



VPSA Storage Array User Guide

Release 21.07

Zadara

Jan 18, 2022

CONTENTS

1	Introduction	3
2	Getting Started	7
3	Managing RAID Groups and Drives	27
4	Configuring Storage Pools	35
5	Understanding Controllers	55
6	Managing Servers	59
7	Managing Volumes, Snapshots and Clones	77
8	Monitoring Performance	113
9	Managing Access Control	117
10	Managing Remote Mirroring	139
11	Managing Remote Clones	153
12	Backup to Object Storage	159
13	Managing Container Services	169
14	Settings	177
15	Diagnostics	183
16	Managing Tech Support Tickets	185

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

Explore our user guide to provision and administrate your VPSA Object Storage.

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 infrustructure: 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 Storage Array and the Flash Array. Sections that are specific to one offering only are marked as follows:

VPSA STORAGE ARRAY

VPSA FLASH ARRAY


Zadara VPSA Storage Array provides:

- Unified storage with both Block volumes and File shares exposure
- Data Protection (RAID-1, 6, 10, 60)
- 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.

GETTING STARTED

This chapter contains step-by-step instructions to create a VPSA and then to configure its storage properties.

2.1 Registering a Zadara Account & Creating a VPSA

- Go to your organizational/cloud service provider **VPSA Provisioning Portal**
- Click on Create account for first time user registration



- Complete the form to create a Zadara Account.
- go to your **VPSA Provisioning Portal**.
- Login using your username / email and password, then press **Create Zadara Storage VPSA®** to create a new VPSA.

Name	Type	Engine	Drives	Status	Provider	URL	Version
Z_20_01_268_R	Object Storage	Mini	2	Created	zadara-qa4	https://vsa-00000147-zadara-qa4.zadara.com/8443	zios-20.01-268.img
Z1908R	Object Storage	-	2	Failed	zadara-qa4		zios-19.08-205.img
Z_TEST_1908_205_R	Object Storage	-	2	Waiting For Approval	zadara-qa4		-
e_R1908_205R	Storage Array	200	2	Created	zadara-qa4	https://vsa-00000145-zadara-qa4.zadara.com	vc-19.08-205.img
Z1908_205R	Object Storage	-	2	Failed	zadara-qa4		zios-19.08-205.img

✓ **Note:** Zadara has a public cloud Provisioning portal. User registration can be performed at: <https://manage.zadarastorage.com/register/>.

- The following dialog will appear:

Select Product

Type
 Name and Region
 Engine and Drives
 Advanced Services
 Review

Creating a new Virtual Private Storage Array (VPSA) takes just a few minutes. Please provide the information requested at each step and watch as your VPSA takes shape in the right-hand column.

Zadara VPSA

Storage Array
 All Flash Array
 Object Storage

Data Services

Backup

Choose... ▾

Your Configuration ✕

Product

Description

Select either Storage Array to create a hybrid storage array or Flash Array to create a performance optimized flash array with built in inline data reduction support.

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

- The VPSA definitions dialog will appear

Create new VPSA® Storage Array

Creating a new Virtual Private Storage Array (VPSA) takes just a few minutes. Please provide the information requested at each step and watch as your VPSA takes shape in the right-hand column.

VPSA Name *

VPSA Description

Cloud Provider

zadara-qa4

Select Region

Your Configuration

Product
Storage Array

VPSA Name
MyVirtualArray

Region
zadara-qa4

Enter the following mandatory fields:

- **VPSA Name** – Give the VPSA a name. This is how it will appear in the Cloud Console and in the VPSA GUI. If you are planning on having multiple VPSAs, you might want to give it as detailed a name as possible.
- **VPSA Description** – Give the VPSA a description.
- **Select Cloud Provider** – Select the Cloud or Co-lo where you have your compute instances. VPSAs are able to simultaneously connect to multiple Cloud Providers and Co-locations (within the same geographical region).

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

- **Select a Region** – Select the Cloud Provider region where your application servers reside. The servers and the VPSA must reside in the same region in order to establish efficient iSCSI or NFS\CIFS connectivity. Available Regions depend on which Cloud Provider you select.
- **Protection Zone** – VPSA supports multiple protection zones in “stretched cluster” configurations, where each VC is in a different zone. In cloud locations that provide protection zones, select in which zone the new VPSA will be built.

For multi-zone configurations select Multiple.

- Click Next

Create New VPSA® Storage Array

Progress: Name And Region | **Engine and Drives** | Advanced Services | Review

Zadara IO Engine
400 - 4 CPUs, 12GB RAM, 20GB Cache (Max. 10 drives) (\$1.99/hr)

Extended Flash Cache - 200GB
\$0.2 Per 200GB per hour

Drive Quantities (Minimum of 2)

10	4656GB (4546GiB) SATA 4656GB 7200RPM (\$0.2/hr)
0	2793GB (2727GiB) SATA 2793GB 7200RPM (\$50.0/hr)
0	2793GB (2727GiB) SATA 2793GB 5940RPM (\$50.0/hr)
0	185GB (180GiB) SSD 185GB 1RPM (\$50.0/hr)
0	278GB (271GiB) SAS 278GB 15000RPM (\$50.0/hr)

Your Configuration

VPSA Name: MyVirtualArray1
Region: US_WEST

Zadara IO Engine: 400
Extended Flash Cache: 200 GB
Drive Quantities: 10 SATA 4656GB 7200RPM

\$4.19
Price Per Hour

Back Next

- **Select the Zadara IO Engine** - 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.
 - **VPSA STORAGE ARRAY** 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 drive 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, providing 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.

- **VPSA STORAGE ARRAY** **Select the Cache size** (for Storage Array engines larger than 200) Use the slide bar to set the amount of flash cache allocated to this VPSA. Note that each VPSA engine comes with a minimum amount of cache. The extended cache is allocated in 200GB increments.
- **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 max 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 the total capacity of all selected drives is within the limit.

VPSA STORAGE ARRAY The following table lists the maximum drives per Storage Array Engine type:

IO Engine Type	Maximum # of Drives	Maximum Raw Capacity
200 (Baby)	5	24 TB
400 (Basic)	10	70 TB
600 (Boost)	20	140 TB
800 (Blast)	30	180 TB
1000 (Blazing)	40	240 TB
1200	60	300 TB
1600	80	400 TB
2400	80	800 TB
3600	80	1000 TB

VPSA FLASH ARRAY The following table lists the maximum drives and capacity per All Flash Array Engine type: **Note that for All Flash Array, due to data reduction, the capacity limit per engine depends on both the physical capacity of the drives and the the customer virtual capacity (as seen by the hosts), before any data reduction.** More about All Flash Array capacities: [Understanding Pool's Capacity](#)

IO Engine Type	Maximum # of Drives	Maximum Raw Capacity	Maximum Customer (host) Capacity
H100	60	280 TB	140 TB
H200	80	440 TB	220 TB
H300	120	800 TB	400 TB
H400	140	1000 TB	500 TB

✓ **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

- Click Next

- **Select the Zadara Container Services (ZCS) Engine** – The Zadara ZCS Engine defines the compute resources of the VPSA’s Virtual Controllers that are allocated for Docker containers within this VPSA. Refer to [Managing Container Services](#) for details about Zadara Container Services.
- **Fibre Channel Support** – Check this checkbox if you will be connecting hosts to this VPSA over FC SAN.

- Click Next

Create New VPSA® Storage Array

○
○
○
●

Name And Region
Engine and Drives
Advanced Services
Review

Name And Region

VPSA Name: **MyVirtualArray1**

Region: **QA3**

[Edit](#)

Engine and Drives

Zadara IO Engine: **400**

Extended Flash Cache: **200 GB**

Drive Quantities: **10 SATA 4656GB 7200RPM**

[Edit](#)

Advanced Services

Zadara Container Services Engine: **01 - 1 CPUs, 256MB RAM**

[Edit](#)

Options

Create RAID-10 Pool(s)

Moderate Request

\$6.19

Calculated VPSA Price Per Hour

Back
Create

Next Steps

After you've created your VPSA, our operations team will perform a quick review of your custom configuration and contact you to schedule an onboarding.

If you have any questions about provisioning your VPSA or the onboarding process, please contact our world-class support team.

Email support@zadarastorage.com

Phone 1-949-284-0713

- Create RAID-10 Pool(s)** – By default, at VPSA creation time RAID-10 pools are automatically created. One pool per each drives type selected. All the drives selected Of each type will be included in the pool. If you want to create different pools setting, uncheck this checkbox, and manually create your RAID groups and pools as described below.
- Once you have completed selecting the above VPSA characteristics, review the displayed summary. You can click Edit to modify your previous selections. Press the Create button to confirm the VPSA creation request. The requested VPSA will appear in the “Awaiting Approval” list.

oded@zadarastorage.com EN

Create Zadara Storage VPSA®

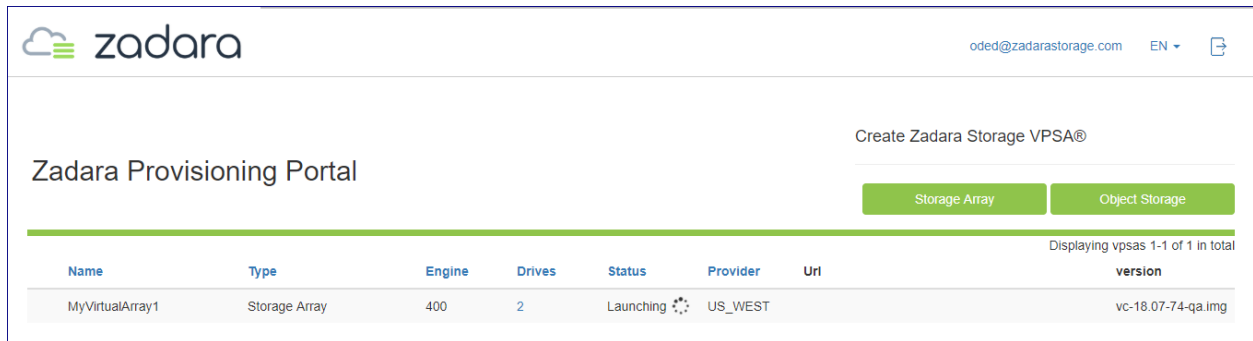
Storage Array
Object Storage

Displaying vpsas 1-1 of 1 in total

Name	Type	Engine	Drives	Status	Provider	Uri	version
MyVirtualArray1	Storage Array	400	4	Waiting For Approval	US_WEST		-

