single VNI is attached to the VPSA.

Navigate to the **Remote VPSAs** page and click the **Discover** button.

Enter the following details:

- **Remote VPSA IP Address:**
 - If the remote VPSA is located in a different Zadara Storage Cloud in a remote Region:
 - * Enter the remote VPSA Public IP address. You can find it in the VPSA details in the Management console or in the remote VPSA GUI, under [Controllers > Public IP](#).
 - * Select the “Discover remote VPSA through Public IP” checkbox.
 - If the other VPSA is located within the same Zadara Storage Cloud:
 - * Enter the remote VPSA Management IP address.
 - * In this case, do NOT check the “Discover remote VPSA through Public IP” checkbox.
- **Username & Password** – For authentication against the remote VPSA you are required to enter the username and password of a valid user in the remote VPSA. A cryptographic hash value (using a one-way SHA-1 hash function) of the entered password is sent to the remote VPSA.

9.2 Remote VPSA Properties


You can review all of the remote VPSAs with which this VPSA has established a trusted relationship. For each VPSA the following details are provided:

Properties

- **Local ID** – The VPSA ID of the Local VPSA.
- **Remote ID** – The VPSA ID of the remote VPSA.
- **Name** – The name of the remote VPSA.
- **Provider** – The name of the Cloud Provider where the remote VPSA is located.
- **Software Version.**
- **IP** – Public or Management IP through which the VPSAs are connected.
- **Rate Limit (MB/s)** – Maximum transfer rate allowed for mirroring data to the remote VPSA.

Pools

- Each VPSA publishes the list of Pools that can be used to provision the remote Volume.

 **Note:** This list does not update automatically. Click the Refresh button to update the remote Pools info from the remote VPSA.

Logs

- The logs related to the selected remote VPSA

REMOTE OBJECT STORAGE

Zadara VPSA provides built in backup and restore capabilities to Zadara Object Storage, AWS S3, Google Cloud Storage, Azure Blob Storage or any other S3 compatible object storage. The backup process involves transporting VPSA Snapshots to the remote Object Storage for safe keeping.

Backup to Object Storage (B2OS) allows you to store a backup of the VPSA volume on Object Storage and later restore it to its original VPSA or to **any** other VPSA in a different location with access to the same object storage bucket.

10.1 Connecting to Remote Object Storage

In order to back up your data to Object Storage you need to connect the VPSA to the Object Storage bucket (container). To do this you will need the following information:

- Bucket/Container name
- Access key ID
- Secret access key



Note:

- In order to keep the data backed up ready for restore, the remote Object Storage bucket must not have any life-cycle policy (such as archiving to Glacier) as all backup objects are required for immediate restore.
- For AWS-S3 the minimal S3 permissions required for the remote Object Storage bucket keys:
 - GetLifecycleConfiguration
 - GetObject
 - PutObject
 - List*
 - DeleteObject

Since public object storage, such as AWS S3, is on a public network and your VPSA is within your private cloud or local network, there are 2 options:

- Connect via a public IP address (see [Assigning Public IPs](#) for assigning a public IP address)
- Connect via a proxy server in your VPC that has access to the Internet

To connect to Remote Object Storage:

1. Go to **VPSA GUI > Remote Object Storage** and click **Connect**.

2. Select between Zadara Object Storage, AWS S3, Google Cloud Storage, Azure Blob Storage or Custom (S3 Compatible Object Storage).
3. Enter the bucket/container name, access key and secret key.
4. Select the connection method – via public IP, or the local management network.
5. If needed, set-up a proxy server and provide the proxy IP address and port, as well as login credentials.

✓ **Note:** For details about setting up the proxy server see this article: [Setup Backup To S3 \(B2S3\) Through a Proxy In Your AWS VPC](#)

If the target Object Storage type is AWS S3, the following options are available:

- Region - the target bucket AWS region (mandatory)
 - Ignore Lifecycle Policies - Could be checked in case Lifecycle cannot be disabled on the target bucket. (not recommended)
 - Use KMS Key ID - default KMS managed private key ID to be used for SSE (Server-Side Encryption). (optional)
6. Press **Submit**.

10.2 Viewing Remote Object Storage properties

The Remote Object Storages details are shown in the following South Panel tabs:

Properties

Each Remote Object Storage includes the following properties:

Property	Description
ID	An internally assigned unique ID
Type	AWS S3, Google Cloud Storage, VP SA Object Storage or Custom
Endpoint	Location (region) of the object storage
Connect Via	The network used for the backup data transfer (Public IP or Management Network)
Bucket	The name of the S3 bucket used to store the backup data
Proxy IP	IP address of the proxy server
Proxy Port	Port used for the proxy connection (typically 3128)
KMS Key	(AWS S3) The KMS Key ID used for SSE
Allow Lifecycle Policies	Whether Lifecycle Policies are ignored for the target Bucket

Backup Jobs Tab - List of all backup jobs using the selected Remote Object Storage

Restore Jobs Tab - List of all restore jobs using the selected Remote Object Storage

Logs Tab - List of event log messages related to that Remote Object Storage

SNAPSHOTS AND SNAPSHOT POLICIES

11.1 Managing Snapshots and Snapshot Policies

Snapshots are Read-Only representations of the Volume's data set at a given point-in-time. Snapshots are very efficiently thinly provisioned, sharing all the unmodified data chunks with the Volume. Write ordering is ensured at Snapshot creation, i.e. all writes that were acknowledged to the Server by the VPSA before the Snapshot was created will be contained in the Snapshot's data set.

11.1.1 Manual creation and deletion of snapshots

To manually create a snapshot:

- Go to the Volumes page, press the **Data Services** button and select **Create Snapshot**.
- Enter a unique name for the snapshot and confirm the operation.

To manually delete a snapshot:

- Go to the Volumes page, select the **Volume** and view the **Snapshots** South Panel tab to display the list of snapshots associated with this Volume.
- Select the snapshot to be deleted in the **Snapshots** tab and press the **Delete Snapshot** button at the top left corner of the South Panel.
- The snapshot will move to a Deleting state and will disappear from the list once the deletion process completes. Please note that Snapshots deletion typically takes less than a minute, but in complex configurations it may extend up to few minutes.



Note: You can not manually delete snapshots related to the volume mirrors. See [Mirroring](#) for details.

11.1.2 Managing Snapshot Policies

Snapshot policies define the Snapshots life cycle via the enforcement of creation and deletion policies. Snapshot Policies are “global” entities, and you can apply instances of the policies to one or more Volumes. Unapplied policies are idle—they do not consume any resources and never create any snapshots. A few points to consider:

- You can apply a Snapshot policy to one or more Volumes.
- You can apply multiple Snapshot Policies to a Volume.

- If two or more Snapshot policies are scheduled to create a Snapshot at the same time on the same Volume, only a single Snapshot will be created. That Snapshot will only be deleted when all relevant Delete Policies approve its deletion.
- Snapshot creation time is a “rounded” time, regardless of the precise policy creation time. For example, if you initialized a Snapshot Policy at 9:02 that has a Creation Policy to create a snapshot every 10 minutes, the Snapshots will be created at 9:10, 9:20, 9:30 and so forth (not at 9:12, 9:22, 9:32, etc.).
- For the predefined snapshots policies like “Every Day” or “Every Hour” the Snapshot creation time is distributed on 10 minutes slots during the hour. The specified interval of one hour is kept, but not necessarily on the hour. Snapshots may be taken every hour 10 minutes after the hour, or 20 minutes after the hour, etc... (For example: 9:10, 10:10, 11:10 , ...) If a precise snapshot creation time is needed, define a custom snapshot policy that specifies the exact time.
- You can decide whether or not empty snapshots are to be created. i.e. if the time has come to create a Snapshot according to the Creation Policy but no data has changed since the previous Snapshot, you can specify whether a new and empty Snapshot will be created. This might be useful if you want to make sure the snapshot policy is enforced and snapshots are taken on time regardless of the data changes.
- The following Snapshots Policies are predefined in the VPSSA.

Name	Create Policy
Hourly Snapshots for a Day	Every hour
Daily Snapshots for a Week	Once per day after midnight
Weekly Snapshots for a Year	Every Sunday after midnight
Daily Backup for a Year	Once per day after midnight

Creating a new Snapshot Policy

1. Go to the Snapshot Policies page and press **Create**.
2. In the **Create Snapshot Policy** dialog, enter:
 - **Name** - Provide a meaningful name to the Policy.

✓ **Note:** Objects names can be up to 128 chars long and can contain letters and digits, dashes “-” and underscores “_”

- **Creation Policy** – Select the appropriate policy from the drop down list.
- **Deletion Policy** – Use these 2 fields to define the maximum number of Snapshots to retain in the Deletion Policy. If you will be using this policy for Remote Mirroring, you can define a different number of Snapshots to retain on the DR site. This field is optional and defaults to the above deletion policy.
- **Allow Empty Snapshot Creation** – Select this checkbox if you’d like Snapshots to be created according to the Creation Policy, even if no data was modified since the previous Snapshot.
- **Set as default policy for newly created volumes** – Select this checkbox if you’d like all new Volumes to default to this Snapshot Policy. Select the appropriate Creation Policy from the drop down list.
- Define the number of Snapshots to retain in the deletion policy.
- Allows Empty Snapshot Creation – Set this checkbox if you’d like snapshots to be created according to the creation policy even if no data was modified since the previous snapshot.
- If you will be using this policy for Remote Mirroring, you can define a different number of Snapshots to retain on the DR site. This field is optional and defaults to the above deletion policy.

Editing a Snapshot Policy

1. Go to the Snapshot Policies page, select the Policy and press **Edit**.
2. You can edit all of the Snapshot Policy's attributes: Name, Creation Policy, Deletion Policy, Allow Empty Snapshots Creation and Set as Default Policy.
3. You can modify a Snapshot Policy even when it is active on one or more Volumes. The modifications in the Policy's behavior will be reflected on all relevant Volumes.
4. If you reduce the number of Snapshots to retain for a Snapshot Policy that is active on one or more Volumes, it will trigger the deletion of all Snapshots that no longer meet the new Deletion Policy.

Applying a Snapshot Policy on a Volume

1. Go to the Volumes page, select the Volume and select **Data Services > Attach Snapshot Policy**.
2. In the **Attach Snapshot Policy to Volume** dialog, select the Snapshot Policy to apply to the selected Volume, and press **Submit**.

Detaching a Snapshot Policy from a Volume

1. Go to the Volumes page, select the Volume and press the Snapshot Policies south tab to view the Volume's applied Snapshot Policies.
2. Select the Snapshot Policy to delete and press **Detach Policy** at the top left corner of the South Panel.

You will be prompted to decide whether or not to delete all the Volume's Snapshots which are associated with this Policy.

Pausing or Resuming a Snapshot Policy

You can pause an active Volume Snapshot Policy. New Snapshots will not be created, but existing Snapshots are not affected. Pausing a Snapshot Policy on one Volume has no impact on other Volumes that also have this Policy active.

- To **pause** a Snapshot Policy:

1. Go to the :Volumes page, select the Volume and press the **Snapshot Policies** tab on the South Panel to view the Volume's active Snapshot Policies.
2. Select the Snapshot Policy and press **Pause Policy** at the top left corner of the South Panel.

The **Policy Status** will change to **Paused**.

- To **resume** a Policy:

✓ **Note:** The **Pause / Resume** button toggles according to the current **Policy Status**.

Select a Policy in a **Paused** state and press **Resume Policy** at the top left of the South Panel.

The **Policy Status** will change to **Active**.

11.2 Filtering Snapshots

Snapshots can be created manually, by using Snapshot Policies, by Remote Mirroring or by Backup to Object Store. This can result in many Snapshots spread across multiple Volumes.

Finding a specific snapshot could therefore take some time. The “Filter Snapshot” option will help you to find the snapshot you need more efficiently.

1. Go to The Volumes page, select a Volume and display the **Snapshots** tab in the South Panel.
2. Press the **Filter** button at the bottom of the page.
3. In the resulting dialog, define one or more of the following parameters:
 - You can define the **From Date/Time** and **To Date/Time** to filter only Snapshots that were created during that interval.
 - You can select the Origin of the Snapshot:
 - **All** – all Snapshots origins.
 - **User** – Snapshot created manually or via a Snapshot Policy which was attached to this Volume.
 - **Mirror** – Snapshots that were created by the Remote Mirroring application (using the Snapshot policy which was defined at the time of the Mirror creation).
 - **Object Storage** – Snapshots that were created by the Backup to Object Store (using the Snapshot policy that was defined at the time of the Backup definition).
 - **Snapshot Policy** – Select a Policy if you’d like to filter only Snapshots that were created by that specific Policy.

MIRRORING

VP SA Asynchronous Remote Mirroring provides the ability to replicate your VP SA's data asynchronously to a different Pool within the same VP SA, to a different VP SA (either locally within the same Zadara Cloud, or remotely to a VP SA located in a remote region), or even to a different cloud provider. You can replicate a single source Volume to any number of remote (or local) Mirrors.

Asynchronous Mirroring has minimal impact on IO throughput and response time from the Server perspective since the Server IO returns immediately after being written to the local VP SA storage (without waiting for acknowledgment from the remote VP SA, like is required with Synchronous Mirroring). Later, the data is synchronized to the Remote VP SA in the background.

The VP SA Remote Mirroring is snapshot-based, meaning that only modified data chunks between two point-in-time Snapshots are synchronized. This has some major advantages:

- If a file/block was modified several times between two consecutive snapshots only the last change will be synchronized, thus saving bandwidth.
- Snapshots are crash-consistent, thus at the Remote Site you always have a crash-consistent point-in-time data set of your application.
- You can easily create many Read/Write Clones of your remote data at various point-in-time snapshots for Test & Dev.

The VP SA manages checkpoints to track the sync progress within a Volume/Snapshot. In case of a transport failure (line failure, VP SA failure etc.) the VP SA has a clear checkpoint from where to resume the sync.

Mirrored data is also compressed before being shipped to a remote VP SA in a different region, for efficient bandwidth utilization.

You can establish a “many-to-many” remote mirroring relationship for different Volumes between different VP SAs. This means that a VP SA can mirror Volumes to many remote VP SAs, while at the same time also be the Destination VP SA for other Volumes in any other VP SA.

12.1 Creating a Local Mirror (on the same VP SA)

To create a Local Mirror, go to the [Mirroring](#) page and click the Create button. Give the new Mirror a name, as well as a name for the new Volume it creates. Select the destination Volume you will be mirroring and then click the Next button.

Create Mirror

Mirror Name: *

Destination Volume Name: *

Volume: *

Name	Capacity	Status	Data Type
VOL1	100 GB	Available	BLOCK

« < | Page 1 of 1 | > » | ↻ Displaying 1 - 1 of 1

On the next screen select your local VP SA as the Destination VP SA and the Pool you want to replicate to. You cannot replicate to the same Pool in which the source Volume resides. Then click the Next button.

Create Mirror

Destination VP SA: *

Destination Pool: *

Name	Total Capacity	Free Capacity	Version
PL1	542 GB	542 GB	2

On the final screen select the Snapshot Policies, you want to apply to the Mirrored Volume, and click Submit.

Create Mirror

Mirroring Policies: *

<input type="checkbox"/>	Name	Create Policy	Delete Policy	Dst. Delete Policy
<input checked="" type="checkbox"/>	Hourly Snapsh...	Every hour	Keep latest 24 ...	Keep latest 24 ...
<input checked="" type="checkbox"/>	Daily Snapshot...	Once per day a...	Keep latest 7 s...	Keep latest 7 s...
<input type="checkbox"/>	Weekly Snapsh...	Every Sunday ...	Keep latest 53 ...	Keep latest 53 ...

12.2 Creating a Remote Mirror

You can create the Remote Mirror from the [Remote Mirroring](#) page by clicking Create. You will see a similar dialog:

Name	Capacity	Status	Data Type
VOL1	100 GB	Available	BLOCK

Name the Mirror and the destination Volume, select the source Volume and click Next.

Hint: Backup job name, similarly to other VP SA entities is limited to:

- 128 characters in length
- ASCII characters only (between 32 and 126) with the exception of the following special characters: ", %, ;, {, }, [,], <, >, \, ', &
- Cannot have a starting or trailing whitespace

Name	Total Capacity	Free Capacity	Version
RAID-10-Pool-1	3.47 TB	3.47 TB	3

Choose the Destination VPSA, WAN Optimization or I/O Performance Optimization (see below) and the destination Pool. Click Next.

✓ **Note:** Mirroring a volume from a VP SA in a software version which is 19.08 and above to a target VP SA with a software version lower than 19.08 is not supported.

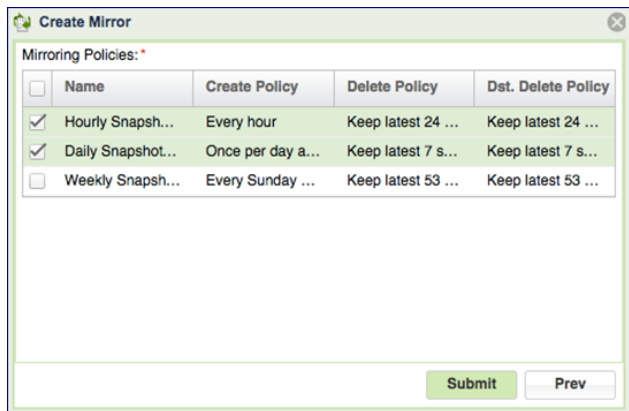
WAN Optimization v. I/O Performance Optimization

The VP SA supports selection between “I/O” and “WAN” data synchronization optimization. When “I/O” is selected, the VP SA synchronizes modified Pool chunks of 25KB or 1MB, depending on the Pool type. When “WAN” is selected, the VP SA

synchronizes only modified 4KB sub-chunks. WAN optimization typically reduces WAN traffic bandwidth at the expense of additional workload on the source Volume, which may affect application performance.

VPSA Flash Array Dedup and Compression

While mirroring to All Flash remote VPSA, you can decide if the mirrored Volume will be dedup'ed and/or compressed, in order to save space on the mirror destination.



Select the snapshot policies you want to apply to the remote Mirrored Volume and click Submit.


12.3 Replicate the same Volume to multiple destinations

It is possible to replicate the same Volume to multiple destinations. Just repeat the above steps, selecting a different destination each time.

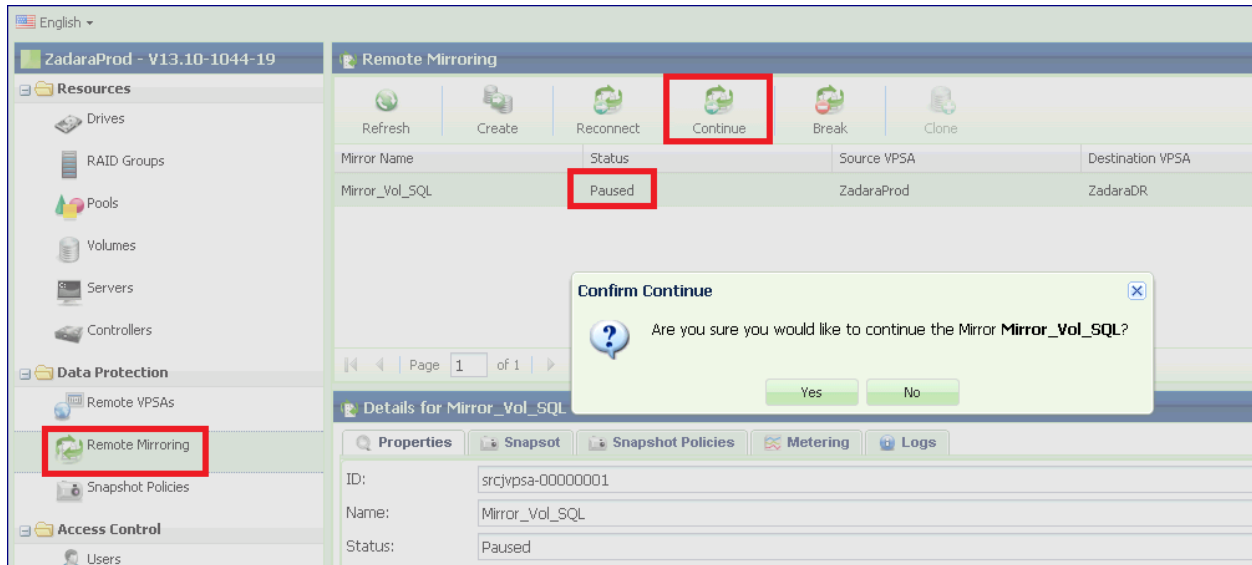
12.4 Pause & Continue Remote Mirror

It is possible to pause a Remote Mirror. A paused Mirror will stop syncing data immediately and stop creating new Snapshots. The status of the Mirror will change to “Paused”.

To pause a Mirror, select the Mirror in the **Mirroring** page and press the Pause button.



To resume the Mirror operation, select the Mirror in the **Mirroring** page and press the Continue button.



12.5 Managing Mirror Lifecycle

The Mirror controls the Remote Volume on the destination VPSA. As long as the Mirror is active, the target volume cannot be attached to any Server, nor can it be modified outside the scope of the Mirror. Hence it is treated as a special kind of "Destination Volume."


You can view Mirror Destination Volumes on the "Dest. Volumes" tab on the [Pools](#) Page of the Destination VPSA, but they do not appear in the [Volumes](#) page.

12.5.1 Clone Destination Volume for Dev & Test of Remote Mirror

For offline processing (e.g. Dev & Test and other purposes) you can Clone the destination Volume using the data set of any Snapshot that was completely synchronized. You cannot create a Clone of the Snapshot that is currently being synchronized.

The Cloned Volume is independent of the Destination Volume or the Mirror (i.e. you can delete both the Destination Volume and the Mirror and the Cloned Volume will not be affected).

To Clone a Mirrored Destination Volume, go to the [Mirroring](#) page on the **Destination VPSA**. Select a Mirror and click the Clone button.




- Select the point-in-time Snapshot that contains the data set that you wish to clone. The VPSA will assign a name to the Cloned Volume which is a concatenation of the Dest Mirror Volume name and the timestamp of the selected Snapshot. You can modify this name at any time.
- You can find the newly created Volume in the [Volumes](#) Page.

12.5.2 Breaking a Mirror

Breaking a Mirror is the process of deleting the Mirroring relationship between the Source Volume and the Destination Volume, while leaving sufficient information for to reconnect the Mirror in future. The Destination Volume then becomes a “regular” Volume and the source and the destination Volumes are now independent of each other. A Mirror can be broken from the source or from the destination VPSAs.

You can perform a future Mirror reconnect in both directions. However, there are implications on the data which is retained and data which will be overwritten depending on which side the Mirror reconnect is initiated from. More details on this are in the next section.

To break a Mirror, go to the [Mirroring](#) page, select a Mirror and click the Break button. After confirming the operation, the Mirror Object in the [Mirroring](#) page will disappear from both the Source and the Destination VPSAs.



While braking a mirror you have the ability to assign the destination volume with the same world wide ID(WWID) as the source volume. WWID is used by some host platforms for volume identification and therefore assigning the source WWID

to a target volume might accelerate Disaster recovery procedure in cases where host environments is the main and DR sites have the same volume metadata. To preserve source WWID for the target mirror volume check the Assign source WWID to target volume box on the mirror brake dialog.

✓ **Note:** To avoid data availability and integrity issues a host should not be exposed to two volumes with the same WWID.

12.5.3 Reconnecting a Mirror

As previously described, the VP SA retains sufficient metadata about each Volume after a Mirror has been broken to enable a future reconnect of the Mirror relationship. Also, all the mirroring snapshot policies (Type = "Remote Mirroring") are kept in place, but in "Paused" State. as shown in the image below. These snapshots policies are used for reconnecting the mirror.

Name	Status	Type	Create Policy	Delete Policy	Dest. Delete Policy
Daily Snaps for DP	Active	Object Store Backup/Restore	Once per day at midnight	Keep latest 14 snapshots	Keep latest 90 snapshots
10 Min Snaps for DR	Active	Remote Mirroring (NAS_SATA_VOL2_MIRR...)	Every 10 minutes	Keep latest 156 snapshots	Keep latest 288 snapshots
10 Min Snaps for DR@...	Paused	Remote Mirroring (NAS_SATA_VOL2_MIGR...)	Every 10 minutes	Keep latest 288 snapshots	Keep latest 156 snapshots
ON-DEMAND@...	Paused	Remote Mirroring (NAS_SATA_VOL2_MIGR...)	On Demand	-	-
Temporary Snaps S3 Catchup	Paused	Remote Mirroring (NAS_SATA_VOL2_MIRR...)	Once per day at midnight	Keep latest 45 snapshots	Keep latest 10 snapshots

This metadata allows the VP SA to identify a remote Volume, on a Remote VP SA, that had a Mirroring relationship with a local Volume anytime in the past and find the most recent Snapshot that is in-sync on both Volumes. This enables it to reconnect the Mirror relationship and resume the sync process from the most recently updated data set. If there is a match between policies on the source and destination the matching snapshot policies will be used to reconnect the mirror.

A mirror reconnect can be done in any direction, regardless of the previous Mirror direction. This provides the required flexibility for a DR plan. In case of a suspected source site disaster, you can break the Mirror, assign the Destination Volume to an application server and work on the DR site. Once the source site is back, you can decide in which direction to resume the mirroring relationship.

⚠ Caution: When resuming Mirroring, the VP SA identifies the most recent point-in-time Snapshot that is completely in-sync on both source and destination Volumes. Any data that was written on the destination Volume after this snapshot will be deleted!

To Reconnect a Mirror:

Mirroring					
Refresh	Create	Reconnect	Pause	Break	Clone
Mirror Name	Status	Source VPSA	Destination VPSA	RPO	
VOLUME10_MIRROR	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	
SRM_MIRROR_cg-0000005b	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	
SRM_MIRROR_cg-0000005a	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	
SRM_MIRROR_cg-0000005d	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	
SRM_MIRROR_cg-00000061	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	
SRM_MIRROR_cg-00000059	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	
SRM_MIRROR_cg-00000070	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	
SRM_MIRROR_cg-00000069	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	
SRM_MIRROR_cg-0000005e	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	
SRM_MIRROR_cg-00000060	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	
SRM_MIRROR_cg-0000005c	Sync complete	VPSA2	VPSA1	2016-04-24T00:00:27+0300	

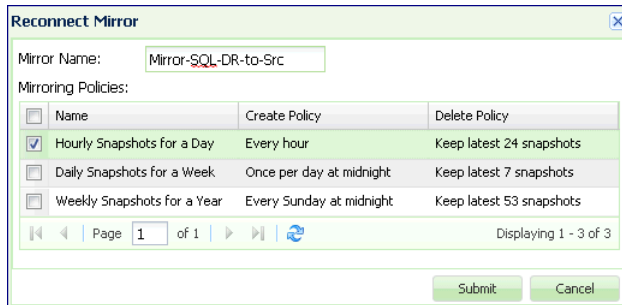
- Go to the **Mirroring** Page and click Reconnect. The system will list candidate volumes with broken mirror. Select the Volume that you wish to act as the Source Volume of the Mirror.
- Select the Remote VPSA that contains a Volume that used to be a Mirror pair of the selected Source Volume in the past. Press Next.
- The VPSA will query the remote VPSA and display suggested Remote Volumes which can be Destination Volumes of the Mirror, with the following information:
 - Remote Volume name.**
 - New Data** – There is new data on the Remote Volume which was written after the last sync point and which needs to be deleted in order to reconnect the Mirror.
 - Last Sync** – The timestamp of the most recent Snapshot. Any data written on the Source Volume after that timestamp will be synchronized to the Remote Volume.
 - Snaps to Del** – Number of snapshots to delete on the Remote Volume. Please note that it is possible that empty Snaps need to be deleted while no new data is lost on the Remote Volume.

Reconnect Mirror			
Suggested Remote Volumes:			
Remote Volume	New Data	Last Sync	Snaps to Del.
vol-5QL	No	2013-11-24 09:20:00	1
Page 1 of 1 Displaying 1 - 1 of 1			
Continue Cancel			

- Press Continue.
- Enter a name for the new Mirror.

✓ **Note:** Objects names can be up to 128 chars long and can contain letters and digits, dashes “-” and underscores “_”

- If the system finds matching policies on the source and destination VPSA’s they are automatically used. If no matching policy can be found, the following dialog is displayed, asking for the Snapshot Policy to be used. Select Snapshot Policy for the new Mirror.
- Press Submit to Reconnect the Mirror.



Reconnect Mirror

Mirror Name:

Mirroring Policies:

<input type="checkbox"/> Name	Create Policy	Delete Policy
<input checked="" type="checkbox"/> Hourly Snapshots for a Day	Every hour	Keep latest 24 snapshots
<input type="checkbox"/> Daily Snapshots for a Week	Once per day at midnight	Keep latest 7 snapshots
<input type="checkbox"/> Weekly Snapshots for a Year	Every Sunday at midnight	Keep latest 53 snapshots

Page 1 of 1 | Displaying 1 - 3 of 3

✓ **Note:** Reconnect Mirror is blocked in the following cases:

- The Destination Volume is attached to a Server
- The Destination Volume has active Snapshot Policies. Remote Mirror does not allow to have any non-remote-mirror snapshot policy attached to the destination volume. You have to detach all non-remote-mirror snapshot policies on the destination volume before reconnecting the mirror (this situation is common when reversing the direction).

12.6 Viewing Remote Mirror Properties

The [Remote Mirroring](#) Page displays the list of Remote Mirrors that the VPSA participates in, either as the Source or the Destination. The Mirrors are not symmetric, so both the source and the destination VPSAs display slightly different info.

Select a Mirror and review the detailed information in the following South Panel tabs:

VPSA2 - 00.00-367

Resources

- Drives
- RAID Groups
- Pools
- Volumes
- Servers
- Controllers

Remote Storage

- Remote VPSAs
- Remote Object Storage

Data Protection

- Snapshot Policies
- Mirroring**
- Backup to Object Storage
- Restore from Object Storage

Container Service

- Images
- Container Groups
- Containers

NAS Access Control

System

- Settings
- User Management
- Logs
- Support

Mirroring

Refresh Create Reconnect Pause Break Clone

Mirror Name	Status	Source VPSA	Destination VPSA	RPO
VOLUME10_MIRROR	Idle	VPSA2	VPSA1	2016-04-2
SRM_MIRROR_cg-0000005b	Idle	VPSA2	VPSA1	2016-04-2
SRM_MIRROR_cg-0000005a	Idle	VPSA2	VPSA1	2016-04-2

Page 1 of 2

Details for SRM_MIRROR_cg-0000005b

Properties Snapshots Snapshot Policies Metering Logs

ID: srcjvpsa-0000008f

Name: SRM_MIRROR_cg-0000005b

Status: Idle

Current Snapshot:

RPO: 2016-04-25T00:00:18+0300

Optimization: I/O Performance

Created: 2016-04-11 18:30:44

Modified: 2016-04-25 00:00:19

Source VPSA

Provider: zadaraga6

VPSA: VPSA2

Pool: pool1

Volume: VOLUME3

Destination VPSA

Provider: zadaraga6

VPSA: VPSA1

Pool: pool1

Volume: VOLUME3_MIRROR

Properties

Each Mirror includes the following properties:

Property	Description
ID	An internally assigned unique ID.
Name	User assigned name. Can be modified anytime.
Comment	User free text comment. Can be used for labels, reminders or any other purpose
Status	<ul style="list-style-type: none"> Idle – Mirror has nothing to Sync. Failed Paused Syncing (X%) – Transferring the modified data of the “Current Snapshot” to the remote Volume. X% stands for the syncing location inside the Snapshot. Ready to sync/Syncing (X%) – Same as “Syncing” but at the destination VPSA.
Current Snapshot	Empty. Snapshots are listed in the Snapshots tab
RPO	Return Point Objective – This is the timestamp of the most recent fully synchronized Snapshot.
Rate Limit	Maximum transfer rate allowed for mirroring data to the remote VPSA.
Optimization	I/O Performance or WAN optimization
Created	Date & time when the Mirror was created.
Modified	Date & time when the Mirror was last modified.
Source VPSA / Provider	The name of the Cloud Provider where the source VPSA resides.
Source VPSA / VPSA	Source VPSA name.
Source VPSA / Pool	Pool name where the Source Volume is provisioned. This parameter is available only at the source VPSA.
Source VPSA / Volume	Source Mirror Volume name.
Destination VPSA / Provider	The name of the Cloud Provider where the destination VPSA resides.
Destination VPSA / VPSA	Destination VPSA name.
Destination VPSA / Provider	Pool name where the destination Volume is provisioned.
Destination VPSA / Pool	Destination Mirror Volume name.

Snapshots

The Snapshots tab for Mirroring, lists the point-in-time Snapshots of the Mirror on this VPSA. Please note that a Mirror configuration supports retaining different numbers of Snapshots on the Source and the Destination VPSAs. Each VPSA will display its own managed list. If you retain many Snapshots, you may want to use the Snapshot Filtering tool to find a specific Snapshot. For more details see [Filtering Snapshots](#).