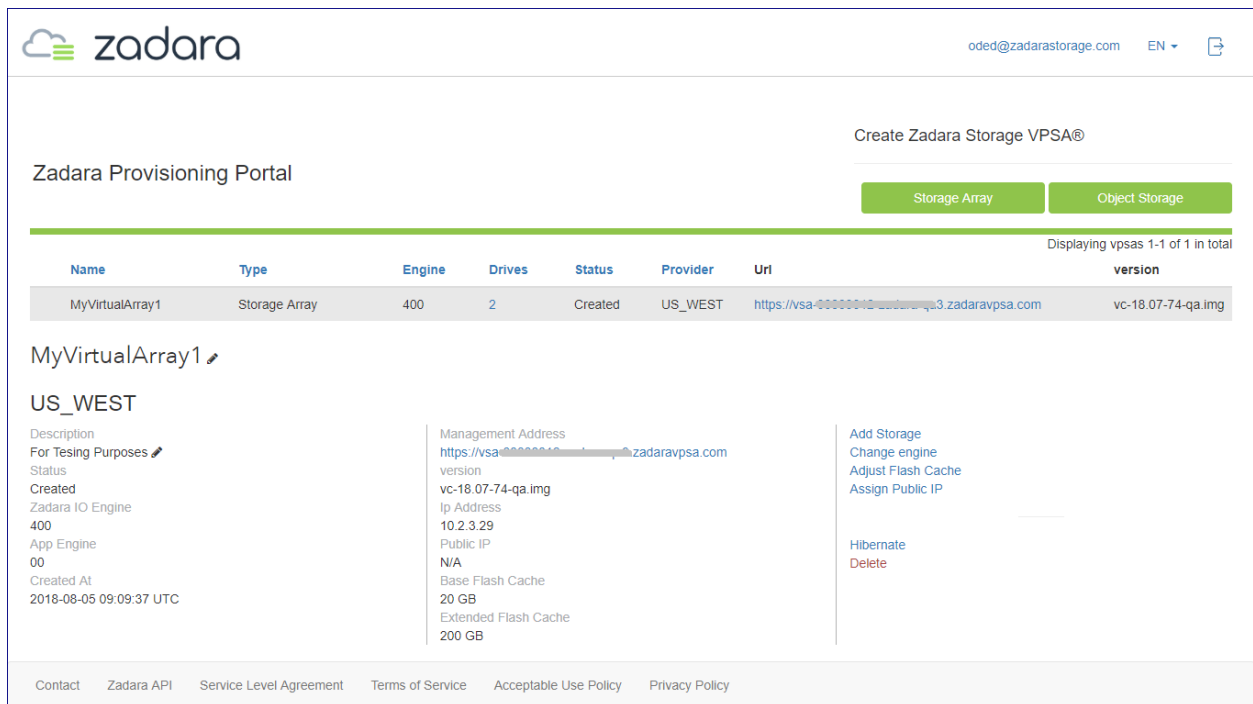
- Completing the VPSA creation requires the approval of a Zadara Storage Cloud admin. Once approved, the new

VPSA only takes a few minutes to launch. During that time you'll see your VPSA in “Launching” status as shown below:



The screenshot shows the Zadara Provisioning Portal interface. At the top, there is a navigation bar with the Zadara logo, the user email 'oded@zadarastorage.com', and a language dropdown set to 'EN'. Below the navigation bar, there is a section titled 'Create Zadara Storage VPSA®' with two buttons: 'Storage Array' (highlighted in green) and 'Object Storage'. Below this is a table with columns: Name, Type, Engine, Drives, Status, Provider, Uri, and version. The table displays one entry: 'MyVirtualArray1' with Type 'Storage Array', Engine '400', Drives '2', Status 'Launching', Provider 'US_WEST', and version 'vc-18.07-74-qa.img'. A small 'Launching' icon with a clock is next to the status.

- Once the VPSA is ready, you'll receive an email with a temporary passcode at your registered email address.
- Use the “Management Address” link to access the VPSA GUI:



The screenshot shows the Zadara Provisioning Portal interface. At the top, there is a navigation bar with the Zadara logo, the user email 'oded@zadarastorage.com', and a language dropdown set to 'EN'. Below the navigation bar, there is a section titled 'Create Zadara Storage VPSA®' with two buttons: 'Storage Array' (highlighted in green) and 'Object Storage'. Below this is a table with columns: Name, Type, Engine, Drives, Status, Provider, Uri, and version. The table displays one entry: 'MyVirtualArray1' with Type 'Storage Array', Engine '400', Drives '2', Status 'Created', Provider 'US_WEST', and version 'vc-18.07-74-qa.img'. Below the table, there is a detailed view for 'MyVirtualArray1' in the 'US_WEST' region. The view includes a description 'For Testing Purposes', status 'Created', and creation date '2018-08-05 09:09:37 UTC'. It also lists configuration details: Management Address 'https://vsa-00000000000000000000000000000000.zadaravpsa.com', version 'vc-18.07-74-qa.img', Ip Address '10.2.3.29', Public IP 'N/A', Base Flash Cache '20 GB', and Extended Flash Cache '200 GB'. On the right side, there are several action links: 'Add Storage', 'Change engine', 'Adjust Flash Cache', 'Assign Public IP', 'Hibernate', and 'Delete'. At the bottom of the page, there is a footer with links for 'Contact', 'Zadara API', 'Service Level Agreement', 'Terms of Service', 'Acceptable Use Policy', and 'Privacy Policy'.

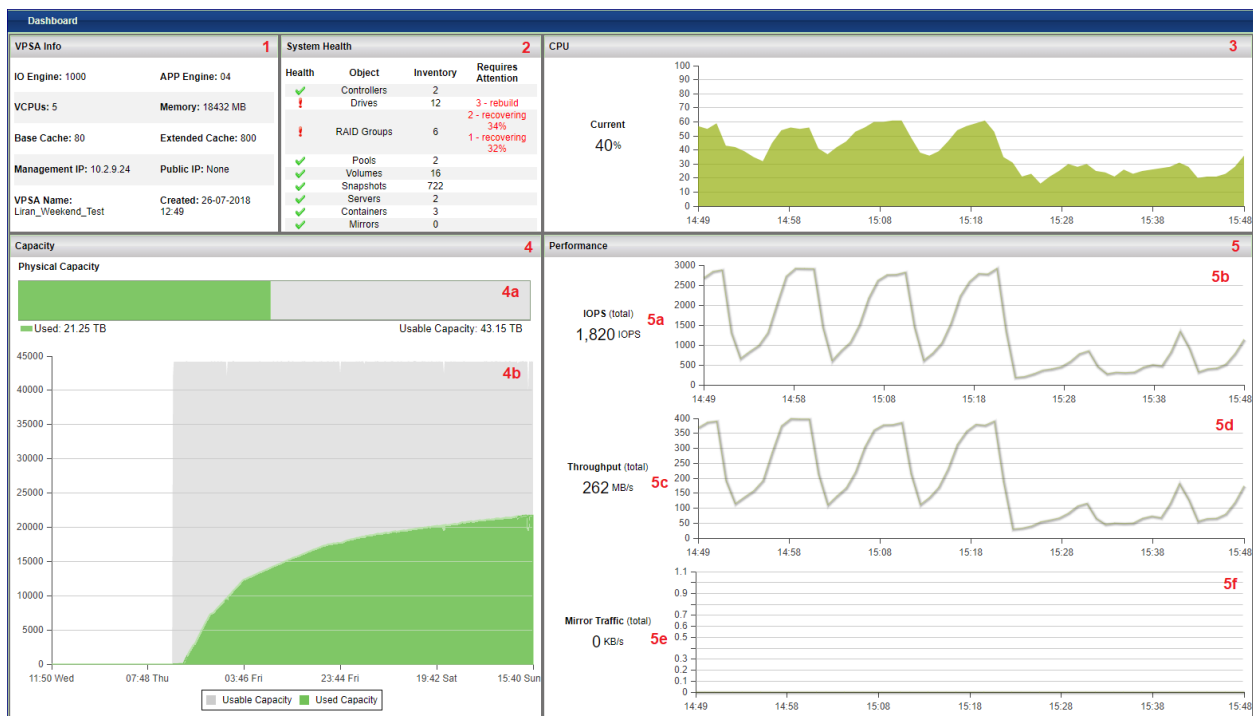
- Use your registered username or email address and the temporary passcode to enter the VPSA GUI. You will be immediately prompted to set a new password for your VPSA User account.

2.2 The VPSA Interface

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

2.2.1 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 is made of the following components:



- 1. VPSA Info:** General information of the VPSA such as 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 etc... If all objects are in “normal” state, there is a green checkmark on the line. If there is situation that needs your attention a red exclamation mark is shown on that specific line, with number of objects that need to be taken care of.
- 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 Storage Array and All Flash Array. For the later it shows the capacity reduction saving. See [Understanding Pool's Capacity](#) for details.
 - Current capacity state
 - 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

2.2.2 Understanding the VPSA GUI

The screenshot displays the VPSA GUI interface. At the top, there is a status bar showing the language as English (5), the display timezone as Asia/Jerusalem (+03:00), and the user as administrator (4). The left navigation panel (1) lists various resources like RAID Groups, Pools, Volumes, Servers, and Controllers. The main content area (2) shows a table of drives with columns for Name, Capacity, Storage Node, Type, Status, RAID Group, and Zone. The 'drive-000' row is highlighted. Below the table, there is a details pane (3) for 'drive-000' with tabs for Properties, Metering, and Logs. The Properties tab is active, showing fields for ID, Name, Capacity, Storage Node, Type, Status, RAID Group, Usage, Added, and Modified.