Snapshot Policies

The Snapshot Policies tab for Mirroring lists the active Snapshot Policies used by this Mirror to manage Snapshots on the Source Volume and the Destination Volume of the Mirror. The snapshot policies only appear on the Source VPSA, not on the Destination VPSA Mirror Snapshot Policies tab.

Details for Mirror_Vol_SQL				
Properties	Snapshot	Snapshot Policies	Metering	Logs
Name	Status	Create Policy	Delete Policy	Dest. Delete Policy
yearly Weekend snapshots	Active	Weekend days at midnight	Keep latest 104 snapshots	Keep latest 104 snapshots
every 5 minutes for an hour	Active	Every 5 minutes	Keep latest 12 snapshots	Keep latest 12 snapshots

The Source VPSA manages the Mirror Snapshot Policies, therefore modifications to the Mirror Snapshot Policies are con-

figured on the Source VP SA.

The Source VP SA updates the Destination VP SA regarding any changes to the Dest Delete Policy.

You may make modifications while the Policy is active on a Mirror and the changes become effective immediately. For example, if you change the policy to retain fewer Snapshots, some older Snapshots will be deleted immediately.

The following information is provided per Snapshot Policy on the **Source VP SA**:

- **Status** – Current state of the mirror.
- **Create Policy** – Minimum time between Snapshots.
- **Delete Policy** – How many Snapshots are retained on Source Volume.
- **Dest Delete Policy** – How many Snapshots are retained on the Destination Volume.

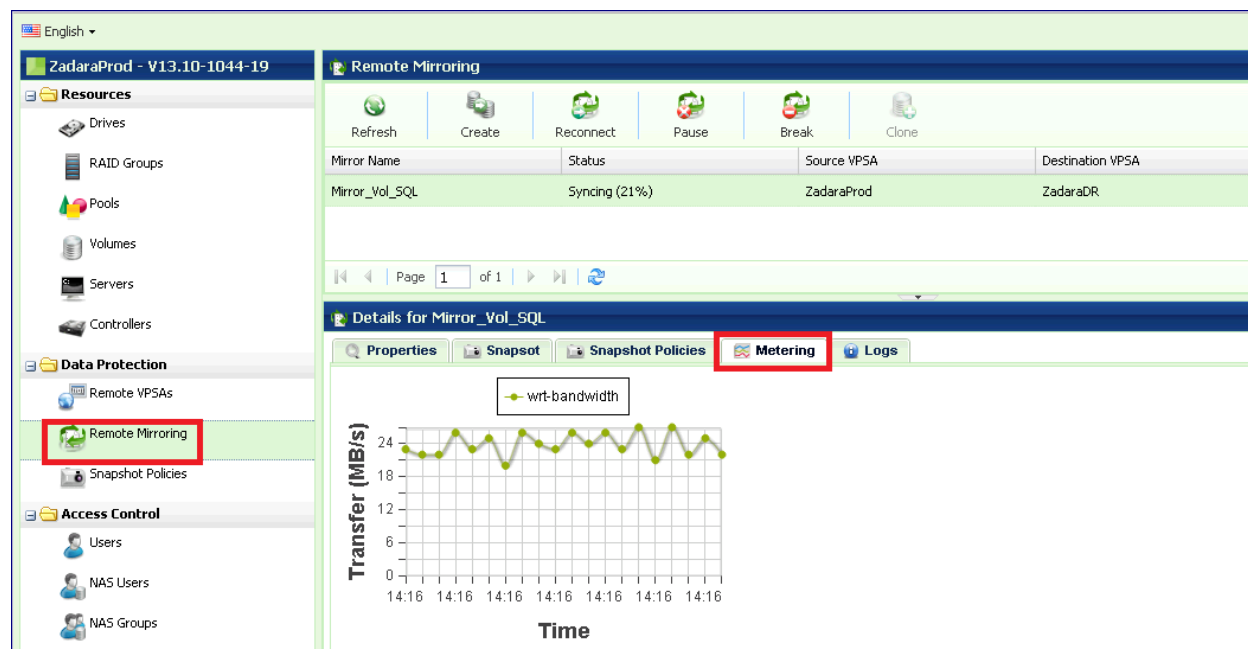
The following information is provided per Snapshot Policy on the **Destination VP SA**:

- **Status** – Current state of the mirror.
- **Create Policy** – N/A.
- **Delete Policy** – How many Snapshots are retained on the Destination Volume. This value is identical to the “Dest Delete Policy” on the Source VP SA.
- **Dest Delete Policy** – N/A.

Metering

The Mirror Metering tab provides live information of the Mirror’s transfer throughput and IO Time associated with the selected Mirror. You can view the Mirror metering information on either the Source or the Destination VP SA.

The charts display the metering data as it was captured in the past 20 “intervals”. An interval length can be set to one of the following: 1 Second, 1 Minute, 10 Minutes, or 1 Hour. The Auto button lets you see continuously updating live-metering information (refreshed every 3 seconds).



Logs

Displays all event logs associated with this Mirror relationship.

REMOTE CLONES

13.1 Remote cloning introduction

The Remote Clone feature makes a specified snapshot of a source volume instantly available as a volume on another VPSA. The snapshot is available before data is copied, in the same cloud or in a different cloud, and over any distance. The cloned new volume is immediately available, in contrast to mirroring that might take a long time to replicate data, which depends on capacity and link bandwidth.

13.2 Remote cloning modes

The remote cloning process can run in one of two modes:

13.2.1 On-demand mode

The cloned volume retrieves blocks of data from the source volume on demand, as needed. During that time, the clone behaves like any other volume, but is dependent on its source volume. After all the data is copied to the clone, the relation between the source and cloned volumes breaks, and the cloned volume becomes a regular volume.

13.2.2 Background mode

The system retrieves all the data from the source volume in the background. In this mode, after all the data is retrieved, the connection to the original snapshot breaks and the volume becomes independent.

13.3 Common remote cloning use cases

The most common remote cloning use cases are:

13.3.1 Instant mobility

Instant mobility is the rapid migration of volumes between VPSAs with minimal downtime for the application accessing the migrating volume.

Instant mobility is useful for:

- Volume migration from a VPSA that runs out of capacity, into a VPSA that has free capacity
- Volume migration between VPSA Storage Array and VPSA Flash Array
- Volume migration between sites
- Volume migration from a private cloud to the public cloud

13.3.2 Offline processing

Offline processing requires the creation of an instant clone of a volume on another VPSA (local or remote). The new cloned volume can then be used for offline processing, without affecting the original production volume.

Offline processing is useful for:

- Development and Test
- Analytics
- Reporting

The following table summarizes an example of a sequence of steps, including high-level steps for remote cloning, for implementing the [Instant mobility](#) or [Offline processing](#) use cases:

Step	Action	For instant mobility	For offline processing
1	If relevant: Stop the application that accesses the volume to be migrated.	✓	☐
2	Take a snapshot of the volume to be cloned from the source VPSA. See Manual creation and deletion of snapshots .	✓	✓
3	Connect source and destination VPSAs .	✓	✓
4	Create a remote clone .	✓	✓
5	Attach a clone to a server .	✓	✓
6	If relevant: Start offline processing on the server accessing the new volume.	☐	✓
6 7	If relevant: Restart the application that will access the new volume.	✓	☐

13.4 Connect source and destination VPSAs

A clone can only be created on a pair of VPSA's that are known to each other. To establish the connection between VPSAs, remote VPSAs are discovered and defined in the same way as for remote mirroring.

If the VPSAs are located in different Zadara Storage Clouds, you will need to first assign a public IP to each VPSA.


Follow the details in [Connect to a remote VPSA](#) to discover the remote VPSA and establish a connection.

13.5 Create a remote clone


A remote cloned volume is created on the destination VPSA as a new type of volume.

To create a remote cloned volume:

1. Go to **Resources > Volumes**.
2. In the **Volumes** view, click **Create** and select **Create Volume from Remote Snapshot** from the dropdown.



3. In the **Create Volume from Remote Snapshot** dialog that opens, select the source VPSA volume's snapshot to clone:



1. **Source VPSA:** From the list of VPSAs in the dropdown, select the remote VPSA from which to clone.

2. **Source Volume:** From the list of volumes in the dropdown, select the remote VPSA's volume to clone. The selected volume's list of snapshots is displayed in the **Source Snapshot** table.
3. **Source Snapshot:** Click the snapshot to use, from the list of snaps of the selected source volume.
4. Click **Next**. The dialog switches, for entering the destination volume details.

Create Volume from Remote Snapshot

New Volume Name: *

Select a Pool: *

Name	Status	Free Capacity
pool1	normal	5.23 TiB Free / 17.35 TiB
pool2	normal	125 GiB Free / 6.88 TiB

Encrypted: ☐ Compress: ☒

Attach Default Snapshot Policies: ☒ Dedupe: ☒

☆ The above configurations are according to the source volume (smb1).
You may change them now. The changes will affect the new volume only.
Snapshot Policies will take effect only after background data transfer is completed

Retrieval Mode: ⓘ ☒ Background ☐ On-demand

Create **Back**

5. **NewVolume Name:** Enter a name for the new volume.
6. **Select a Pool:** From the list of available storage pools, click the pool from which to allocate storage resources for the new volume.
7. **Encrypted:** To encrypt the new volume, mark the checkbox.

✓ **Note:** A cloned volume can be encrypted, even if the source volume is not encrypted.

8. VPSA Flash Array To dedupe or compress a new volume on a VPSA Flash Array, mark the checkboxes.

✓ **Note:** A volume that is being cloned on a VPSA Flash Array can be deduped or compressed, even if the source volume is not.

9. Select the **Retrieval Mode** for the cloning process:
 - **On-Demand:** Blocks of data are retrieved as needed, until all the source volume data is copied to the destination volume.
 - **Background:** All the source volume's data is retrieved in the background.
10. Click **Create**.

13.6 Monitoring a remote clone

13.6.1 Destination VPSPA

VOLUMES

[REFRESH](#)
[CREATE](#)
[DELETE](#)
[EXPAND](#)
[SERVERS](#)
[QUOTAS](#)
[DATA SERVICES](#)

Filter Add Filter:

Name	Capacity	Status	Attributes	Data Type	Compress	Dedupe	Pool	Server(s)
vol1	1 TiB	Available		File-System	Enabled	Enabled	RAID-10-Pool-1	
vol2	50 GiB	In-use		BLOCK	Enabled	Enabled	RAID-10-Pool-1	TEMPLATE
MyNewClone	50 GiB	Cloning		BLOCK	Enabled	Enabled	RAID-10-Pool-1	

<< < PAGE 1 OF 1 > >>
Displaying 1 - 3 of 3

Details for MyNewClone

[Properties](#)
[Remote Snapshot Status](#)
[Snapshots](#)
[Object Storage Snapshots](#)
[Snapshot Policies](#)
[Servers](#)
[Containers](#)
[Metering](#)
[Logs](#)
[Performance Alerts](#)
[Tags](#)

BREAK

General		Source Information	
Name:	dstrclone-00000005	VPSPA Name:	my_VPSPA
Status:	Syncing	Volume ID:	
Mode:	On-demand	Volume Name:	smb1
		Snapshot Name:	s-r2-1701702050581
		Created At:	2023-12-14 14:37:59
		Modified At:	2023-12-14 14:37:59

- Cloned volume properties are identical to those of a regular volume.
- Capacity of a clone shows the virtual capacity of the original volume, and the physical capacity of the cloned volume.
- While retrieving data, the system detects whether the clone is connected. If connectivity to the source snapshot is lost during data retrieval, the clone becomes unavailable.
- During the cloning process, the destination volume is listed in the **Volumes** table with its **Status** column entry displaying the value **Cloning** together with a connectivity icon, indicating the activity in progress.
- The cloning data retrieval's status is displayed on the south panel's **Remote Snapshot Status** tab, while data is copied from the source volume to its clone.
- During the cloning process, you can break the connection between the clone and its origin, by clicking the **Break** button in the **Remote Snapshot Status** tab.
- During the cloning process in **Background** mode, there are also controls to **Pause** and **Resume** the data transfer. This feature is useful when experiencing load problems on the system.
- After all the data is retrieved, the relationship between the source and the cloned volumes is broken, the cloned volume becomes a regular volume and the **Remote Snapshot Status** tab is no longer displayed.

13.6.2 Source VPSA

VOLUMES

REFRESH | CREATE | DELETE | EXPAND | SERVERS | QUOTAS | DATA SERVICES

Filter

Name	Capacity	Status	Attributes	Data Type	Pool
smb1	50 GiB	Available		BLOCK	RAID-10-Pool-1
nas1	1 TiB	Available		File-System	RAID-10-Pool-1

<< < PAGE 1 OF 1 > >>

Details for smb1

Properties | **Remote Clones** | Snapshots | Object Storage Snapshots | Snapshot Policies | Servers | Containers | Metering | Logs | Performance Alerts | Tags

VPSA Name	Provider Name	Destination Volume Name	Snapshot Name	Compress	Dedupe
vpsa1_FLASH	dummy	MyNewClone	snap-00000006	Yes	Yes
vpsa6	dummy	clone6	snap-00000006	No	No

- On the source VPSA during the cloning process, the south pane of a source volume being cloned displays the additional **Remote Clones** tab. This tab lists the names and details of the destination volumes in the process of being cloned from the selected source, and the destination VPSA on which they are being cloned.
- On completion of a cloning process, the entry of the successfully cloned remote volume disappears from the table in the source volume's **Remote Clones** tab. On completion of all cloning processes for a source volume, there are no longer any entries in the table in its **Remote Clones** tab, and the tab is no longer displayed.

13.7 Attach a clone to a server

A remote clone is attached to servers on the target VPSA, in the same way as a regular volume. Similarly, a cloned volume can be detached from a server in the same way as a regular volume. Follow the instructions in [Attaching & detaching Volumes to Servers](#).

13.8 Data services

The following data services are not available on destination volumes while the source data is being retrieved:

- Snapshots
- Clones
- Mirrors
- B2OS

After all the data is retrieved and the relationship between the source and the clone has been broken, all data services become available.

13.9 Delete a clone

Deletion of a remote cloned volume is done on the destination VPSA, and it breaks the relations with the source volume. You can delete a remote clone in the same way that you delete any other volume. Follow the instructions in [Creating and Deleting a Volume](#).

BACKUP TO OBJECT STORAGE

Zadara VPSA provides built in backup and restore capabilities to Zadara Object Storage, AWS S3, Google Cloud Storage, Azure Blob Storage or any other S3 compatible object storage. The backup process involves transporting VPSA Snapshots to the remote Object Storage for safe keeping.

Backup to Object Storage (B2OS) allows you to store a backup of the VPSA volume on Object Storage and later restore it to its original VPSA or to **any** other VPSA in a different location with access to the same object storage bucket.

14.1 Creating New Backups

In order to create a Backup for a given Volume, you must first have the Remote Object Storage connected as explained in [Connecting to Remote Object Storage](#).

To create a Backup:

- Open the [VPSA GUI > Backup to Object Storage](#) and click the Create button.
- Give the new Backup Job a name

Hint: Backup job name, similarly to other VPSA entities is limited to:

- 128 characters in length
 - ASCII characters only (between 32 and 126) with the exception of the following special characters: ", %, ;, {, }, [,], <, >, \, ' , &
 - Cannot have a starting or trailing whitespace
-

- Select the Volume to be backed up
- Select the Remote Object Storage to be used
- Select a Snapshot Policy. Snapshots created by the selected Policy are stored in the Object Storage bucket

✓ **Note:** Snapshot Policies used for backup purposes are the same Snapshots used locally within the VPSA.

- (AWS S3 Only) Select the SSE (Server-Side Encryption) - AES256, KMS(Default KMS Key), KMS Key ID(User defined KMS Access ID) (AWS S3 Only)
- (AWS S3 Only) Select Storage Class for backup data placement. Besides S3 standard storage class Backups can be also sent to S3 Intelligent Tiering or S3 Infrequent Access storage class.

✓ **Note:**

1. S3 Storage classes can optimize overall S3 costs for specific data types and retention policies. Please consult AWS documentation and consider your backup retention policy before selecting a storage class.
 2. In case the object storage backup is planned to be used as a seed for remote mirroring via Import Seed mode (i.e. initiate remote mirroring based on two VPSAs using the same volume) ensure volume encryption is disabled on the volume. In case the volume is encrypted, Import Seed mode will be disabled for the B2OS data upon recovery
-

- Check the Compress Data box if you want to compress the data in flight. This may save on the traffic fees
- Press Submit

14.2 Monitoring Backups

Remote Object Storage Backups can be managed and monitored from the VPSA GUI.

Go to the Backup to Object Storage page. It lists all of the jobs that have been configured. From this page you can perform the following actions on each Backup Job, regardless of the parameters given when the Backup Job was created:

- Top menu buttons:
 - Delete the Backup Job
 - Pause / Resume
 - Enable / Disable compression
 - Rate Limit - Limit the Backup Job's bandwidth (MB/s)
 - Change the Snapshot Policy of the Backup Job
- In a Backup Job's **Properties** tab:
 - Add a **Comment** to a Backup Job
 - Change a Backup Job's target S3 **Storage Class** (AWS S3 Only)

✓ **Note:** If the target S3 **Storage Class** setting is modified for a specific Backup Job, the new **Storage Class** will be applied to backups taken after this change was performed. Previously created backups copies will not be modified.

The Backup Job details are shown in the following South Panel tabs:

Properties

Each job includes the following properties:

Property	Description
ID	An internally assigned unique ID.
Name	Name that was given at creation time
Comment	User free text comment. Can be used for labels, reminders etc...
Status	Current job status: Idle / Running
SSE	(AWS S3 Only) Server side encryption type
Storage Class	(AWS S3 Only) S3 target storage class for backup copies
KMS Key ID	(AWS S3 Only) AWS KMS key ID (for SSC with KMS Key ID)
Snapshot Policy	The Snapshot Policy used by this job.
RPO	Time stamp of the most recent successfully backed up Snapshot.
Compression	Compression enabled: Yes / No
Created	Creation time stamp.
Modified	Last modify time stamp.
Source Volume	Name of the protected Volume.
Destination Type	Type of the Remote Object Storage.
Account	Account on the Remote Object Storage.
End Point	Location of the Remote Object Storage.
Bucket	Bucket in the Remote Object Storage where the backups are kept.

Local Snapshots

The Local Snapshots tab lists the point-in-time Snapshots of this Volume that were created for backup purposes by the selected job.

The following Properties are provided per Local Snapshot:

Attribute	Description
ID	Snapshot ID
Name	Display Name.
TimeStamp	Snapshot creation time stamp
Status	Normal/Pending Deletion/Deletion

Object Storage Snapshots

The Object Storage Snapshots tab lists the point-in-time Snapshots of this Volume as stored in the Remote Object Storage. These snapshots were created by the selected job.

The following Properties are provided per Object Storage Snapshot:

Attribute	Description
ID	Snapshot ID
Name	Display Name.
TimeStamp	Snapshot creation time stamp.
Status	Normal\Pending Deletion\Deleting

Metering - The Metering Charts provide live metering and statistics of the IO workload associated with the selected Backup Job.

The following charts are displayed:

Chart	Description
Bandwidth (MB/s)	Total throughput (in MB) of backup data transferred to the Remote Object Storage.
IO Time (ms)	Average response time IO commands issued by the Backup Job during the selected interval.


Logs – The Logs tab displays a list of event log messages related to that Backup Job.

RESTORE FROM OBJECT STORAGE

Zadara VPSA provides built in backup and restore capabilities to Zadara Object Storage, AWS S3, Google Cloud Storage, Azure Blob Storage or any other S3 compatible object storage. The backup process involves transporting VPSA Snapshots to the remote Object Storage for safe keeping.


Backup to Object Storage (B2OS) allows you to store a backup of the VPSA volume on Object Storage and later restore it to its original VPSA or to **any** other VPSA in a different location with access to the same object storage bucket.

In order to restore a Volume from a Snapshot in Remote Object Storage, open the [VPSA GUI > Restore from Object Storage](#) page and click Create. In the dialog that opens select the Remote Object Storage, and navigate to the bucket (VPSA / Volume / Snapshot) to restore from. Click Next.



✓ **Note:** Since listing of large buckets may be time consuming there is an option to specify the full path of the snapshot to restore from (if known). The path should be given in the following format:

<cloud_name.cloud_uid/vpsa_name.vpsa_id/volume_name.volume_id/object_snapshot_name>



The Restore Job creates a new Volume from the selected Snapshot. Restore supports three modes of operation:

1. **Restore** – This mode is useful for creating a full copy of the Volume from the Snapshot, to be used for offline processing. In this mode there is no need to wait for all of the data to be transferred back. The new Volume can be immediately attached to the Host. If the Host needs data that is not yet restored the system will get it on demand.
2. **Clone** – This mode is useful for restoring a small amount of data (a few files) without needing to copy the entire Volume capacity from the Object Storage. Again, the new volume can be immediately attached to the host, but data is only transferred on demand.
3. **Import Seed** – This mode is useful for restoring data from a given point-in-time, subsequently enable synchronization via Mirroring. In this mode a full capacity Volume is created, but you have to wait until all of the Volume's capacity is restored before you can use it.



Warning: In case the object storage backup is planned to be used as a seed for remote mirroring via Import Seed mode (i.e. initiate remote mirroring based on two VPSAs using the same volume) ensure volume encryption is disabled on the source volume. In case the volume is encrypted, Import Seed mode will be disabled for the B2OS data upon recovery

To create a new Restore Job:

- Give the new Volume a name.
- Select the restore mode.
- If you want the new Volume to be encrypted check the Encrypted box.
- Select a Pool to contain the new Volume.
- Press Submit.

A Restore job is then generated and begins working according to the selected mode. You may switch between Restore and Clone mode while the job is running by clicking the Switch to... button. This button toggles depending on its current status.

IMAGES

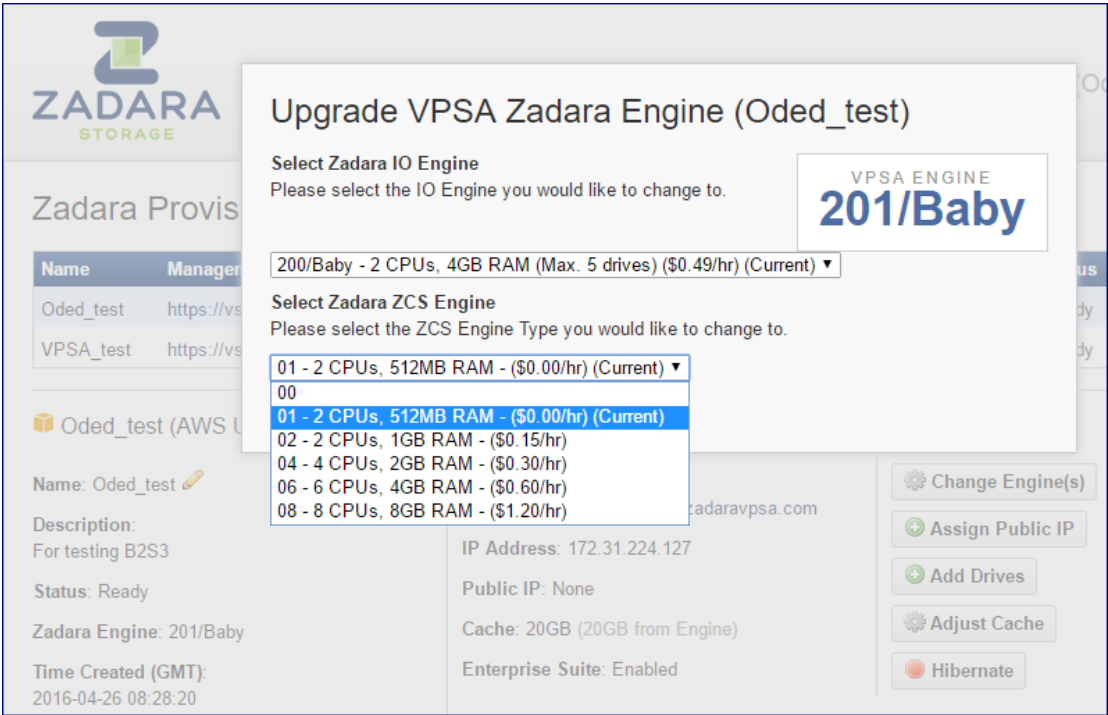
Zadara Container Service (ZCS) makes it possible to run arbitrary processing tasks from directly inside the storage. This is possible due to Zadara’s convergence of Docker Container technology into the Zadara Engines. The benefit of data processing inside the storage, rather on a connected Server, is the direct, low latency access to the data Volumes.

16.1 Adding ZCS Engines

In order to run ZCS Containers within a VPSA a ZCS engine is needed in addition to the IO engine. The ZCS engine contains the compute resources of the VPSA’s Virtual Controllers that are allocated for the Docker Container.

The ZCS can be added when the VPSA is originally created, as described in [Creating a VPSA](#), or it can be added at a later time.

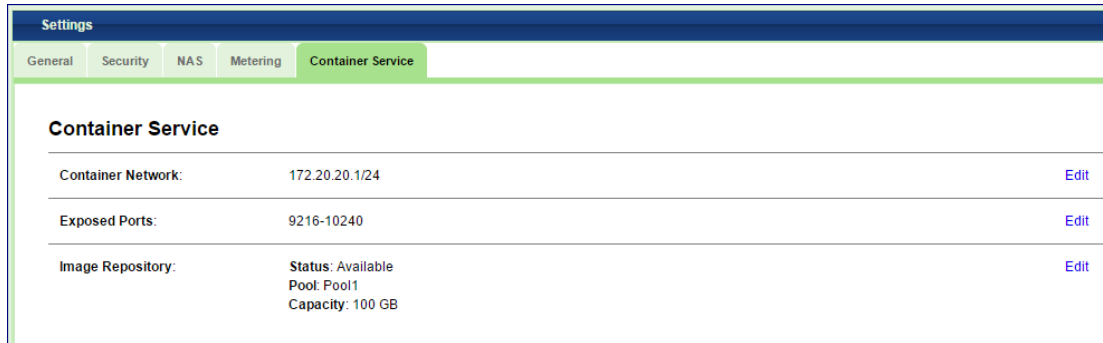
To add a ZCS engine go to the Zadara Provisioning Portal, select the relevant VPSA, click Change Engines and select the engine size that fits the needs of the application that will run in the Container.



16.2 Creating an Image Repository

This one-time operation is needed in order to reserve some storage space for storing all of the Container images you plan to use.

To create an Image Repository, open the [VPSA GUI > Settings > Container Service](#) tab and click Edit on the Image Repository section. Select the Pool that will host the Image Repository.

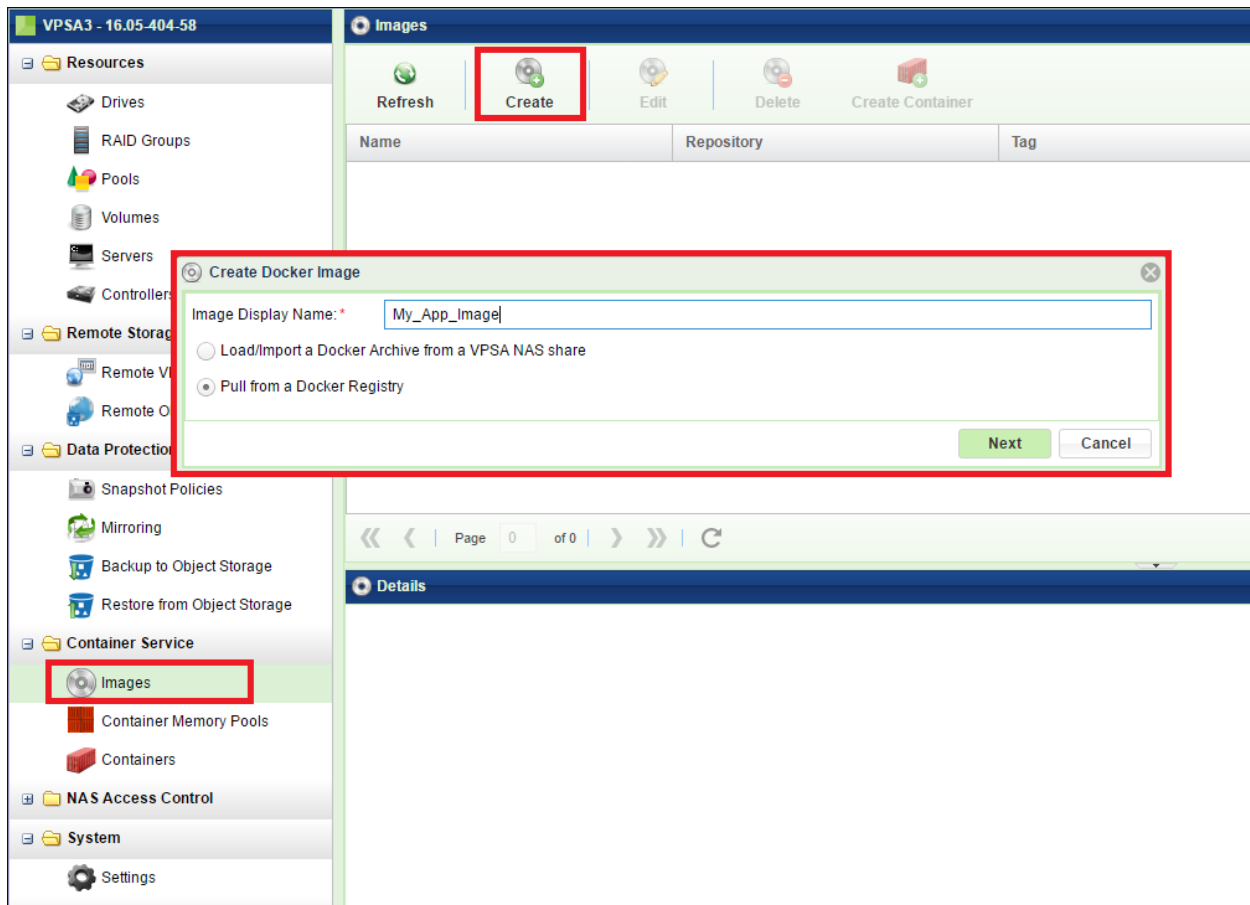


Settings		
General Security NAS Metering Container Service		
Container Service		
Container Network:	172.20.20.1/24	Edit
Exposed Ports:	9216-10240	Edit
Image Repository:	Status: Available Pool: Pool1 Capacity: 100 GB	Edit

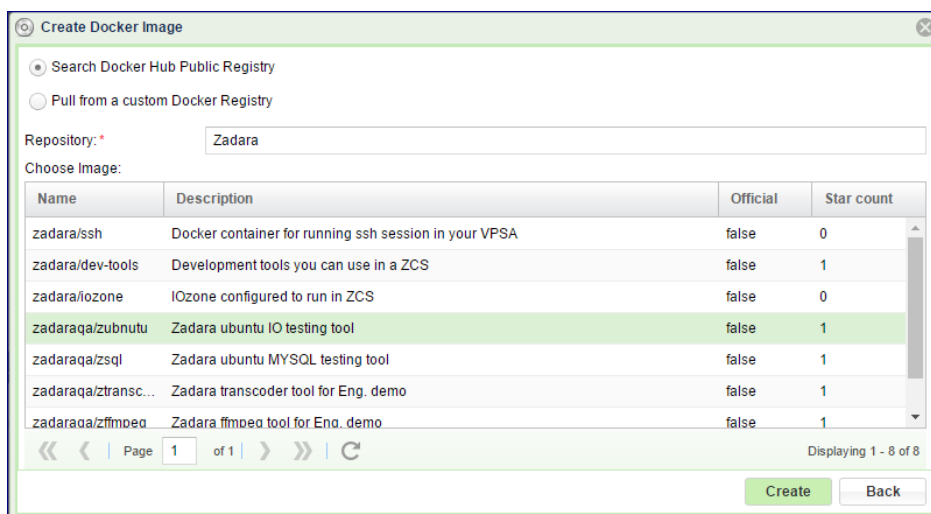
16.3 Creating a Container Image

Before you can create a Container its Image must be entered into the Image Repository. You can take the image from any NAS share, or download it from Docker Hub (<https://hub.docker.com>).

To place an Image into the Image Repository open the [VPSA GUI > Images](#) and click Create.



In the dialog that opens, give a name for the new Docker Image and select if you want to download the image from a Docker Hub or if you want to load it from a NAS share on this VPSA. Click Next .



Search for and select the Image and click Create .

✓ **Note:** It might take a while to download the image, depending on its size and the Internet connection's speed. Wait for the image status to become "normal." Your image is now ready for use.

Images

Refresh

Create

Edit

Delete

Create Container

Name	Repository	Tag	Status	Capacity
My_first_image	zadaraqa/zubnutu	latest	normal	490 MB

CONTAINERS

Zadara Container Service (ZCS) makes it possible to run arbitrary processing tasks from directly inside the storage. This is possible due to Zadara's convergence of Docker Container technology into the Zadara Engines. The benefit of data processing inside the storage, rather on a connected Server, is the direct, low latency access to the data Volumes.