Name	Capacity	Storage Node	Type	Status	RAID Group	Zone
drive-000	279 GB	qa2-sn1	SAS	Normal	rg1	zone_0
drive-001	279 GB	qa2-sn3	SAS	Normal	rg1	zone_0
drive-002	279 GB	qa2-sn2	SAS	Normal	rg2	zone_0
drive-003	279 GB	qa2-sn1	SAS	Normal	rg2	zone_0

Details for drive-000:

- ID: volume-000000ca
- Name: drive-000
- Capacity: 279 GB
- Storage Node: qa2-sn1
- Type: SAS
- Status: Normal
- RAID Group: rg1
- Usage: In-use
- Added: 2016-03-24 16:24:45
- Modified: 2016-03-27 12:47:25

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 at the top left corner. You can use this drop down to change the displayed language.

✓ **Note:**

- From version 20.12 TLS v1.1 is no longer supported by the VPSA GUI.

2.3 Creating RAID Group, Pools, and Volumes

By default a new VPSA 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.

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

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

<table border="1" class="docutils"> <thead> <tr> <th></th> </tr> </thead> <tbody> <tr> <td></td> </tr> </tbody> </table>

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

2.4 Provisioning your VPSA

You create, add, change, delete and manage the resources composing your VPSAs via **the Zadara Provisioning Portal**.

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

2.5 Adding and removing Disk Drives

To add drives to your VPSA go to the Provisioning Portal, select the VPSA and then press the Add Storage button.

Add Storage to MyVirtualArray1

Select drive quantities

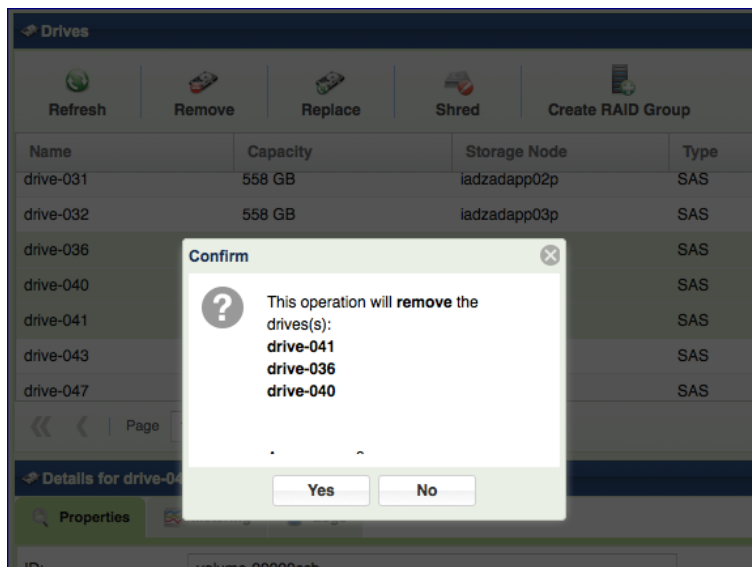
<input type="text" value="0"/>	SAS 278GB 15000RPM SAS 278GB 15000RPM (\$50.0/hr)
<input type="text" value="0"/>	SATA 2793GB 7200RPM SATA 2793GB 7200RPM (\$2.0/hr)
<input style="border: 2px solid #00aaff;" type="text" value="10"/>	SATA 5000GB 7200RPM SATA 5000GB 7200RPM (\$0.2/hr)
<input type="text" value="0"/>	SATA 3000GB 5940RPM SATA 3000GB 5940RPM (\$0.13/hr)
<input type="text" value="0"/>	SSD 185GB 1RPM SSD 185GB 1RPM (\$50.0/hr)

Moderate Request?

- Select the number of drives from each available drive type you wish to add to your VPSA, and press Submit. 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 with status "Available") from within the VPSA.

Go to the [VPSA GUI > Drives](#), select the Drive you wish to remove and press the Remove button:



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 here: [Replacing a Drive](#)

2.6 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.
- **VPSA STORAGE ARRAY Flash Cache Size** - 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 comes with a base SSD cache partition size. You can request additional flash capacity for caching. For more details see “Adjusting Cache.”
- **Maximum number of drives** - The maximum number of drives that can be allocated to each VPSA engine type.

VPSA STORAGE ARRAY The following Zadara IO Engines are available For Storage Array:

IO Engine Type	Dedicated Compute Resources	Base Flash Cache	Max # of Drives	Max Raw Capacity
200 (Baby)	2 CPU, 6 GB RAM	20 GB	5	24 TB
400 (Basic)	4 CPU, 12 GB RAM	20 GB	10	70 TB
600 (Boost)	6 CPU, 20 GB RAM	40 GB	20	140 TB
800 (Blast)	8 CPU, 28 GB RAM	60 GB	30	180 TB
1000 (Blazing)	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

VPSA FLASH ARRAY The following Zadara IO Engines are available For All Flash Array: **Note that for All Flash Array, due to data reduction, the capacity limit per engine depends on both the physical capacity of the drives and the the customer virtual capacity (as seen by the hosts), before any data reduction.** More about All Flash Array capacities: [Understanding Pool's Capacity](#)

IO Engine Type	Dedicated Compute Resources	Maximum # of Drives	Maximum Raw Capacity	Maximum Customer (Host) 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:** 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

The following Zadara Container Services Engines (see: [Managing Container Services](#)) 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, press the Change Engine link in the Zadara Provisioning Portal:

Upgrade Engine MyVirtualArray1

Zadara IO Engine (current: 400)

1200 - 12 CPUs, 52GB RAM, 100GB Cache (Max. 60 drives) (\$4.99/hr) ▼

Zadara App Engine (current: N/A)

04 - 2 CPUs, 1GB RAM (\$4.0/hr) ▼

Moderate Request?

Cancel
Change

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

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, your VPSA status will change to "Upgrade Pending".


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

2.7 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. click Network Management in the Zadara Provisioning Portal’s Management Menu.



Support Management **admin**

My Profile
Network Management

Create Zadara Storage VPSA®

Zadara Provisioning Portal

All Flash Array

Storage Array

Object Storage

To create another virtual network press Create , 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 VP SA. The VP SA internally maintains a “Virtual Networks Interface” (VNI) that connects into each virtual networks.

To add Virtual Network Interface, press the Add Virtual Network Interface link in the Zadara Provisioning Portal VP SA Operations.

To remove Virtual Network Interface, press the Remove Virtual Network Interface link in the Zadara Provisioning Portal VP SA Operations.

zadara-qa10

<p>Description None </p> <p>Status</p> <p>Created</p> <p>Zadara IO Engine 800</p> <p>App Engine 04</p> <p>Created At 2019-09-16 06:02:06 UTC</p> <p>Version vc-19.08-168-qa.img</p>	<p>Primary Management Interface</p> <p>URL https://vsa-00000001-zadara-qa10.zadara.com</p> <p>IPv4 10.2.10.45</p> <p>IPv6 N/A</p> <p>Public IP N/A</p> <p>Base Flash Cache 60 GB</p> <p>Extended Flash Cache 0 GB</p>	<p>VP SA Operations</p> <p>Add Storage</p> <p>Add Proxy Virtual Controller</p> <p>Change Engine</p> <p>Adjust Flash Cache</p> <p>Assign Public IP</p> <p>Add Virtual Network Interface</p> <p>Hibernate</p> <p>Delete</p> <p>Global Operations</p> <p>Network Management</p>
---	---	--

Note:

- Number of VNIs per VP SA is limited to 5.
- VP SA REST API/GUI is accessible through any VNI.
- Only Primary VN IP is registered in DNSimple
- VP SA 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”.

2.8 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, simply click the Remove Public IP button 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.

2.9 **VPSA STORAGE ARRAY** Adjusting Flash Cache


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 (Baby)	20 GB	0 GB	0 GB
400 (Basic)	20 GB	200 GB	400 GB
600 (Boost)	40 GB	400 GB	800 GB
800 (Blast)	60 GB	600 GB	1200 GB
1000 (Blazing)	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, go to the Provisioning Portal and press the Adjust Flash Cache link:



2.10 Hibernating your VPSA

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

To hibernate a VPSA, go to the VPSA Provisioning Portal and press the Hibernate link:

Zadara Provisioning Portal

Create Zadara Storage VPSA®

Storage Array | Object Storage

Displaying vpsas 1-1 of 1 in total

Name	Type	Engine	Drives	Status	Provider	Uri	version
MyVirtualArray1	Storage Array	400	2	Created	US_WEST	https://vsa-00000000-zadara-vpsa.com	vc-18.07-74-qa.img

MyVirtualArray1 [✎](#)

US_WEST

Description

For Testing Purposes [✎](#)

Status

Created

Zadara IO Engine

400

App Engine

00

Created At

2018-08-05 09:09:37 UTC

Management Address

<https://vsa-00000000-zadara-vpsa.com>

version

vc-18.07-74-qa.img

Ip Address

10.0.0.0

Public IP

N/A

Base Flash Cache

20 GB

Extended Flash Cache

200 GB

Add Storage

Change engine

Adjust Flash Cache

Assign Public IP

Hibernate


Delete

Contact | Zadara API | Service Level Agreement | Terms of Service | Acceptable Use Policy | Privacy Policy

To resume access to the VPSA, go to the Provisioning Portal and press the Restore link. (The Hibernate and Restore toggle depending on the current state of the VPSA.)