17.1 Adding ZCS Engines

In order to run ZCS Containers within a VPSA a ZCS engine is needed in addition to the IO engine. The ZCS engine contains the compute resources of the VPSA's Virtual Controllers that are allocated for the Docker Container.

The ZCS can be added when the VPSA is originally created, as described in [Creating a VPSA](#), or it can be added at a later time.

To add a ZCS engine go to the Zadara Provisioning Portal, select the relevant VPSA, click Change Engines and select the engine size that fits the needs of the application that will run in the Container.

ZADARA STORAGE

Zadara Provisioning Portal

Name	Manager
Oded_test	https://vs...
VPSA_test	https://vs...

Oded_test (AWS U...)

Name: Oded_test

Description: For testing B2S3

Status: Ready

Zadara Engine: 201/Baby

Time Created (GMT): 2016-04-26 08:28:20

IP Address: 172.31.224.127

Public IP: None

Cache: 20GB (20GB from Engine)

Enterprise Suite: Enabled

Upgrade VPSA Zadara Engine (Oded_test)

Select Zadara IO Engine
Please select the IO Engine you would like to change to.

200/Baby - 2 CPUs, 4GB RAM (Max. 5 drives) (\$0.49/hr) (Current)

Select Zadara ZCS Engine
Please select the ZCS Engine Type you would like to change to.

01 - 2 CPUs, 512MB RAM - (\$0.00/hr) (Current)

02 - 2 CPUs, 1GB RAM - (\$0.15/hr)

04 - 4 CPUs, 2GB RAM - (\$0.30/hr)

06 - 6 CPUs, 4GB RAM - (\$0.60/hr)

08 - 8 CPUs, 8GB RAM - (\$1.20/hr)

VPSA ENGINE 201/Baby

Change Engine(s)

Assign Public IP

Add Drives

Adjust Cache

Hibernate

17.2 Creating a Container

A Docker Container provides a layer of abstraction and automation of operating-system-level virtualization on Linux. It uses the resource isolation features of the Linux kernel to allow independent “containers” to run within a single Linux instance, avoiding the overhead of starting and maintaining virtual machines. Zadara’s VP SA is utilizing this technology to allow user applications to run within VP SA in an effective and controlled manner. For more details on Docker Containers please refer to the Docker documentation at <https://docs.docker.com>

A Container can access any Volume from its hosting VP SA, excluding NAS Volumes defined as “SMB Only” (see [Creating and Deleting a Volume](#)). A Container can be attached to a single Block Volume or to multiple NAS Shares.

When creating a Container you will need to specify its operating environment such as Memory Pool assignment, Volumes it can access and communication ports.

To create a Container open the [VP SA GUI > Containers](#) and click Create .

The screenshot displays the VP SA GUI interface. On the left, a sidebar menu lists various system components. The 'Containers' option under the 'Container Service' section is highlighted with a red rectangle. The main content area shows the 'Containers' management page. At the top, there are action buttons: Refresh, Create (highlighted with a red rectangle), Start, Stop, Delete, and a CPU Utilization graph. Below these is a table listing three containers: c1, c2, and c3, all using the 'zubuntu' image and running on IP 110.10.2.106. Below the table, the 'Details for c3' are shown, with fields for ID, Name, Image ID, Image Name, Memory Pool ID, Memory Pool Name, Status, Started, IP, and Use Public IP.

Name	Image name	Status	Running	IP	Entry point
c1	zubuntu	normal	Yes	110.10.2.106	/usr/sbin/sshd
c2	zubuntu	normal	Yes	110.10.2.106	/usr/sbin/sshd
c3	zubuntu	normal	Yes	110.10.2.106	/usr/sbin/sshd

Details for c3

Properties

ID: container-00000003

Name: c3

Image ID: img-00000001

Image Name: zubuntu

Memory Pool ID: dgroup-00000001

Memory Pool Name: Default_Memory_Pool

Status: normal

Started: Yes

IP: 110.10.2.106

Use Public IP: No

In the dialog that opens up do the following:

Create Container

Load from Existing

Name: *

Image: *

Volumes

Add Edit Delete

Name	Volume Type	Access	Path
B1	BLOCK	rw	/b1

Container Ports

Add Edit Delete

User Start	User End	Internal Start	Internal End
22		10000	

- Give the Container a name.
- Select the Image for this Container.

✓ **Note:** You must provide a full Container Image. Container files are not supported.

- Assign Volumes that this Container can access.
- Select the Port Ranges this Container will use.


✓ **Note:** Available external ports range is defined in the system settings as described here [Container Service](#).

- Set environment variables
- Set arguments to the entry point (see below)
- Set links to other Containers, so that this Container will only run while the others are running too.
- Select a Memory Pool or leave it empty to use the default Memory Pool.
- Entry point is the program or the daemon to execute in the Container.
- Select whether the Container will start immediately following its creation.
- Allow the Container to use the public IP of the VPSA (if any).

Creating a Container from an existing one

You can avoid repeatedly entering the same Container parameters over and over again for each Container created. When creating a Container similar to an existing one you can use the Load from Existing option and just modify parameters as required.

Create Container

 Load from Existing

Name: * c3

Image: * zubuntu

Volumes

Add Edit Delete

Name	Volume Type	Access	Path
B1	BLOCK	rw	/b1

Container Ports

Add Edit Delete

User Start	User End	Internal Start	Internal End
22		10000	

17.3 Monitoring Containers

The Containers details are shown in the following South Panel tabs:

Details for c1

Properties Volumes Port Ranges Environment Variables Args Links Logs Metering

ID: container-00000001

Name: c1

Image ID: img-00000001

Image Name: zubuntu

Memory Pool ID: dgroup-00000001

Memory Pool Name: Default_Memory_Pool

Status: normal

Started: Yes

IP: 110.10.2.106

Use Public IP: No

Entry Point: /usr/sbin/sshd

Properties

Each Container includes the following properties:

Property	Description
ID	An internally assigned unique ID
Name	Name that was given at creation time
Comment	User free text. Can be used as label, reminder, ect...
Image ID	An internally unique ID of the Container Image
Image Name	Name of the Container Image
Memory Pool ID	An internally unique ID of the assigned Memory Pool
Memory Pool Name	The name of the assigned Memory Pool
Status	Normal / Failed / Creating / Deleting
Started	Yes /No
IP	IP address assigned to the container
Use Public IP	Yes / No
Entry Point	The entry point program/daemon

Volumes

The Volumes tab lists the Volumes that the selected Container can access.

Port Ranges

The Port Ranges tab lists all of the Ports that are assigned to the selected Container.

Environment Variables

This tab lists all of the Environment Variables to be used in the Container.

Args

The Args tab lists all of the Arguments for the entry point execution.

Links

The Links tab lists all of the Links from the selected Container to other Containers. These other Containers must run for the selected Container to run.

Logs

The Logs tab lists all of the event log messages related to that Backup Job.

Metering

The Metering Charts provide live metering of the Container's memory consumption (Only appears when the Container is running).

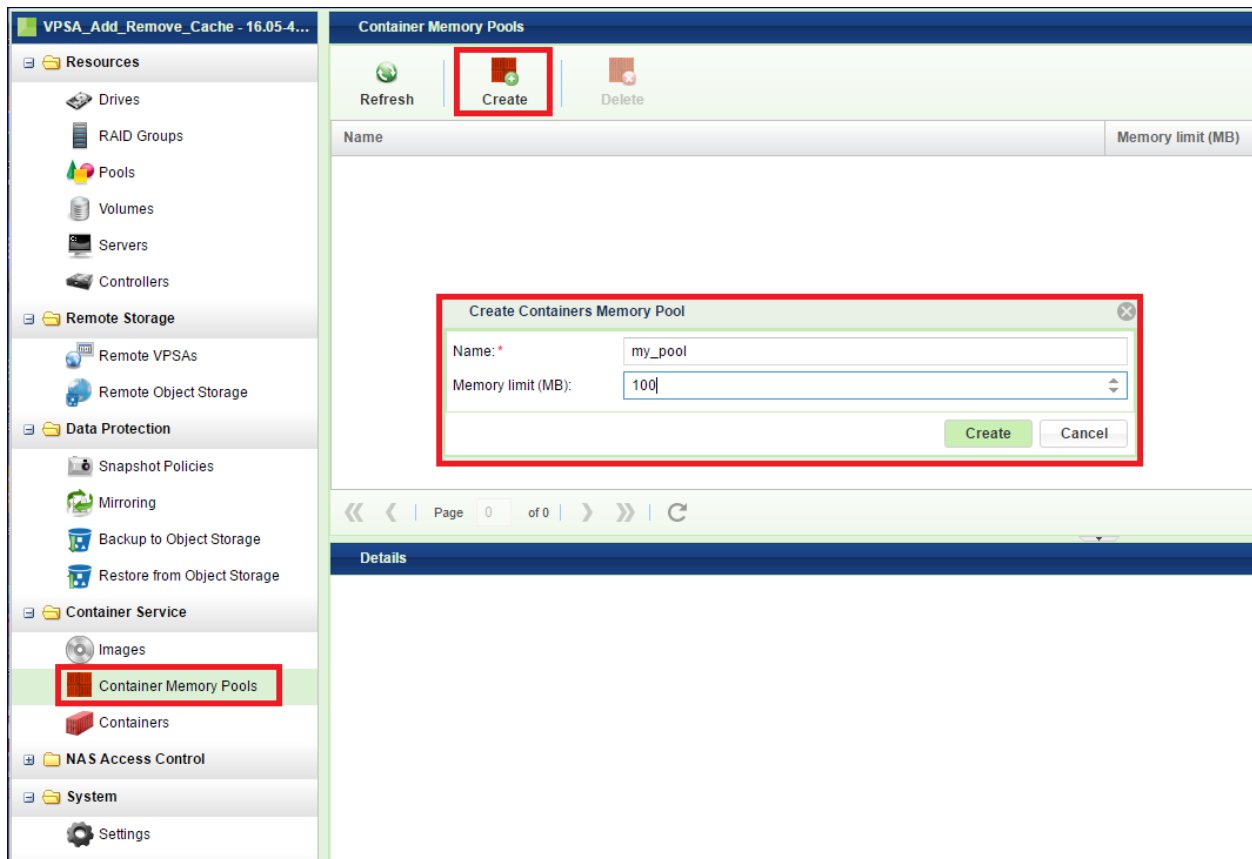
✓ **Note:** It is not possible to update/edit the configuration of an existing Container. The Container must be deleted and recreated with the required settings.

CONTAINER MEMORY POOLS

A Container Memory Pool helps with managing and controlling the memory allocated to Containers. Containers run in the ZCS engine and compete with each other. To avoid a situation where some Containers consume all of the memory resources potentially leaving other Containers unable to run, you can create Container Memory Pools. Each Container can be assigned to a Memory Pool, limiting it only to consume memory from that Pool. Containers that are not assigned to any specific Container Memory Pool consume memory from the default Memory Pool, which holds all the engine memory not allocated to any specific Container Memory Pool.

18.1 Creating a Container Memory Pool

To create a Container Memory Pool open the [VPSA GUI > Container Memory Pools](#) and click Create .



1. Give the new Memory Pool a name.

2. Select the amount of memory to allocate to this Container Memory Pool.

✓ **Note:** The combined total of all of the Pools' limits must be less than or equal to the memory size of the ZCS engine.

3. Click Create .

FILE LIFECYCLE

19.1 Understanding File Lifecycle Management Analytics

Zadara VPSA provides the ability to collect data over the lifetime of files and report analytics by various parameters and perspectives. File lifecycle management analytics provide visibility into customer NAS environments, by delivering utilization and trend reporting capabilities, allowing insights into current usage of an environment's resources and supporting future optimization. When File Lifecycle Index Management is enabled on a specific NAS share volume, the VPSA performs a single full scan of NAS share files. Subsequent filesystem changes are detected and updated in periodic indexing.

The available file lifecycle analytics reports:

File Lifecycle Analytics Report	Report by
Growth trends	<ul style="list-style-type: none">• All files• File type
Capacity utilization	<ul style="list-style-type: none">• File size• Age• Access• File type• Owner, by:<ul style="list-style-type: none">- Top users- Top groups

The VPSA is preconfigured with file types, grouped into categories and identified by their filename extensions. Categories and recognized filename extensions can be customized. See [Managing File Categories](#).

To enable the analytics reporting, the VPSA is configured to index file lifecycle data.

✓ **Note:** Additional SSD volumes for the file lifecycle management's indexing repository are allocated to the VPSA, to support its infrastructure. File lifecycle management indexing consumes some VPSA compute and memory resources.

File lifecycle management can be enabled on creation of a VPSA, or enabled later or disabled on existing VPSAs, via the Zadara Provisioning Portal. See [Creating a VPSA](#) for more details.


On VPSAs that are enabled for file lifecycle management, it is possible to pause or resume file lifecycle management globally, using the toggle in the [File Lifecycle Management](#) tab in the Settings page.

On VPSAs that are configured with file lifecycle management enabled, individual NAS share volumes can be created with file lifecycle indexing as enabled or disabled. At any time, a NAS share volume's file lifecycle indexing can be enabled, paused, resumed or disabled. See [Creating and Deleting a Volume](#) and [Volume File Lifecycle Management](#) for more details.

✓ **Note:** Disabling file lifecycle indexing for a volume removes all existing data collected for that volume.

19.2 Reporting File Lifecycle Analytics

The VPSA can display a large single report chart filling the full screen, or a user-selectable number of charts, up to three rows down by three across.



To report analytics of file lifecycles according to one or more of the available perspectives:

1. To choose the number of report panes on the screen, in the Analytics page under File Lifecycle Management, open the **Layout** combobox at the lower left.
2. At the top left of the report pane, select a specific volume or all volumes from the **Report Scope** combobox.
3. Click the **File Analytics Report** folder icon to expand and display available analytic reports. The most recent date and time that data collection (indexing) took place is displayed as **Last update** in the report header. Updates occur once every 24 hours.
4. Click and drag the selected report to the desired panel. From the **Report Scope** combobox on individual report panes, you can change the display to report analytics for a different volume or all volumes.
5. To switch between graphic and tabular report styles, at the top right of the report pane, click the **Display** combobox to choose one of the styles.

✓ **Note:** Some reports offer display styles that are not available for other reports.

6. On graphic reports, mouse-over points of interest or shaded areas to display details of specific metrics.


FILE CATEGORIES

20.1 Managing File Categories

The VPSA is preconfigured with file types, grouped into categories and identified by their filename extensions.

The File Categories page under File Lifecycle Management displays a table where each row is a file category and the list of filename extensions associated with that category. A filename extension can be associated with only one category. The analytics feature reports based on data collected per filename extension and its associated category.

File categories and filename extensions can be customized and configured.



After customizing file categories using the Create, Edit or Delete options, click **Refresh** to display the updated configuration in the File Categories table.

✓ **Note:** Changes applied to file categories are reflected in analytics reporting after the next volume index cycle.

20.1.1 Creating a File Category

The VPSA supports a maximum of 12 categories of files. If there are already 12 categories, consider merging two categories by removing filename extensions from one category and adding them to another category, using the **Edit** option. Then, use the **Delete** option to remove the category that remained without any associated filename extensions.

1. In the top menu, click **Create** to add a new file category.
2. In the **Create Category** dialog box, enter:
 1. **Category Name:** A meaningful name indicating the type of files (based on the filename extensions) that are grouped together for association with the category.
 2. **File Extensions:** For each filename extension that will be associated with the category, enter the extension and click **Add**. The file extension will be added to the File Extensions table below.
3. Click **Submit** to save the new file category.

20.1.2 Editing a File Category

To change the name of a file category or to add or delete filename extensions associated with a category:

1. Mark the category in the File Categories table by clicking on it.
2. In the top menu, click **Edit**.
3. In the **Update Category** dialog box:
 - To change the category's name, edit the entry in the **Category Name** field.
 - For each new filename extension to associate with the category, in the the **File Extensions** field enter the extension and click **Add**. The file extension will be added to the File Extensions table below.
 - **To delete an associated file extension from the category:**
 1. Scroll the File Extensions table to locate the file extension to delete.
 2. Mark the file extension row in the File Extensions table by clicking on it.
 3. Click **Delete** above the File Extensions table.
4. Click **Submit** to save the updated file category.

20.1.3 Deleting a File Category

Deleting a file category will also remove all of the filename extensions associated with the category.

 **Note:** To delete an individual filename extension from a category, follow the [Editing a File Category](#) procedure.

1. Mark the category in the File Categories table by clicking on it.
2. In the top menu, click **Delete**.
3. In the confirmation dialog box, confirm the deletion.

ACTIVE DIRECTORY

21.1 Active Directory Authentication

By joining the VPSA to the Active Directory (AD), Users can use the same credentials that are stored in the AD to login to the SMB shares.

✓ **Note:** Microsoft Active Directory requires the following ports for users and computers authentication:

- Kerberos - 88(UDP/TCP)
 - Microsoft-DS - 445(UDP/TCP)
 - LDAP - 389(UDP/TCP)
 - RPC Endpoint mapper - 135(UDP/TCP)
 - RPC - Dynamically-assigned unless restricted, 49152-65535(TCP)
 - DNS - 53(UDP)
-



Warning: Active Directory cannot be used while the VPSA is configured to use the LDAP service. The transition from LDAP to Active Directory based authentication should be handled carefully, as existing NAS permissions may be affected.

21.1.1 Joining the VPSA to Active Directory

To join the VPSA to a Microsoft Active Directory Go to [VPSA GUI > NAS Access Control > Active Directory](#) and click the Join button.

Enter the following information:

1. **Active Directory Server Name**
2. **Domain Name**
3. **Domain NetBIOS Name**
4. **Administrator Name (of the AD Domain)**
5. **Administrator Password (of the AD Domain)**
6. **DNS IP** - Up to three DNS servers IPs, used for domain name resolution.

Advanced options:

1. Active Directory UID Mapping - Use RFC2307 attributes, the UID/GID will be taken from Active Directory attributes(uidNumber, gidNumebr). In case UID mapping is required, it is required to specify the valid id range. In the case of trusted domains enabled it is required to specify the ID range for each trusted domain after joining the Active Directory and trusted domain discovery.
 2. Allow trusted domains - allow users from a trusted domain to access SMB Volumes.
 3. DNS Lookup realms - Indicate whether DNS TXT records should be used to determine the Kerberos realm of a host.
 4. DNS Lookup KDC - Indicate whether DNS SRV records should be used to locate the KDCs and other servers for a realm. Once disabled, the KDC server IP should be provided manually.
- Click the Submit button and then press OK to confirm the following warning message, which requests that you ensure proper permissions of files and folders created on the VP SA shares, prior to joining the AD Domain.

✓ **Note:** The joining of the VP SA to the Active Directory may fail if the time on the VP SA and the Active Directory Domain Controller is out of sync by more than a few minutes. Sync the time and try again. Different time zones are not an issue.

✓ **Note:** As of version 22.06-SP1, in order to avoid access issues - on top of the existing scheduled connectivity check (once in 10 minutes), in case a DC connectivity issue was detected the system will retry to establish connectivity in a 1 minute interval. After 10 attempts the VP SA AD service will be automatically restarted and in case the connectivity wasn't restored the VP SA administrator will be notified.

21.1.2 Modifying an existing Active Directory connection

On an existing Active Directory connection, the following parameters can be modified:

- **DNS IP** - Up to three DNS servers IPs, used for domain name resolution.

Advanced options:

- **DNS Lookup realms** - Indicate whether DNS TXT records should be used to determine the Kerberos realm of a host.
- **DNS Lookup KDC** - Indicate whether DNS SRV records should be used to locate the KDCs and other servers for a realm. Once disabled, the KDC server IP should be provided manually.



Warning: While a VP SA maintains an SMB connection, changing this setting might impact new authentication requests.

If you are considering this type of transition, contact the Zadara support team for additional information.

21.1.3 Changing Active Directory DNS

You can update the DNS servers associated with your Active Directory without leaving the domain. To update the DNS server Go to [VP SA GUI > NAS Access Control > Active Directory](#) , Select the Domain you want to change and click the Configure button. Edit the DNS server(s) IP address(s).

21.1.4 Leaving an Active Directory

To leave the Active Directory, Go to [VP SA GUI > NAS Access Control > Active Directory](#) , Select the Domain you want to leave and click the Leave button (the Join and Leave button toggles depending on the current status).

Enter the Domain Administrator's Name and Password and press Submit.

Press OK to confirm the following warning message, which requests that you ensure proper permissions of files and folders created using AD, before leaving it.

Sometimes there is a need to temporary leave the Active Directory, and re-join the domain at later time. In this case check the Keep Configuration. The domain's details will be kept for future use.


LDAP

22.1 Enabling LDAP Authentication

By integrating the VPSA with an LDAP service, NAS Users can use the same credentials that are stored in the directory service to login to SMB shares.

Starting from VPSA version 19.08, the VPSA SMB service can be configured to authenticate users against LDAP service (JumpCloud or similar).

✓ **Note:** LDAP service requires port 389 for directory connectivity. The communication with the LDAP service would be done encrypted(TLS).

 **Warning:** Using LDAP authentication cannot be used while the VPSA configured to use Active Directory, the transition from Active Directory to LDAP based authentication should be handled carefully, as existing NAS permissions may be affected. If you are considering such transition, contact Zadara support team for additional information.

22.1.1 Configuring LDAP service for NAS authentication

To enable the LDAP service, navigate to NAS Access Control > LDAP and click **Join**.

In the **Join LDAP Server** dialog, enter the following information:

1. **Interface** - the VPSA network interface that will be used for LDAP connectivity. If public service (like JumpCloud), the interface selected must have a direct Internet connectivity. Select one of the following interfaces - Frontend, Public IP (if assigned to the VPSA), Outnet interface (if assigned).
2. **LDAP Server** - the directory service FQDN or IP. FQDN must be resolved by the default public DNS server. (the ldap:// prefix is mandatory).
3. **LDAP WORKGROUP** - as defined in the directory service.
4. **LDAP Search Base** - LDAP search scope DN.
5. **LDAP Bind username** - the DN for the bind user (samba service account)
6. **LDAP Bind Password** - password for the bind user (samba service account)

✓ **Note:** In a case of JumpCloud integration, Samba authentication should be enabled in the target directory. See

<https://support.zadara.com/hc/en-us/articles/360036369912> for a KB article covering JumpCloud specific integration.

Click **Submit**.

22.1.2 Updating LDAP configuration

If the existing configuration needs to be changed, the directory parameters can be updated directly from the VP SA GUI.

Navigate to NAS Access Control > LDAP and click **Configure**.

Once the configuration is submitted, the file services will be restarted in order to apply the new configuration.

22.1.3 Disable LDAP service SMB authentication

If LDAP authentication is no longer needed, the LDAP authentication can be disabled from the VP SA GUI.

Navigate to NAS Access Control > LDAP and click **Leave**.

If you intend to disable LDAP SMB authentication temporarily, you may want to keep the existing configuration for later.

You can restore the configuration by navigating to NAS Access Control > LDAP and clicking **Restore**.

LOCAL NAS USERS AND GROUPS

23.1 Creating NAS Users

By default “root” User and Group at the NFS client are mapped to “root” User and Group in the VPSA. To prevent remote “root” access to the Volume enable the “NFS Root Squash” setting, either at the time the Volume is created or later under Volumes > Properties. All other client-side Users are mapped to User “nobody” and Group “nogroup”.

To configure a basic NAS authentication so that Users and Groups on the NFS client will be mapped to the corresponding Users and Groups at the VPSA, perform the following steps:

- Go to **VPSA GUI > Settings > NAS** tab and press Edit for **NFS Domain**. The NFS Domain dialog will appear:

The screenshot shows the VPSA GUI Settings window with the 'NAS' tab selected. The 'NFS Domain' dialog is open, displaying the 'Domain' field set to 'localdomain' and a 'Save' button. Below the dialog, a table lists various NFS settings:

NFS		Close
Domain:	localdomain	
Save		
NFS4 ID Mapping:	Disabled	Edit
Allow Trusted Domains:	No	Edit
SMB Character Set:	UTF-8	Edit
Defragmentation:	Automatic defragmentation: Enabled Defragmentation Status: Standby	Edit

- Enter the NFS domain name identical to the domain name set in the Client and press the Update button. Typically the default domain name on a Linux client is “localdomain” and is therefore also the default value in the VPSA.

✓ **Note:** On a Linux client the domain name is usually set in the `/etc/ldap.conf` file. It is mandatory to have this value set.

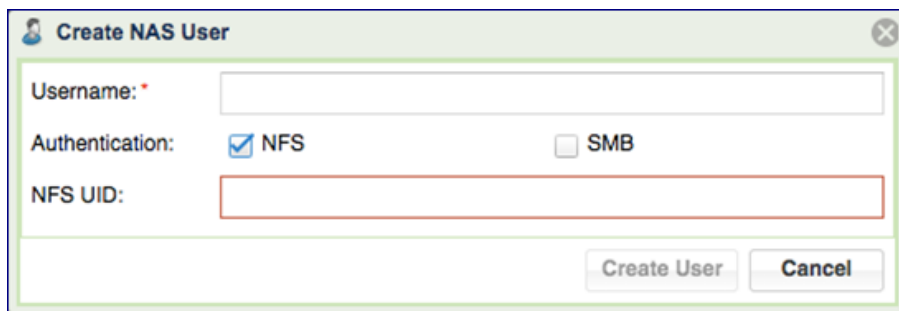
✓ **Note:** Make sure the “idmapd” service is running (Ubuntu = ‘imapsd’, RHEL = ‘rpcidmapd’), and that `/sys/module/nfs/parameters/nfs4_disable_idmapping` is set to “N”. To make this setting persistent, set the following

in '/etc/default/grub' and then run 'update-grub':

```
GRUB_CMDLINE_LINUX_DEFAULT="nfs.nfs4_disable_idmapping="N"
```

- Go to **VP SA GUI > NAS Users** and press the Create button.
- Enter a Username.
- Select the NFS checkbox for Authentication.
- Enter a NFS UID (in the range 1-999,999).
- **If you wish to grant this User access to SMB shares as well, also**
select the SMB checkbox and enter a password (which will be used later when mounting the NAS Volume on a Windows Client).

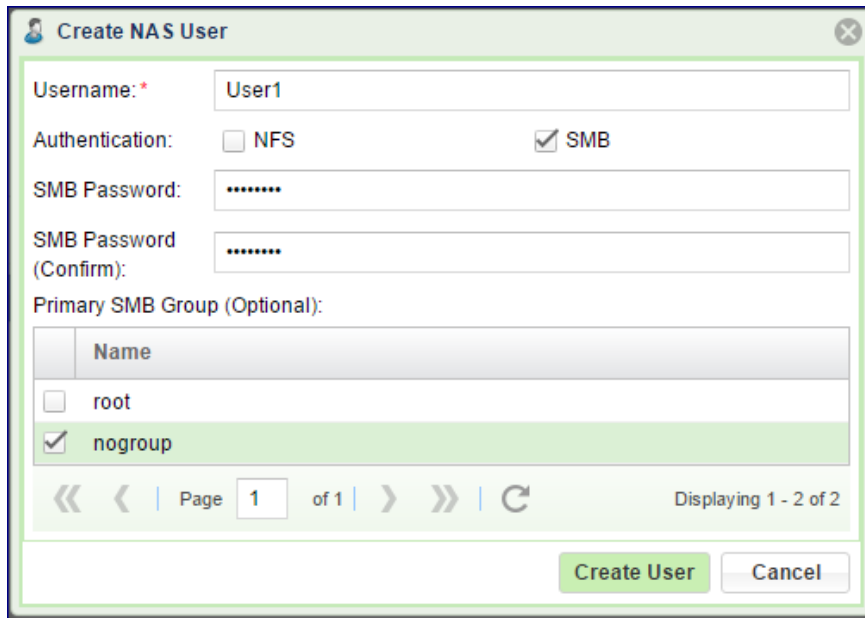
✓ **Note:** This can only be done at the time the User is created, it cannot be changed or added later.



23.2 Creating SMB Users

- Go to **VP SA GUI > NAS Users** and click the Create button.
- Enter a Username.
- Select the SMB checkbox for Authentication.
- Enter a password. You will be asked to provide this username and SMB password when mapping a network drive on the Windows Client.
- If you wish to grant this user access to NFS shares as well, also check the NFS checkbox and enter a NFS UID (in the range of 1-999,999).

✓ **Note:** This can only be done at the time the User is created, it cannot be changed or added later.



Create NAS User

Username: *

Authentication: ☐ NFS ☒ SMB

SMB Password:

SMB Password (Confirm):

Primary SMB Group (Optional):

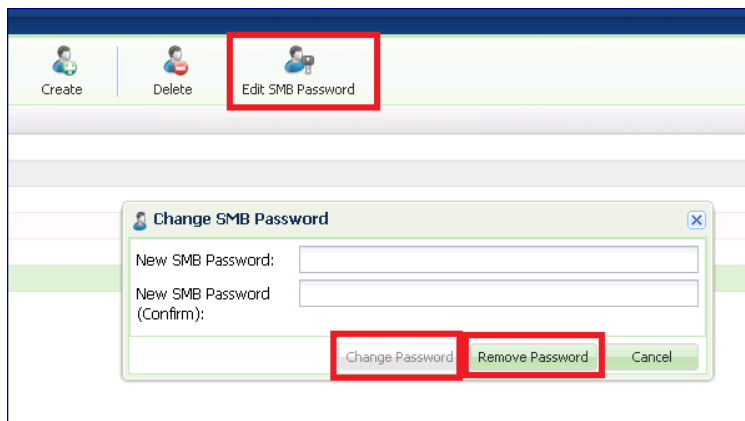
	Name
<input type="checkbox"/>	root
<input checked="" type="checkbox"/>	nogroup

« < | Page 1 of 1 | > » | ↺

Displaying 1 - 2 of 2

23.3 Editing SMB Users Password

It is possible to edit the Password of a SMB User at any time. Go to the NAS Users page and select Edit SMB Password:



Create Delete **Edit SMB Password**

Change SMB Password

New SMB Password:

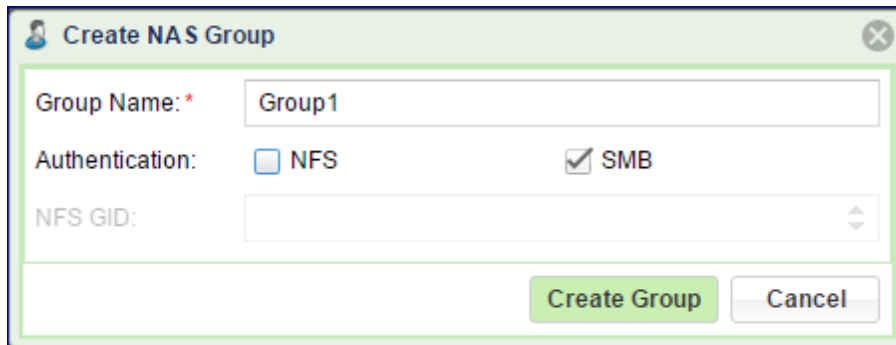
New SMB Password (Confirm):

- To change the SMB Password enter a new SMB Password, confirm the password and click the Change Password button.
- If the User is also defined with a NFS ID you can press the Remove Password button to erase the User SMB Password.

23.3.1 Creating NAS Groups

You can create and view NAS Groups via the NAS Groups page.

To create a NAS Group go to [VPSA GUI > NAS Groups](#) and click the Create button.



Create NAS Group

Group Name: *

Authentication: ☐ NFS ☒ SMB

NFS GID:

- Enter a name for the NAS Group. This should match the Group name on the NFS client.
- Select either NFS or SMB checkbox (or both) for Authentication.
- If you are creating a NFS group also add a valid NFS Group ID (in the range of 1-999,999) that matches the Group Name and GID on your Linux Server.

23.3.2 Managing NAS Quotas

23.4 Enabling or Disabling User/Group/Project Quotas

To enable/disable Quotas on a given NAS share, open the [VPSA GUI > Volumes](#) and select the Volume on which you want to set Quotas. In the South Panel, scroll down to the User Quotas and Group Quotas lines and click the edit icon.

The screenshot shows the 'Volumes' management interface. A table lists three volumes: 'RecoPH' (5 GB, In-use), 'NFS-RecoPH' (5 GB, Available), and 'test_nas_volume1' (2 GB, Available). The 'test_nas_volume1' row is highlighted. Below the table, the 'Details for test_nas_volume1' dialog is open, showing various properties. The 'User Quotas' section is highlighted, showing 'Off' selected among 'Off', 'On', and 'Account only' options. The 'Submit' button is visible.