MANAGING RAID GROUPS AND DRIVES

3.1 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 (RAID-6 allows for a 2-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 Stripe Size. Choosing a Stripe Size other than the default value of 64K is only applicable to RAID-6 and depends on the performance needs specific to your workload.
- Select whether to allocate a drive as a Hot Spare for this RAID group. Adding Spare drive for RAID-1 groups is recommended. See more details about managing **Hot Spares** [here](#).
- 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 RAID-6: in a VPSA Storage Array a minimum of 4 drives is required and a VPSA Flash Array supports 4+2 or 8+2 RAID group configurations.
- 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.
- RAID-5 is no longer supported.

Create RAID Group

Name: *

Protection: RAID1 RAID6

Stripe Size: 4KB 16KB 32KB 64KB 128KB 256KB

Hot Spare:

Drives: *

<input type="checkbox"/>	Name	Type	Capacity	Status	Storage Node	Zone
<input checked="" type="checkbox"/>	drive-002	SAS	558 GB	Available	iadzadapp...	zone_0
<input checked="" type="checkbox"/>	drive-006	SAS	558 GB	Available	iadzadapp...	zone_0
<input checked="" type="checkbox"/>	drive-010	SAS	558 GB	Available	iadzadapp...	zone_0
<input checked="" type="checkbox"/>	drive-011	SAS	558 GB	Available	iadzadapp...	zone_0
<input checked="" type="checkbox"/>	drive-012	SAS	558 GB	Available	iadzadapp...	zone_0
<input type="checkbox"/>	drive-016	SAS	558 GB	Available	iadzadapp...	zone_0
<input type="checkbox"/>	drive-020	SAS	558 GB	Available	iadzadapp...	zone_0
<input type="checkbox"/>	drive-021	SAS	558 GB	Available	iadzadapp...	zone_0
<input type="checkbox"/>	drive-022	SAS	558 GB	Available	iadzadann...	zone_0

3.1.1 Understanding RAID levels

The following RAID level 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-6	RAID-6 uses multiple parity sets to store data and can therefore tolerate up to 2 drive failures simultaneously. It requires a minimum of 4 drives. It offers the best data protection. Usable capacity is N-2 where N is the number of physical drives in the logical array.
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.
RAID-60	RAID-60 are achieved (similar to RAID-10) by creating RAID-6 RAID Groups and striping them together at the Pool level.

3.2 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 or RAID-6.
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.
Stripe Size	Stripe size (per drive) for RAID-6.
Mirror Number	Number of mirror copies for RAID-1.
Protection Width	Number of Drives participating in a RAID-6 RAID Group (including parity).
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.

3.3 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 > 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 here: [Replacing a Drive](#).

3.4 Managing RAID Group Sync Speed

RAID Group Sync Speed 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 I/Os. Therefore, the VPSA allows you to control two parameters impacting the sync Speed:

- **“Max Speed During Host I/Os”** – Controls the RAID sync speed when there are Server I/Os. You will want to set it low if the Server’s I/Os are the priority. Set it high if you want to prioritize the RAID rebuild process.
 - Default value: 10 MB/s
 - Range: 1 - 500 MB/s
- **“Max Speed w/o Host I/Os”** – Controls the sync speed when there are no Server I/Os. 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 I/Os).



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 (see below).

3.5 Understanding Media Scans

Media Scan is the process of checking RAID-6 parity integrity. It reads data and parity from all devices and automatically fixes any inconsistent parity.

This process runs automatically once a month in order to identify and handle any possible silent data integrity issues.

You may want to trigger a Media Scan on a RAID Group manually if, for example, there was an event that is suspected of putting data integrity at risk, such as the failure of two or more drives in a RAID-6 group.

The RAID status will change to “Repairing X%” during the Media Scan. At the end of the Media Scan the results are saved in an event-log message.

You cannot abort a Media Scan that’s in process, but you can pause it by pressing the Pause Media Scan button. The Media Scan button toggles to Pause Media Scan for RAID Groups that are currently being scanned. The RAID status will change to “Repairing Paused” when the Media Scan is paused.

3.6 Force Recovery

You can only issue Force Recovery on failed RAID-6 RAID Groups, after one (or more) of the failed drives has been recovered.


If all the drives were recovered, the VPSA will have enough information to determine how to recover the RAID automatically. If two drives are permanently gone in a RAID-6 RAID Group, the VPSA will be unable to determine if the available drives contain the most up-to-date data and hence will be unable to safely decide to automatically recover the RAID Group.

You can instruct the VPSA to perform a “Force Recovery” of the RAID Group which marks all drives as consistent and in-sync and moves the RAID to Normal state.

It is recommended that you run Media Scan following Force Recovery, which will ensure RAID parity consistency (although data may still be inconsistent from the application perspective).

⚠ Caution: This operation may result in application data loss. It must be used only when drives are permanently lost and when there are no other alternatives to recover the data.

3.7 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.

3.8 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!**

3.9 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 VPSA 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 VPSA 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 VPSA.
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.

CONFIGURING STORAGE POOLS

4.1 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.

4.2 **VPSA** **FLASH ARRAY** Understanding Pool's Capacity

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

Capacity metrics to consider:

4.2.1 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

4.2.2 Virtual View

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

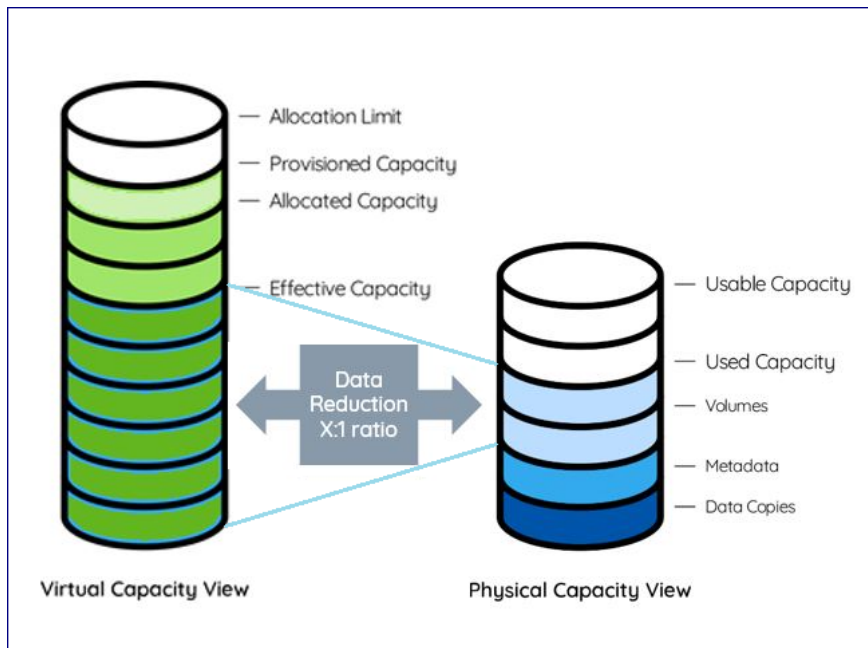
Allocated Capacity - Pool's allocated address space of all Volumes, Snapshots and Clones

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

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



4.2.3 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 / Data Reduction Ratio)

e.g.

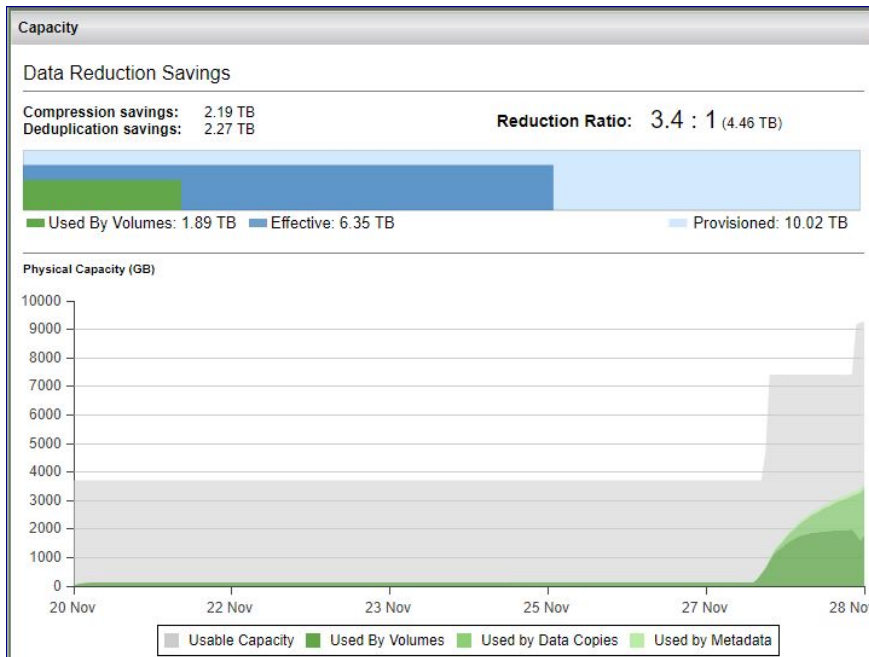
Data reduction ratio 2:1 , Data Reduction Percentage 50%

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

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

4.2.4 Pool capacity Monitoring

The All Flash VPSA **Dashboard** show the capacity consumption and data reduction saving



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** shows 2 bars.

The physical capacity bar shows the usable vs. used capacities.

The virtual capacity bar shows the allocated capacity vs. the allocation limit.

Physical Capacity	Virtual Capacity
42.57 TB Free / 52.09 TB	108.34 TB Free / 200 TB

<table border="1" class="docutils"> <thead> <tr> <th></th> </tr> </thead> <tbody> <tr> <td></td> </tr> </tbody> </table>

4.3 Creating and Managing Pools

4.3.1 Creating a Pool

✓ **Note:** By default when a new VPSA is created, a default pool is automatically created for each type of drives 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.

- a. Create a Pool from RAID Groups
- b. 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

To create a Pool from RAID Groups, you will see the following dialog:

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 (GB): 5000 Calculate Max

Type: Repository Storage
 Transactional Workloads
 Archival 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** – The VPSA supports Transactional, Repository and Archive Pool types. These Pool types use different chunk sizes for the mapping of virtual LBAs to Physical Drive addresses. The following table describes the tradeoffs for each type and the recommended use cases:



Storage Array Pool types:

	Transactional Pool	Repository Pool	Archive Pool
Chunk size	256KB	1MB	2MB
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
Cons	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)
Use Case	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 (> 100TB)
Limit	Transactional Pools have a maximum size of 20TB	<ul style="list-style-type: none"> • Repository Pools have a maximum size of 100TB 	<ul style="list-style-type: none"> • Archive Pools have a maximum size of 200TB



All Flash Array Pool types:

	Transactional Pool	Repository Pool	Archive Pool
Thin Provision Chunk size	512KB	1MB	2MB
Deduplication Chunk size	8KB	16KB	32KB
Pros	<ul style="list-style-type: none"> • Faster COW operation • Space efficiency on Random writes to Snapshots • Better Deduplication 	<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 • Better sequential workload throughput • Better Compression ratio
Cons	<ul style="list-style-type: none"> • Increased metadata size • Lower throughput 	<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)
Use Case	<ul style="list-style-type: none"> • Transactional Workload with Snapshots • High IOPS • Database (OLTP) 	<ul style="list-style-type: none"> • Analytics • High Throughput • General Purpose • File System 	<ul style="list-style-type: none"> • Relatively static data • Archive type workloads • Sequential workloads like Video Streaming • Very large pools/volumes (> 200TB)
Limit	<ul style="list-style-type: none"> • Transactional Pools have a maximum size of 50TB 	<ul style="list-style-type: none"> • Repository Pools have a maximum size of 100TB 	<ul style="list-style-type: none"> • Archive Pools have a maximum size of 200TB

In case there are number of pools in a given VPSA, there is a limit to aggregated total size of all pools. The following table lists the capacity limits of **all** pools per each VPSA All Flash engine:

Pool Type Engine	F800	F1200	F2400	F3600	F4800
Transactional	20	40	60	80	100
Repository	40	60	100	160	200
Archival	60	100	160	200	250

- **Cached** – Check this box to use SSD to Cache Server’s reads and writes.
 - All Pools that are marked as “Cached” share the VPSA 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 or RAID-6 creates RAID-10 or RAID-60 configurations respectively. Use striping to improve performance of random workloads since the IOs will be distributed and all drives will share the workload.

VPSA FLASH ARRAY

All Flash Array pools are always striped, and this checkpoint is hidden.

To create a Pool from Drives, you will see the following dialog:

Create Pool

Name: * Pool2

Drives: *
Select even number of drives of the same type

<input checked="" type="checkbox"/>	Name	Type	Capacity	Status	Storage ...	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: Repository Storage
 Transactional Workloads
 Archival Storage

Cached:

[Use Raid Group Selection](#) Create Cancel

The parameters are same as above. Just Check the box(es) of drives that will be allocated for this Pool.

4.3.2 Expanding Pool Capacity

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

Expand Pool Pool1

Additional Capacity (GB): * 556

Raid group(s): *

<input checked="" type="checkbox"/>	Name	Protection	Status	Available
<input checked="" type="checkbox"/>	RG3	RAID1	Normal	556 GB

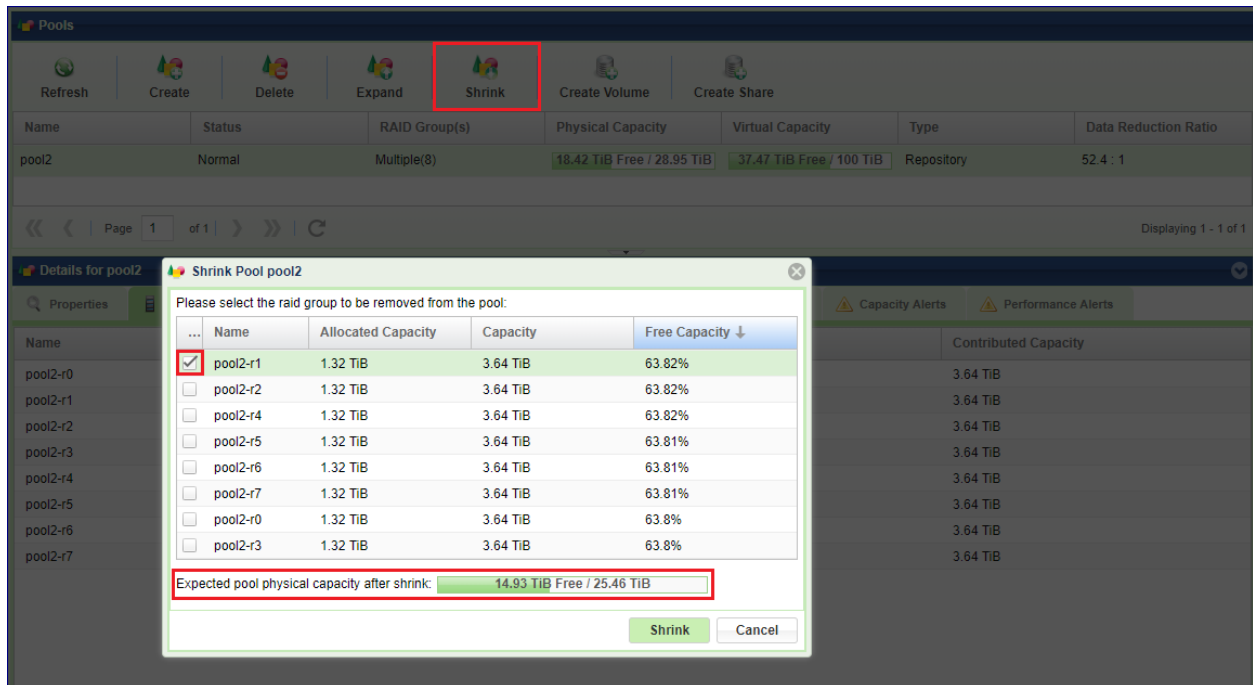
Expand Cancel

You can use capacity from any RAID Group to expand a Pool. If the RAID Group from which the new capacity is added doesn't match the protection type or drive type of the existing capacity you'll see a warning message pop up asking you to confirm the mismatch. Keep in mind that continuing with the mismatched types may impact the pool performance and protection QoS.

4.3.3 **VPSA FLASH ARRAY** Shrinking Pool Capacity

✓ **Note:** Pool shrink is only supported in All Flash VPSAs.

If the pool capacity is not fully used you can shrink it's size by removing RAID Group(s) (one 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.




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 Event once done.

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.

4.3.4 VPSA STORAGE ARRAY Disabling SSD cache on a pool



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.

4.4 Viewing Pool properties

The Pools details are shown in the following South Panel tabs:

Properties

VPSA STORAGE ARRAY Each Pool of Storage Array 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"> • 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.
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	Metadata Capacity
Capacity State	<ul style="list-style-type: none"> • Normal • Alert • Protected • Emergency See Managing Pool Capacity Alerts for more details.
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.
Type	<ul style="list-style-type: none"> • Transactional Workloads • Repository Storage • Archival Storage
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.
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. Disable only on rare cases where frequent snapshots cause extreme load of metadata operations. Consult Zадara support.
Raid Group(s)	RAID Group name, or "Multiple (X)" where X denotes the number of RAID Groups in the Pool.
Created	Date & time when the object was created.
Modified	Date & time when the object was last modified.

VPSA **FLASH ARRAY**

Each Pool of All Flash Array 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"> • 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.
Type	<ul style="list-style-type: none"> • Transactional Workloads • Repository Storage • Archival Storage
Raid Group(s)	RAID Group name, or “Multiple (X)” where X denotes the number of RAID Groups in the Pool.
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
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 See Managing Pool Capacity Alerts for more details.
	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
Effective Capacity	Amount of data written in the pool by all volumes and can be accessed by hosts. Not including space taken by snapshots
Virtual Capacity State	<ul style="list-style-type: none"> • Normal • Alert • Protected • Emergency See Managing Pool Capacity Alerts for more details.
	Capacity Savings
Data Reduction Ratio	Capacity savings by all data reduction techniques. Data Reduction Ratio = Effective Capacity / Used by Volumes
Deduplication 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

RAID Groups

In the [RAID Groups View](#) This tab lists the RAID Groups allocated to the selected Pool. Each RAID Group includes the following information:

- Name
- Protection (RAID-1, RAID-5or RAID-6)
- Status
- Contributed Capacity

In the [Segments View](#) This tab shows the structure of pool made of concatenated or striped segments

The screenshot displays the 'Pools' management interface. At the top, there are several action buttons: Refresh, Create, Delete, Expand, Disable Cache, Create Volume, and Create Share. Below these is a table listing the pools:

Name	Status	RAID Group(s)	Capacity	Type	Cached
repository pool	Normal	r1	390 GB Free / 489 GB	Repository Storage	Yes
pool1	Normal	Multiple(3)	10.59 TB Free / 16.26 TB	Transactional Workloads	Yes
pool2	Normal	Multiple(2)	1.52 TB Free / 5.35 TB	Archival Storage	Yes

Below the table is a pagination control showing 'Page 1 of 1' and 'Displaying 1 - 3 of 3'. The 'Details for pool1' section is expanded, showing the 'RAID Groups' tab. It displays a tree view of the pool's structure, totaling 16749GB (concatenation):

- RaidGroup-145 (rg1) 5583GB
- RaidGroup-146 (rg2) 5583GB
- RaidGroup-147 (rg3) 5583GB

At the bottom of the details panel, there are two tabs: 'Raid Groups View' and 'Segments View', with 'Segments View' currently selected.