Name	Capacity	Status	Protection
RecoPH	5 GB	In-use	
NFS-RecoPH	5 GB	Available	
test_nas_volume1	2 GB	Available	

Details for test_nas_volume1

Properties | Snapshots | Object Storage Snapshots | Snapshot Policies | Servers | Containers | Metering | Logs

Data Copies Capacity: 10.24 MB
 Status: Available
 Data Type: File-System
 Pool: pool1
 Server(s):
 NFS Export Path: 10.0.0.1:/export/test_nas_volume1
 SMB Export Path: \\10.0.0.1\test_nas_volume1
 Access Type:
 atime Update: No
 SMB Only: No
 SMB Guest Access: No
 Enhanced Windows ACLs: No
 Directory Creation Mask: 0755
 File Creation Mask: 0744
 Map Archive: Yes
 User Quotas: Off
 Group Quotas: Off

Quotas
 User Quotas: * ☒ Off ☐ On ☐ Account only
 Submit Cancel

In the dialog that opens, select the Off or On option.

✓ **Note:** It is not possible to change the state of Quotas (on/off) when the Volume is attached to a Server. The Volume must be detached from any Servers first.

✓ **Note:** This can also be done on the Volumes tab, select the required Volume, then select Quotas. In here, select Settings > Change Quotas State. In here you can also import and export a Quotas configuration file. See below the format of the Quotas configuration file.

The same process applies for enabling Group and Project quotas.

✓ **Note:** Group quotas and Project quotas cannot coexist on the same Volume.

Quotas Configuration File Format

This is a CSV file where each line sets the quota for a specific user or group.

The line syntax is the following:

```
type,is_user,id,usage,soft,hard,warns,name
```

Where:

- **type:** 1-nfsid or 2-nasuser or 3-aduser
- **is_user:** 0-groups or 1-users
- **id:** uid or gid (if type='aduser' and id is still unknown, set to 0 and name will be translated to id)
- **usage:** 0
- **soft:** 0
- **hard:** hard limit in MB
- **warns:** 0
- **name:** AD name or NAS name

e.g.:

```
1,1,50001,0,0,28,0,-  
3,1,2015348,0,0,24,0,ZADARA\user1  
3,1,0,0,0,24,0,ZADARA\user2
```

23.5 Setting User/Group Quotas

To set quota limits on a given NAS Volume go to the [VPSA GUI > Volumes](#), select the Volume where you want to set up Quotas and click Quotas

The screenshot shows the VPSA GUI interface. At the top, the 'Volumes' section is active, displaying a table with columns: Name, Capacity, Status, Protection, and Data Type. The first row, 'NAS1', is highlighted with a red box. Below the table, the 'Quotas' dialog box is open, also highlighted with a red box. The dialog has tabs for 'Users' and 'Groups'. The 'Users' tab is selected, showing a 'Default Limit' of 0 and an 'Update defaults' button. Below this is a table with columns: Type, Id, Usage (MB), and Limit (MB). The table is empty, and a message at the bottom right says 'No data to display'. The 'Last Updated' timestamp is '2 minutes ago'.

In the dialog that opens, you can set the Quotas for Users, Groups and Projects (as applicable).

If you want to define a default Quota for all Users on the selected Volume, enter the default limit and click the Update defaults button.

Automatic Users discovery:

Press the Refresh button. If this VPSA is connected to an Active Directory the system will scan the AD to find users that have data on this volume. They will all be added and given the default limit. You can edit and change the default value.

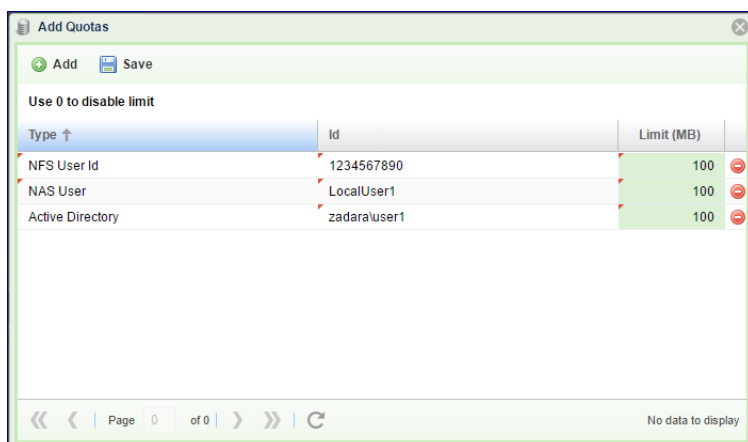
✓ **Note:** Limit set to 0 (zero) means no limit.

If the VPSA is not connected to an Active Directory, a similar scan will be done against all locally defined Users.

Adding User Quotas manually:

In addition, other Users can be added to the Quotas list even if they don't currently have any files on the given volume. Click Add Quotas and then fill in the User details in the line that opens. The User ID should be entered according to the User type. There are 3 User types:

1. **Active Directory user** – the ID is the user name in this format: Domain\username
2. **NAS user** – the ID is the same name as defined in NAS Users.
3. **NFS User** – use the UID as defined in UNIX/Linux systems



Setting Groups Quotas is the same as described above for Users. Click the Groups tab and repeat the same process.

✓ **Note:** For Group Quota accounting the capacity consumed by any individual user is counted against the user's primary group

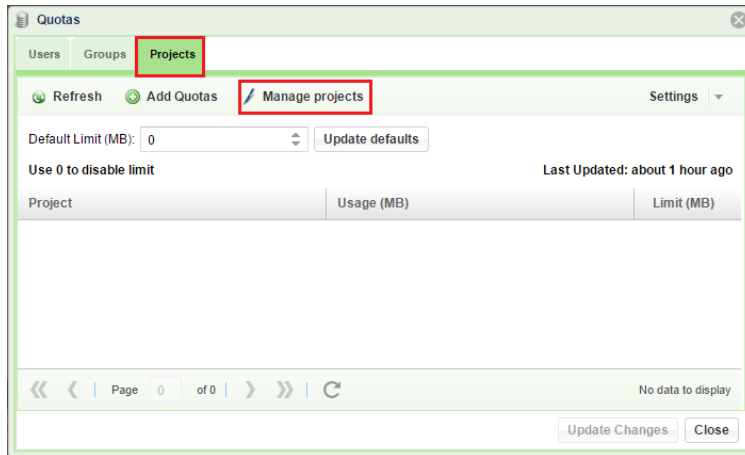
After making any additions or changes to Quota Limits, on the Quotas dialogue box press 'Refresh' to update the figures displayed.

23.6 Setting Project Quotas

Project Quotas are quotas set on a group of one or more folders. Setting these quotas is done in 2 steps: Defining the Projects and then setting the limits.

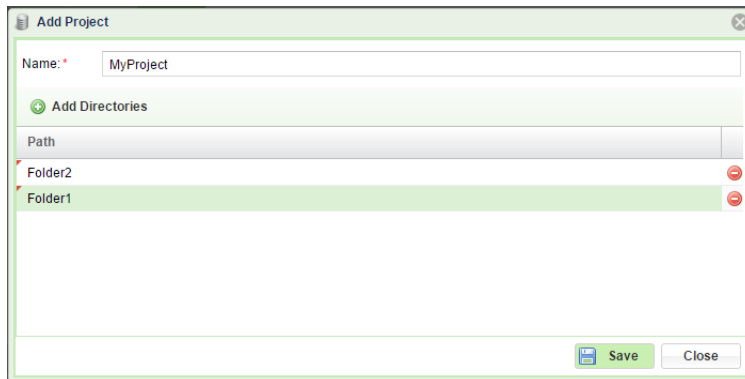
Defining Projects

To define a Project on a given NAS Volume open the [VPSA GUI > Volumes](#) page, select the Volume you want to set Quotas on and click Quotas. On the dialog that opens select the Projects and click Manage projects.



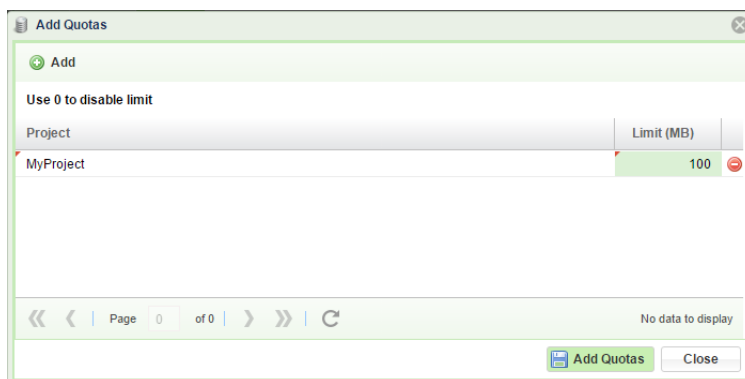
Click Add Project and add directories to this Project. When done click Save and close.

✓ **Note:** The Folders must exist in the Volume, otherwise you will get an error at this point.



Setting Projects Quotas

Click Add Quotas, select the project of interest and set its quota limit. When done click Add Quotas and close.



Finally, on the Quotas dialogue box, press Refresh to update the Quota Limits displayed.

LOCAL USERS

The VPSA's User Management system supports multiple users. There are two distinct user types:

- **Admin** - When the VPSA is created via the Provisioning Portal a default 'admin' user is created. This default 'admin' user cannot be deleted and the password associated with this account should be complex and stored securely. The email address could be a single person, but might be better if it was a distribution list. This Admin user can add, update and delete other Users and reset Users' passwords through the VPSA GUI. It also has full control over all VPSA functions. This should not be confused with a standard User account which has been assigned the 'Admin' privilege.
- **User** - A User who was added by the Admin User. This User has rights to manage the VPSA either through the GUI or REST API, according to their assigned Roles. Each User has their own Password and Access Key.

24.1 User Roles

User Roles define the access rights given to a User. By default, all Users have read rights to all Objects. In addition, the roles define the User's create/update/delete rights for each object type (Pools, Volumes, Backups, etc.). Roles are assigned to each User at creation time and can also be updated later.

24.1.1 Creating a new User Role

When creating a new User Role, give it a name and select the access rights to be granted to the new role. Press **Add Role**.

24.2 Adding and Deleting Users

24.2.1 Adding a new User

Log in to the VPSA the 'admin' user credentials, or as a User who has been assigned the 'Admin' privilege. Go to the Users page and click **Add User**.

Enter the Username and Email address and specify if this new User will be assigned the 'Admin' privilege (full control), or select specific Roles. Select the **Notify on events** checkbox if you want this User to receive email notifications from this VPSA. Then press the **Add User** button to complete the operation.

Once the new User is created, a dialog with a temporary passcode will appear. This passcode is also sent to the Admin User's email. The new User will need to use this temporary passcode when logging into the VPSA for the first time.

24.2.2 Changing a User's Role

The Roles of any given User can be changed at any time. Log in to the VPSA with the 'admin' user credentials, or as a User who has been assigned the 'Admin' privilege. Go to the Users page, select the User from the list and click the Change Roles button.

24.2.3 Deleting a User

Log in to the VPSA with 'admin' user credentials, or as a User who has been assigned the 'Admin' privilege. Go to the Users page, select the User from the list and click the Delete User button.

The User will be deleted, but this operation will not affect any other entities that were created or managed by that User.

24.3 Managing User Passwords

The VPSA stores a cryptographic hash value (using a one-way SHA-3 hash function) of the VPSA User Password. When you log in to the VPSA the entered **password's** hash value is compared with the one stored.

24.3.1 Changing your password

Log in to the VPSA and click your user name on the right upper corner of the screen. Your account page will open. Click **Change Password**.

Enter your current password, a new password and confirm the new password. Click **Change Password** to submit the operation.

✓ **Note:** This operation is available to Admin and to all regular Users. Each User can only change their own password.

24.3.2 Resetting User Password

This operation is available only to the Admin User. The Admin User (or User with Admin privilege) can reset any User's password. A new temporary passcode will be created and sent to the User's email. The User will be requested to set a new password on next log in.

Log in to the VPSA with Admin User credentials. Go to the Users page, select a User from the list, and click **Reset Password**.

24.3.3 Resetting API Key

Zadara Storage employs a session-based authentication mechanism as a means to identify a user for every HTTP request to a VPSA.

You initiate a session by logging in with the VPSA User Password. Upon successful authentication a Secret API Token is sent back to the client application for any subsequent REST API communication with the VPSA to identify the authenticated User and validate the session.

At any time you can generate a new Secret API Token, thus invalidating the previous token and any sessions using it.

Log in to the VPSA, click your user name on the upper right corner of the screen. Your account page will open. Click **Reset Access Key**.

24.4 Managing Password Policy

The VPSA Admin can control the VPSA Password Policy. For details, see VPSA Settings > [Security](#).

24.5 Dual Factor Authentication

The VPSA's User Management system supports Dual Factor Authentication (DFA) using Authenticator mobile application. It is a common practice to protect access in case of compromised password, as a password is not enough in order to login. Each user can turn Dual Factor Authentication on/off for herself. The VPSA admin can force Dual Factor Authentication on all users.

24.5.1 Enabling Dual Factor Authentication

To enable DFA open the current User Properties by clicking the user name on the upper right corner of VPSA GUI screen.

Click **Activate** or **Deactivate**. Close the properties dialog, and log out.

The first time you log in again, a **Confirmation** dialog with a QR code opens.

Install an Authenticator mobile app. (e.g. Google Authenticator) from Google Play or Apple AppStore, and scan the QR code. Enter the code you get on the Authenticator. You are now set.

Every login, from now on will require the temporary code from the Authenticator app.

Important: The mobile device that runs the Authenticator app is needed for login. If the device was lost or replaced, the user must ask the VPSA admin to reset their DFA settings. The VPSA admin must contact Zadara support for resetting the DFA.

24.5.2 Enforcing Dual Factor Authentication

A VPSA administrator can force DFA for all users. In **Settings > Security** click **Edit** on the **Dual Factor Authentication**, check the **Enforce dual factor for all users** checkbox and click **Save**.

This setting change does not have immediate effect. The next time each user will login, they will be required to set their mobile device Authenticator app as described above.

✓ **Note:** When DFA enforcement is removed, the users with DFA configured are still required to use the temporary code when logging in. However each user can change their settings in the user properties as described above.

PERFORMANCE

25.1 Understanding Performance Monitoring

This chapter contains instructions for monitoring the storage performance. The VPSA Performance Monitor allows you to check and monitor the behavior of each element that can affect the overall storage performance, from a single drive to the whole VPSA system and the Servers attached to it.

Each element of the data path can impact the overall performance if not configured and operated properly. The VPSA performance Monitor is a tool for pinpointing a storage performance bottlenecks. The following metrics are of interest in measuring the performance of a storage system:

- **Bandwidth (Throughput):** This value is how much read or write throughput a certain Resource (disk, pool, volume, etc.) delivers. Usually expressed in Megabytes/Second (MB/s)
- **IOPS:** IO operations per second, which is the amount of read or write operations completed in a one second interval. A certain amount of IO operations will also give a certain throughput of Megabytes each second, so these two are related.

Average IO size x IOPS = Throughput

- **Response time (Latency):** is the time it takes each IO operation to complete. Latency is measured in milliseconds (ms) and should be as low as possible.

25.2 The Performance Monitor

To view the VPSA Performance Monitor, go to **System > Performance**.

The Performance Monitor screen consists of the following elements:



1. **Resources Tree:** The Resources Tree lists all the data path objects currently exist in the VPSA

- Pools (including the RAID Groups and Drives that each pool is made of)
- Volumes
- Servers mapped to this VPSA
- Controllers
- System Cache

✓ Note: VPSA Flash Array

For the VPSA Flash Array, there are additional performance parameters for the Pool regarding the data reduction activities and other elements of the data path, such as write buffer activity, dedup accuracy, activity distribution between tiers, and more.

2. **Resource Tile:** The Performance Monitor has 1 to 9 resource tiles depending on the chosen layout. Each tile contain either table or chart.
3. **Layout Selector:** Selector for the number of rows of tiles and tiles per row.
4. **Interval Selector:** Allows switching between different intervals. The interval is a sampling period. Each interval is a single point in the chart. This point represents the average value during that interval. The chart always shows 60 intervals.

For example: If a 1 minute interval is selected, 60 points are displayed, and each point is the average value for that specific minute. In total, the last hour is displayed.

The interval selection affects all tiles.

25.3 Customizing the Performance Monitor

25.3.1 Customizing the Layout

- Go to **System > Performance** and click **Layout**.
- Select the layout of your choice. Note that if the selected layout has fewer tiles than the original, the other tiles will be lost.
- Drag the object of interest from the resources tree, and drop it into a tile. Repeat for all tiles.

25.3.2 Customizing a Tile

Each tile represent a single resource, and provides a number of display options related to the specific resource. The display can be either a table of the most current performance figures, or a chart over time of the recent history.

- To display a chart click **Charts** at the top right corner of the tile, and select the metric of interest.
- To display a table click **Tables** at the top left corner of the tile, and select the table of interest. The table provides performance information as well as other parameters such as **data reduction ratio**.

✓ **Note:** Some of the performance metering charts and table are for Zadara support use only.