Volumes and Dest Volumes

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. Each Volume includes the following information:

- Name
- Capacity (virtual, not provisioned)

- Status
- Data Type (Block or File-System)

Recycle Bin

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.

Logs

Displays all event logs associated with this Pool.

Metering

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 .

Capacity Alerts

The Capacity Alerts tab lists the configurable attributes of the Pool Protection Mechanism. See [Managing Pool Capacity Alerts](#) for more details. You can modify the following attributes:

- **Physical Pool Alert Mode Threshold** - “Alert me when it is estimated that the Pool will be at full physical capacity in X Minutes.”
 - Default Value: 360 minutes
- **Physical Pool Protection Mode 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.”

- Default Value: 60 minutes
- **Physical Pool Calculation Window** - “Calculate the estimated time until the Pool is full based on new capacity usage in the previous X minutes.”
 - Default Value: 60 minutes
- **Physical Pool Emergency Mode Threshold** - “Delete snapshots, starting from the oldest, when there is less than the following physical capacity left in the Pool”
 - Default Value: 50 GB

VPSA FLASH ARRAY

- **Allocated Capacity Alert Mode Threshold** - “Alert me when it is estimated that the Pool’s address space will be at full capacity in X Minutes.”
 - Default Value: 360 minutes
- **Allocated Capacity Protection Mode 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.”
 - Default Value: 60 minutes
- **Allocated Capacity Calculation Window** - “Calculate the estimated time until the Pool’s address space is full based on new capacity usage in the previous X minutes.”
 - Default Value: 60 minutes
- **Allocated Capacity Emergency Mode Threshold** - “Delete snapshots, starting from the oldest, when there is less than the following free address space left in the Pool”
 - Default Value: 5 GB

Performance Alerts

The Performance Alerts tab lists the Pool’s ability to send alerts when performance drops below expectations. See [Managing Pool Performance Alerts](#) for more details.

4.5 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.

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

- **Physical Pool Capacity Alert Threshold** – The estimated time (in minutes) before running out of free space or percentage used. 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.
 - Default time: 600 minutes (10 hours)
 - Minimum: 1 minute (0 means disable this alert by time)

or

- Default Percentage: 90% full
- Minimum: 1 % (0 means disable this alert by %)
- **Physical Pool Capacity Protection Threshold** – 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.
 - Default: 180 minutes (3 hour)
 - Minimum: 1 minute (0 means disable this alert by time)

or

- Default Percentage: 95% full
- Minimum: 1 % (0 means disable this alert by %)
- **Physical Pool Capacity Emergency Threshold** – When the Pool’s free capacity drops below this fixed threshold (in GB) or below the specified % threshold, 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.
 - Default: 50 GB
 - Minimum: 1 GB

or

- Default Percentage: 99% full
- Minimum: 1 % (0 means disable this alert by %)
- **Physical Pool Capacity Alert Interval** – 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 hours)
 - Minimum: 1 minute

Details for RAID-10-Pool-1

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

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

Physical	Virtual
<p>Physical Capacity Alert Threshold Alert me when:</p> <p>It is estimated that the Pool will be at full capacity in <input type="text" value="600 minutes"/> or The pool is <input type="text" value="90 % full"/></p> <p>Physical Capacity Protection Threshold Do not allow new Volumes, Shares, or Snapshots to be created when:</p> <p>It is estimated that the Pool will be at full capacity in <input type="text" value="180 minutes"/> or The pool is <input type="text" value="95 % full"/></p> <p>Physical Capacity Emergency Threshold Delete snapshots, starting from the oldest, when:</p> <p>There is less than the following capacity left in the Pool <input type="text" value="50 GB"/> or The pool is <input type="text" value="99 % full"/></p> <p>Physical Capacity Alert Interval Time estimations above are based on capacity usage during the last: <input type="text" value="60 minutes"/></p>	<p>Allocated Capacity Alert Threshold Alert me when it is estimated that the Pool Allocated Capacity will be at full capacity in: <input type="text" value="360 minutes"/></p> <p>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: <input type="text" value="60 minutes"/></p> <p>Allocated Capacity Emergency Threshold Delete snapshots, starting from the oldest, when there is less than the following Allocated capacity left in the Pool: <input type="text" value="5 GB"/></p> <p>Allocated Capacity Alert Interval Time estimations above are based on capacity usage during the last: <input type="text" value="60 minutes"/></p>

VPSA FLASH ARRAY

In addition to the physical capacity alerts, The All Flash VPSA provides alerts in case 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 Pool Protection mechanism:

- Allocated Capacity Alert Threshold** – 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** – 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 Alert Interval** – 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
- **Allocated Capacity Emergency Threshold** - When the Pool's free address space drops below this fixed threshold (in GB), 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 GB
 - Minimum: 1 GB

4.6 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.

UNDERSTANDING CONTROLLERS

Controller Objects in the VPSA represent the clustered virtual controller pair of a VPSA.

Controllers					
Refresh		Failover			
Name	Storage Node	Management IP	iSCSI IP	Status	Zone
vsa-00000017-vc-1	qa2-sn1	180.80.2.103		standby	zone_0
vsa-00000017-vc-0	qa2-sn3	180.80.2.102	180.80.2.102	active	zone_0

5.1 Failover

Failover halts the active controller and activates the standby controller with minimal I/O interruption.

5.2 Viewing Controller Properties

Properties

Each Controller displays the following Properties:

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

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.

Details for vsa-00000006-vc-0				
Properties Paths System Usage Cache Metering Logs Performance Alerts				
Initiator	Target	Number of sessions	Server name	Server ID
21000024ff7bc174	21efb07e7001a6ed	1	SERVER-204	srv-00000001
21000024ff7bc175	21efb07e7001a6ec	1	SERVER-204	srv-00000001
21000024ff8fc9fe	21efb07e7001a6ed	1	Server-206	srv-00000002
21000024ff8fc9ff	21efb07e7001a6ec	1	Server-206	srv-00000002
iqn.1991-05.com.microsoft.vm90-...	iqn.2011-04.com.zadarastorage.v...	1	VM90-WIN2012R2	srv-00000004

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.

Virtual Networks

If the VPSA was assigned multiple Virtual Networks (See: [Managing Virtual Networks](#)) they are list in the Controller's south panel.

Each network is displayed with its IP address and the associated VLAN ID.

Controllers							
Refresh		Failover					
Name	Storage Node	Management IPv4	Management IPv6	iSCSI IPv4	iSCSI IPv6	Status	Zone
vsa-00000064-vc-1	qa10-sn3	10.2.10.25				standby	zone_0
vsa-00000064-vc-0	qa10-sn4	10.2.10.21		10.2.10.21		active	zone_0

Page 1 of 1 | Displaying topics 1 - 2 of 2

Details for vsa-00000064-vc-0				
Properties Paths Virtual Networks System Usage Container Service CPU Usage Logs Performance Alerts				
Virtual Network Interface	IP	Network Mask	Gateway	VLAN ID
vni1	192.168.80.10	255.255.255.0	192.168.80.1	51

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 VPSA performance dashboard.

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.

Logs

Displays all event logs associated with the selected Controller.

Performance Alerts

The 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

MANAGING 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/CIFS protocols.

6.1 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).

6.1.1 Adding a Server for NAS access

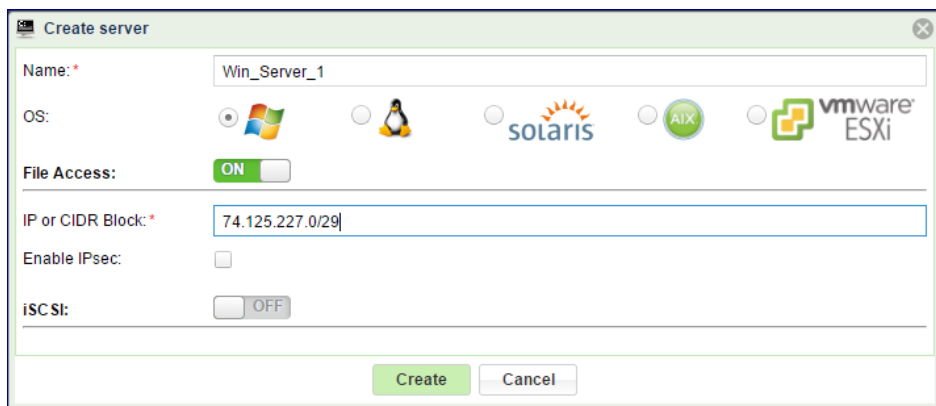
Adding servers to the VPSA in order to access files over NFS/SMB requires introducing the server's IP address to the VPSA.

- 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. 192.168.1.1/24)



- 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 VP SA. 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.

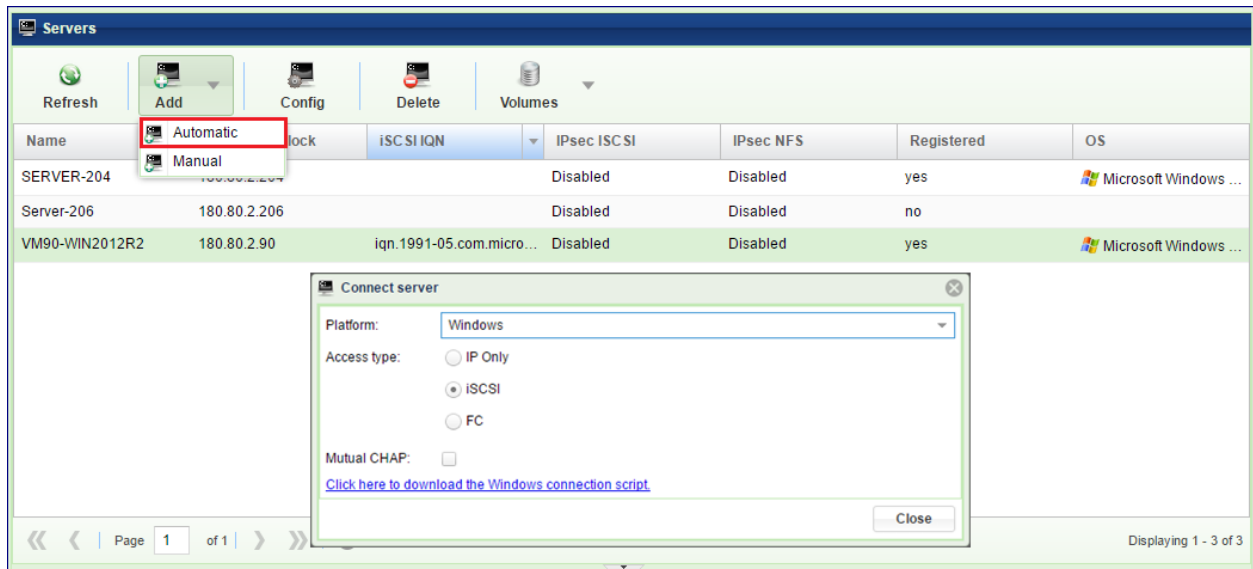
✓ **Note:** The VP SA NAS services will require the following ports and protocols to be accessible from the servers to the VP SA:

- 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:

6.1.2 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:



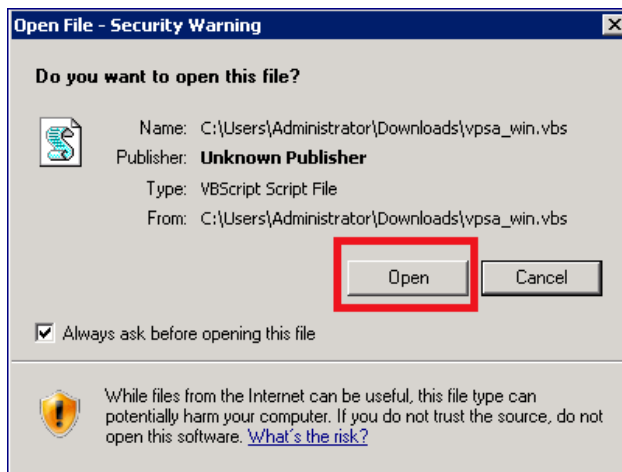
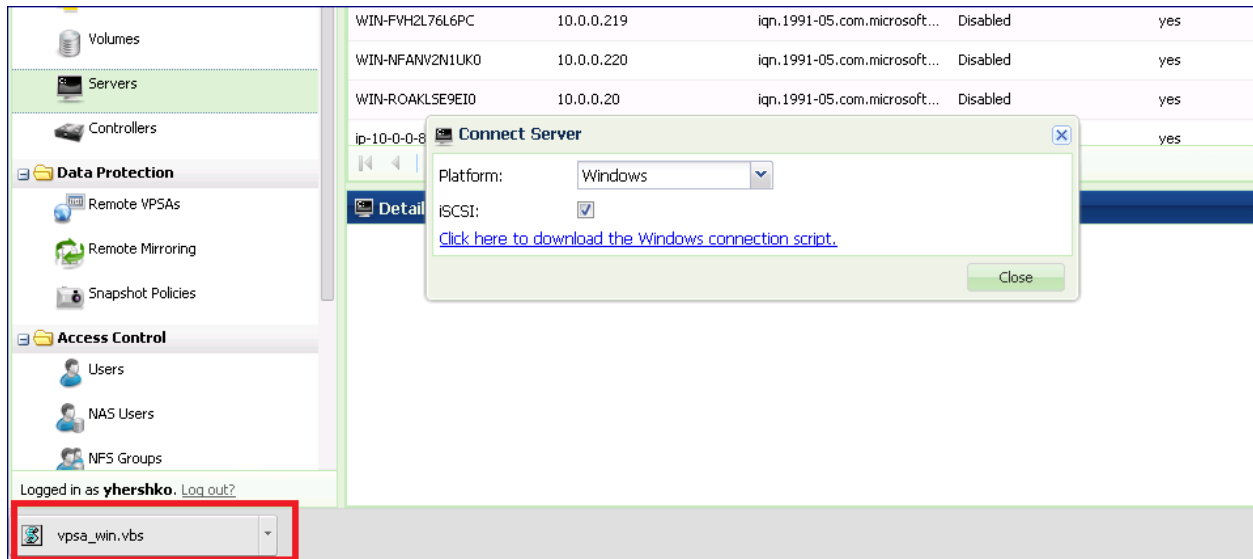
Adding a Server automatically over iSCSI

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 VPSA 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 VPSA 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:

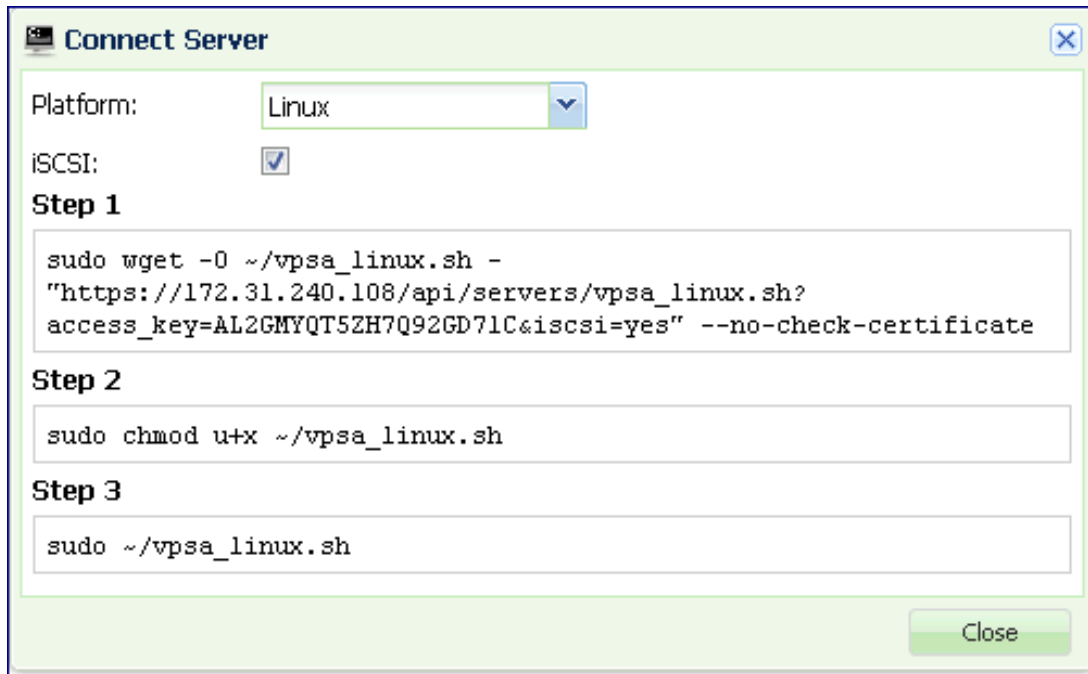
- On RedHat Servers do:

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

- On Ubuntu Servers do:

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

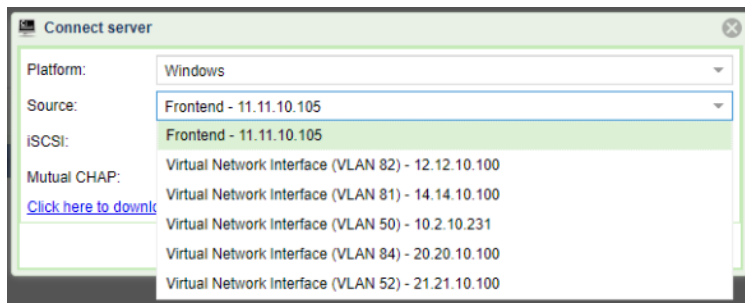
- On the **VPSA GUI > Connect Server** dialog, select platform: Linux.
- Select the iSCSI checkbox if you wish to expose VPSA Block Volumes to this Server.
- Run the three steps as detailed in the connect server dialog to execute the vpsa_linux.sh script.