SETTINGS

26.1 General

26.1.1 Volumes Recycle Bin

The Recycle Bin is enabled by default, but you can also disable it. When enabled, deleted Volumes are kept in the Pool's Recycle Bin and can be restored. If the Recycle Bin is disabled deleted Volumes are immediately destroyed and cannot be recovered.

26.1.2 Public IPs

This displays any Public IPs assigned to the Controllers. A Public IP allows host connectivity from outside of the VPSA VPN.

26.1.3 Datamover Concurrency Level

You can control the load allowed for datamovers such as mirroring, cloning, etc... by setting the concurrency level. Default is Medium

26.1.4 Zadara Container Services Engine

This displays which, if any, ZCS Engine has been configured via the Provisioning Portal.

26.1.5 Zadara IO Engine

This displays which VPSA IO Engine (Model) has been configured via the Provisioning Portal.

26.1.6 Server Connectivity Monitoring

Server connectivity monitoring allows the VPSA administrator to modify the default settings for all servers records that are set with “Connectivity monitoring” option.

Available configuration:

- Status:
 - Enabled - feature is enabled and notifications (emails) are allowed
 - Enabled (Tickets Disabled) - the feature is enabled however notifications will not be sent to the administrator. Connectivity status can be reviewed in the scope of the server in the VPSA servers view
 - Disabled - feature is disabled globally and will not be present per server
- Connection attempts - number of connection attempts to a server record (default: 10)
- Success threshold - the number of successful attempt that will be considered as normal (default: 6)

26.2 Security

⚙️
SETTINGS

General
Security
NAS
Metering
Container Service
Network
File Lifecycle Management

Security

Passwords policy:	Enforce password expiration: No Password history: 8 Min password length: 8 Password must contain letters: Yes Password must contain numbers: Yes Password must contain uppercase letters: No Password must contain special characters: No	Edit
Dual Factor Authentication:	Enforce Dual factor Authentication: No	Edit
Global VPSA CHAP:	User: admin1 Secret: *****	Edit
Encryption:	No password set	Edit
Support Privilege Access:	Enabled	Edit
Cloud Admin Access:	Enabled	Edit
File Access Audit:	Status: Enabled Audited Operations: User Logon/Logoff, Object Creation, File Open, File Lifecycle Management, Object Attributes Change, Object Rename/Remove, Object Deletion Audit Volume Location: pool-00010003 Object Storage Extraction: None	Edit
IPSec Key:	IPSec Key: *****	Edit

26.2.1 Password Policy

The VP SA Admin can control the VP SA Password expiration policy and password history policy.

26.2.2 Dual Factor Authentication

The VP SA Admin can force all users to login to the VP SA GUI using dual factor authentication. For details see [Dual Factor Authentication](#).

26.2.3 Global VP SA CHAP

This gives you a uniform username and password to use when you create Servers.

26.2.4 Encryption

This sets the Volume encryption password for the VP SA's data-at-rest encryption.

For more information on managing encrypted volumes see [Managing Encrypted Volumes](#).

26.2.5 Support Privilege Access

This controls the ability of Zadara support engineers to access the VP SA virtual controllers with privileged rights. Only the VP SA Admin can change this setting. If enabled, the VP SA Admin gets notification every time the privileged access is used.

26.2.6 Cloud Admin Access

This sets the cloud admin's VP SA GUI access (via the Command Center) to Enabled/Disabled status.

26.2.7 File Access Audit

The VP SA provides the option to enable audit logging of specific file system events. File access audit can be leveraged to comply with organizational security demands and to assist in root cause analysis of specific data access related events.

To implement file access auditing, it must first be enabled globally, and then applied to each volume that should be audited.

To enable file access auditing globally, or to edit the existing file access audit policy:

1. In the **Security** tab in the **Settings** page, click **Edit** on the right of **File Access Auditing**. The **File Access Auditing** configuration dialog box opens.
2. Mark the checkboxes of the access operations to be audited.
3. Click **Save**.

To apply file access auditing to a volume, see [Creating a NAS share](#) under [Creating and Deleting a Volume](#).

26.2.8 IPsec Key

This displays the key to be used when configuring IPsec tunneling for secured host connections.

26.2.9 Allow iSCSI over Public IP

Allow VPSA iSCSI connectivity over Public IP interface (default = False)

26.3 NAS

26.3.1 NFS Domain

This sets the domain name for NFS shares. This defaults to localdomain. (NFS4 Only)

26.3.2 NFS ID mapping

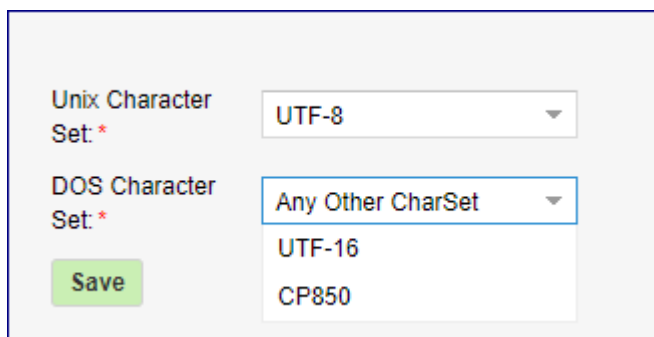
If enabled, each UNIX (Linux) User must be defined as a NAS User in the VPSA. If disabled, UNIX Users are authenticated on the UNIX host side.

26.3.3 SMB NetBios Name

This gives the VPSA Admin the option to change the default name of the VPSA as it appears in Active Directory. This field must be modified **before** the VPSA joins AD.

26.3.4 SMB Character Set

This gives you the default Character Sets used by the SMB service for SMB Volumes. Unix charset - indicates the local character set used by the System. DOS Charset - indicates the Character sets used to communicate with DOS(windows) clients connecting to SMB shares. If you plan to use filenames with different encoding in the filename (other than English), you may want to change the Character Set. The default value for the unix character set is UTF-8 and the default value for the DOS character set is CP850. It is important to note that some character sets can be selected using the listbox items in the setup dialog but all other character sets can be also specified by directly editing the settings field.



The screenshot shows a configuration window for SMB character sets. It contains two dropdown menus. The first, labeled 'Unix Character Set: *', has 'UTF-8' selected. The second, labeled 'DOS Character Set: *', has 'Any Other CharSet' selected. Below the second dropdown is a listbox containing 'UTF-16' and 'CP850'. To the left of these controls is a green button labeled 'Save'.

Changing this value while clients are connected will cause them to temporarily lose access to all SMB shares.

26.3.5 Defragmentation

On HDD media, files are written and deleted in blocks. When files are deleted, data can become fragmented and stored in non-contiguous blocks. Defragmentation reorganizes these blocks, placing file fragments closer together. This reduces seek time and improves file access speed on mechanical drives.

You can enable or disable background file system Defragmentation. You can also run Defragmentation on demand.

By default, Defragmentation is enabled and is scheduled to run every weekend on Sundays at 00:00 (midnight Saturdays night).

To view or change the Defragmentation schedule:

1. In the navigation tree, select **System > Settings**.
2. In the **NAS** tab, next to **Defragmentation**, click **Edit**.
3. Optionally, update the settings:

- **Defragmentation status:**

Displays the current Defragmentation status:

- **Standby:** Waiting to start the next Defragmentation run according to the configured schedule or an on-demand request.
- **Running:** Defragmentation is currently in progress.

- **Start Now / Stop** toggle: Immediately start Defragmentation

- on demand, or stop an in-progress Defragmentation run.

- **Allow automatic Defragmentation:** Enable or disable periodic Defragmentation runs in the background, according to the configured schedule. Select one of:

- **Weekly** (default):

- * **Day of Week:** From the dropdown, select a day of the week.

Sunday (default) - Saturday.

- * **Hour:** From the dropdown, select the time of day that the Defragmentation will run.

00:00 (default) - 23:00, on the hour only.



Note: The default schedules for Defragmentation and File System Trim vary in both day and time, to distribute system overhead more efficiently.

- **Monthly:**

- * **Day of month:** From the dropdown, select a day of the month.

The range is restricted to 1 - 28, to avoid discrepancies between longer and shorter months.

- * **Hour:** From the dropdown, select the time of day that the Defragmentation will run.

00:00 - 23:00, on the hour only.

1. Optionally, click **Default** to reset the settings to their defaults.
2. Click **Submit** to save the changes, or **Cancel** to exit without applying changes to the settings.

26.3.6 File System Trim

On SSD media, files are written and deleted in memory cells. When files are deleted, the SSD might not immediately recognize the freed space, leading to performance degradation over time. The File System Trim (FSTRIM) operation overcomes this by informing the SSD which blocks are no longer in use, allowing them to be erased and reused efficiently. This helps maintain consistent performance and extends the lifespan of the drive.

You can enable or disable background File System Trim. You can also run File System Trim on demand.

By default, File System Trim is enabled and is scheduled to run every weekend on Saturdays at 00:00 (midnight Friday night).

To view or change the File System Trim schedule:

1. In the navigation tree, select **System > Settings**.
2. In the **NAS** tab, next to **File System Trim**, click **Edit**.
3. Optionally, update the settings:

- **File System Trim status:**

Displays the current File System Trim status:

- **Standby:** Waiting to start the next File System Trim run according to the configured schedule or an on-demand request.
- **Running:** File System Trim is currently in progress.
- **Start Now / Stop** toggle: Immediately start File System Trim
- on demand, or stop an in-progress File System Trim run.
- **Allow automatic File System Trim:** Enable or disable periodic File System Trim runs in the background, according to the configured schedule. Select one of:
 - **Weekly** (default):
 - * **Day of Week:** From the dropdown, select a day of the week.
Sunday - Saturday (default).
 - * **Hour:** From the dropdown, select the time of day that the File System Trim will run.
00:00 - 23:00, on the hour only. Default: 01:00.



Note: The default schedules for Defragmentation and File System Trim vary in both day and time, to distribute system overhead more efficiently.

- **Monthly:**

- * **Day of month:** From the dropdown, select a day of the month.
The range is restricted to 1 - 28, to avoid discrepancies between longer and shorter months.
 - * **Hour:** From the dropdown, select the time of day that the File System Trim will run.
00:00 - 23:00, on the hour only.
1. Optionally, click **Default** to reset the settings to their defaults.
 2. Click **Submit** to save the changes, or **Cancel** to exit without applying changes to the settings.

26.3.7 Default Filesystem Write Policy

Set the default write policy for new volumes (can be applied for existing volumes in a bulk operation as well). The Write Policy refers to different ways in which data is written to the underlying VP SA volume filesystem during filesystem operations.

- **Asynchronous Writing (default)** When the filesystem is mounted with the “Asynchronous Writing” option, data modifications are not immediately synchronized with the volume file-system. Instead, the system buffers these changes in memory and may perform the actual write to the filesystem at a later time. This can lead to faster write performance as the system doesn’t need to wait for each individual write to complete before proceeding with other tasks.
- **Synchronous Writing** When a filesystem is mounted with the “Synchronous Writing” option, all data modifications (writes) are immediately synchronized with the storage device. This means that before a write operation is considered complete, the data is physically written to the VP SA filesystem, ensuring that changes are safely stored on stable storage. This can ensure data integrity but can also lead to slower write performance, as the system waits for the filesystem to confirm the write operation before proceeding.

26.4 Metering

The VP SA provides an option to download its performance metering database which contains per-minute performance statistics about all active and monitored Objects – Drives, RAID Groups, Pools, Volumes and Servers. The database is downloaded in a binary format and is accompanied with a tool (meter2csv) to convert the raw binary database to a csv formatted file.

26.5 Container Service

26.5.1 Container Network

This displays the internal IP range of the ZCS and is accessible only by the host VP SA.

26.5.2 Exposed Ports

This displays the Ports ranges that are exposed for host access.

26.5.3 Image Repository

This displays the Status, Pool and Capacity of the Image Repository that stores all of the ZCS containers and images.

26.6 Network

Maximum Transmission Unit (MTU) is the largest size, in bytes, of a single packet that can be sent over a network interface without fragmentation.

Important: Best Practice

While the FE and Public networks may use different MTU values from each other, all components within each network must use the same MTU setting.

- For the **FE Network**, configure an identical MTU across all its connected zStorage components — VPSAs, Object Storages, Storage Nodes, and any associated customer-side hosts or switches.
- For the **Public Network**, also ensure a consistent MTU across all of its connected components.

Inconsistent MTU settings across zStorage components for a network can cause packet fragmentation, increased latency, or dropped connections, which could impact performance and data availability.


The MTU size can be set by selecting values from 1500 (default) to 9000 from the dropdown for both the Front End (data) network, and the public network & VNIs (Public & VNI network MTU).

The VPSA also supports customizing the Front End Network MTU, by overwriting the displayed MTU value with a custom MTU size in the range of 1420 - 9000.


✓ **Note:** Changing the MTU can be disruptive for ongoing traffic. Existing iSCSI server sessions may require a restart for the new MTU setting to take effect.

26.7 File Lifecycle Management

The VPSA supports file lifecycle management and analytics. When file lifecycle management is enabled for a VPSA, a toggle is provided to pause or resume global file lifecycle management indexing on volumes.



In the File Lifecycle Management tab on the Settings page, click **Edit**.



In the dialog that opens, click **Pause** or **Resume**, according to the context.

✓ **Note:** To disable file lifecycle management globally for the VPSA, use the VPSA's **Disable File Lifecycle** operation in the **Provisioning Portal**.

DIAGNOSTICS

27.1 Network Diagnostics

A common issue storage administrators face is the ability to verify connectivity between the storage system and the servers using the storage. Connectivity might be even harder to verify in a cloud environment, depending on the network topology. VPSA Network Diagnostics allows you to check connectivity over the selected network to any server:

1. Go to the **System > Diagnostics**.
2. Select the VPSA network **Interface**.
3. Select the **Diagnostics type**:
 - **Reachability**
 1. Enter the **Target IP Address** of the server in question.
 2. Select whether you want to do one or both of:
 - **Ping** the server, and set the **Count** of echo requests.
 - Run **Traceroute**, and set the number of **TTL** hops.
 3. Click **Run**.

It takes about a minute until the results display in the output frame.

- **Packet capture (tcpdump)**

Click **Run**.

Packet capture (tcpdump) runs as a background task. Runtime is limited to about 2 minutes and the size of the output file is limited to 200MB. After tcpdump completion you will be able to download the output file.

LOGS

28.1 Access Log

Access log lists all operations done by any user, either using the GUI or the REST API. Each operation is listed with all given parameters.

The list can be filtered by:

- User who took the action
- Action type (e.g. create account)
- Date and time

28.2 Events Log

The events log lists all the events reported by the system. The list can be filtered by:


- Message contains - lists only events that contain the given string
- Min severity - lists only events at the given severity level and more severe

The **Advanced Mode** provides additional options:

- Doesn't contain - excludes events in the given string
- Date - lists events in a date and time range
- Source Type - lists events from a selected source type, e.g. Disk, Volume, Server, Controller

SUPPORT

You can manage your Zadara Tech Support tickets directly from your VPSA. Support requests are redirected to the Zadara Support portal at <https://support.zadarastorage.com/home>.



To Open a Support Ticket

- Open the VPSA GUI > Support > Tickets page and click Create.
- Enter the Subject and Description and press Create Ticket.
- Select if the ticket should include a full set of logs (ZSnap) or not.
- This creates a ticket along with (or without) a set of logs (ZSnap), and is uploaded to the Zadara portal for analysis of the issue.

A screenshot of a "Create Support Ticket" dialog box. It contains three input fields: "Subject: *" with the text "Help! Can't connect", "Description: *" with the text "Something went wrong...|", and "Create zsnap:" with a checked checkbox. At the bottom right, there are two buttons: "Create Ticket" (highlighted in green) and "Cancel".

To Manage Support Tickets

- You can view the list of open support tickets, with each ticket displaying its ticket number, date, status, and subject per ticket.
- You can Comment on a ticket or Add Zsnap to an existing ticket.
- Finally, if you feel an issue is resolved you can close it.

CSI DRIVER

30.1 Zadara VPSA CSI for Kubernetes

The VPSA supports Persistent [Kubernetes Container Storage Interface \(CSI\)](#). CSI is a standard for exposing arbitrary block storage to containerized workloads on Container Orchestration Systems (COS) like Kubernetes.

This gives Kubernetes users the ability to use Zadara block and file volumes as persistent storage for containers and makes the system more secure and reliable.

The Zadara VPSA CSI provider implements an interface between the Container Storage Interface (CSI) and Zadara VPSA Storage Array and VPSA Flash Array, for a dynamic provisioning of persistent Block and File volumes.

For further information and implementation examples, see the [Zadara CSI repository on Github](#).