- 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 VPSA 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:

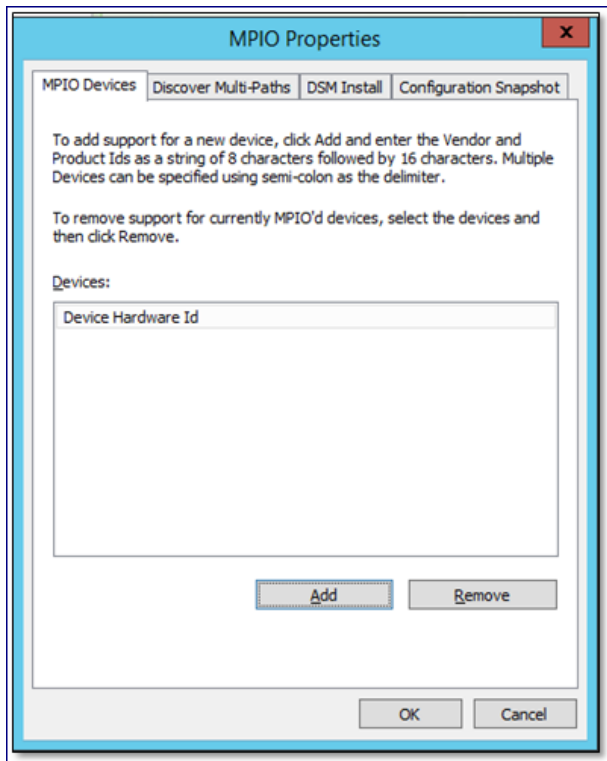


Adding a Server automatically over Fibre Channel

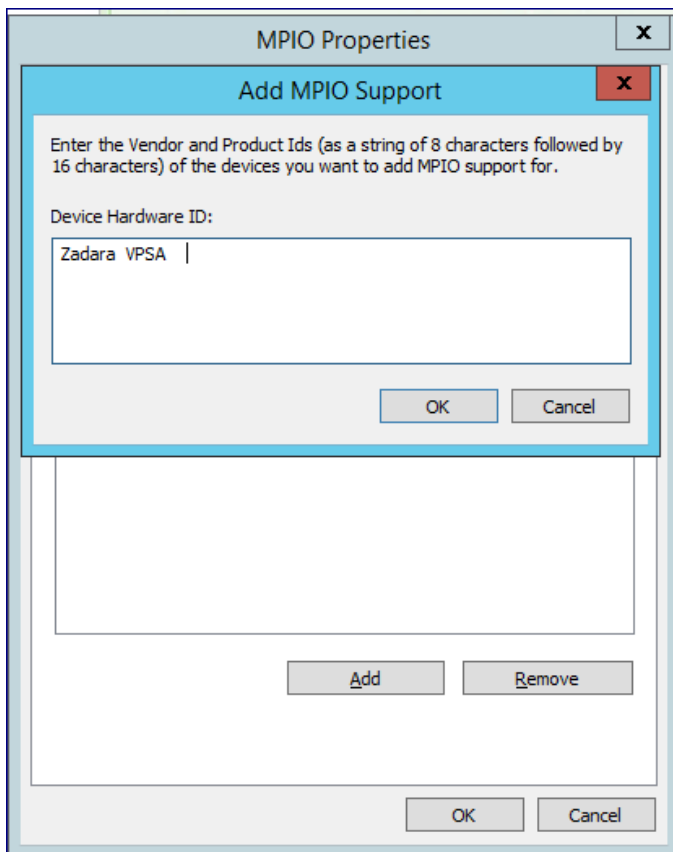
Before defining a server using FC connectivity, make sure to install and configure multipathing software on the server. Also make sure to setup the zoning on the FC switch to allow connectivity between the connecting server and the VPSA FC ports.

To Add a Windows Server:

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

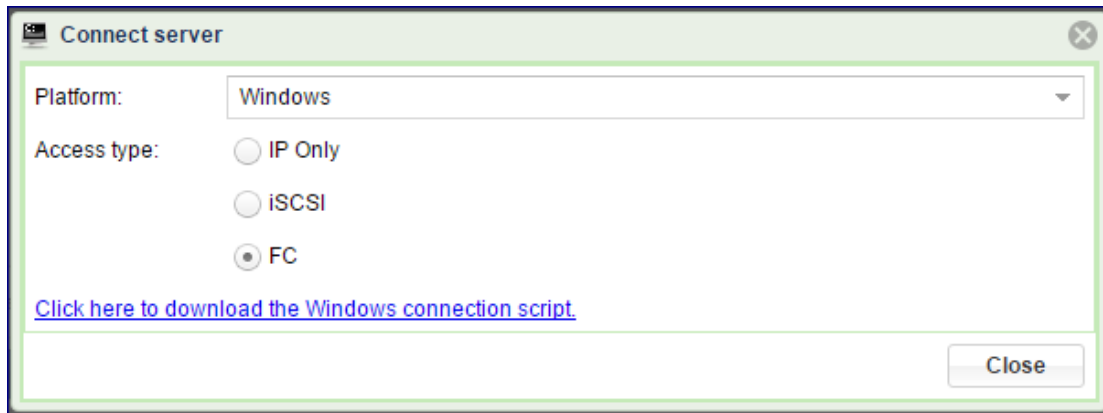


Enter the device HW ID as follows: "Zadara VPSA " ("Zadara" followed by 2 blank spaces, and then "VPSA" followed by 4 blank spaces. 16 characters total) and press OK.



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

- Open the VPSA GUI on the Windows Server
- On the [VPSA 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.zadarastorage.com/entries/21664397-How-To-setup-Multiple-FC-sessions-and-MultiPath-on-your-Linux-Cloud-Server>

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

- On the VPSA 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.



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

6.1.3 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:

- Enter the Server Name.

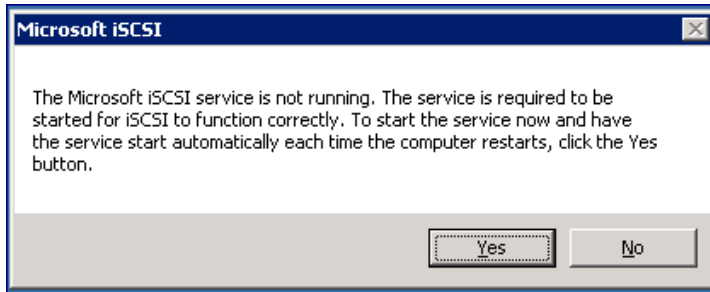
✓ Note: Objects names can be up to 128 chars long and can contain letters and digits, dashes “-” and underscores “_”

- 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](#)

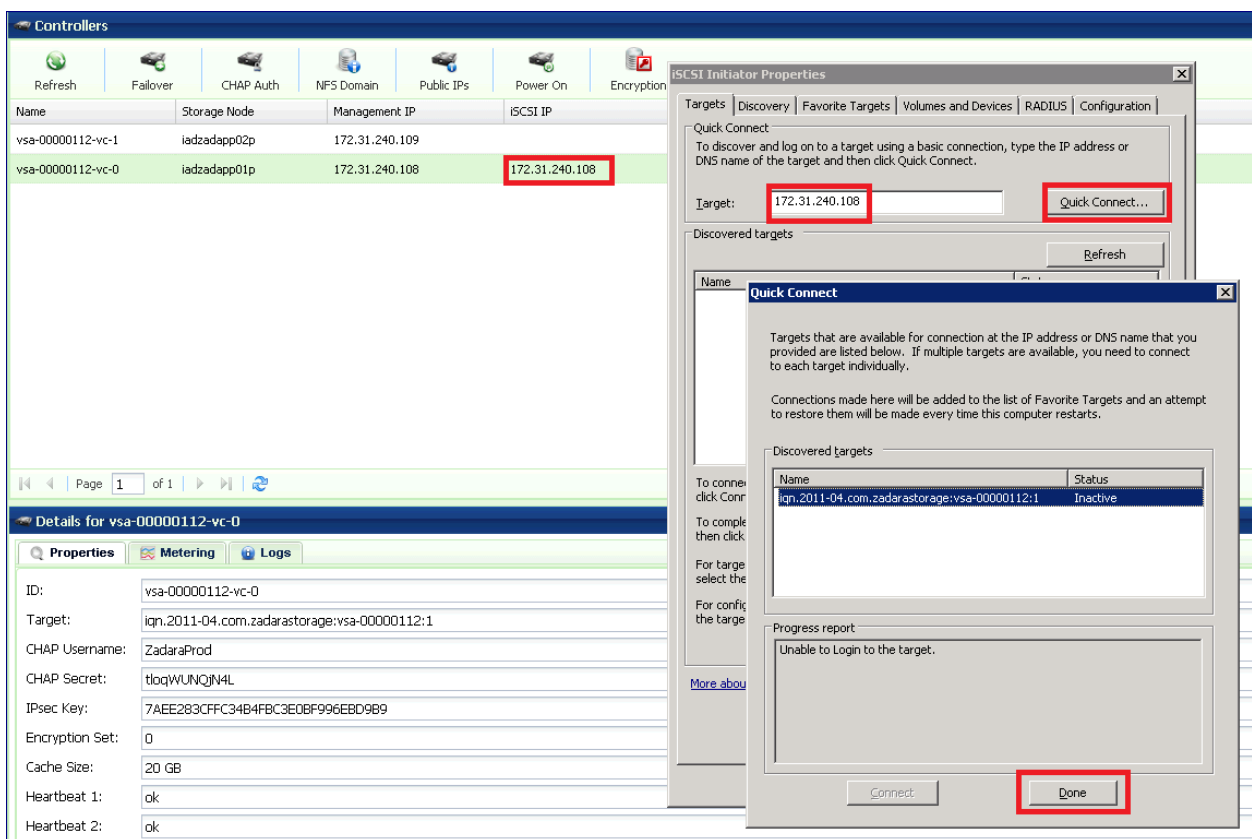
After manually adding a Server you need to establish an iSCSI connection between the Server and the VPSA. Please note that you can skip this step if the Server was added automatically or if the Server is only consuming NFS/SMB type Volumes.

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 VPSA (which is displayed in the [VPSA 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 VPSA discovered iSCSI target in an “Inactive” status. Press Done.



- To activate the connection, select the VPSA target and press the Connect button. Please note that if you have multiple targets listed you can identify the VPSA 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 VPSA 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 VPSA GUI, under the [Controllers](#) page, in the properties tab. Press OK to confirm the operation.

The screenshot displays the VPSA GUI interface. In the foreground, the 'ISCSI Initiator Properties' dialog box is open, showing the 'Advanced Settings' tab. The 'Enable CHAP log on' checkbox is checked. The 'Name' field contains 'ZadaraProd' and the 'Target secret' field is filled with dots. The background shows a table of controllers and a 'Details for vsa-00000112-vc-0' page with fields for ID, Target, CHAP Username, CHAP Secret, and IPsec Key.

- 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

<http://zadarastorage.zendesk.com/entries/20925646-how-to-enable-mpio-and-set-multiple-iscsi-sessions-on-windows-server-2008-r2>

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: <http://support.zadarastorage.com/entries/21664397-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.

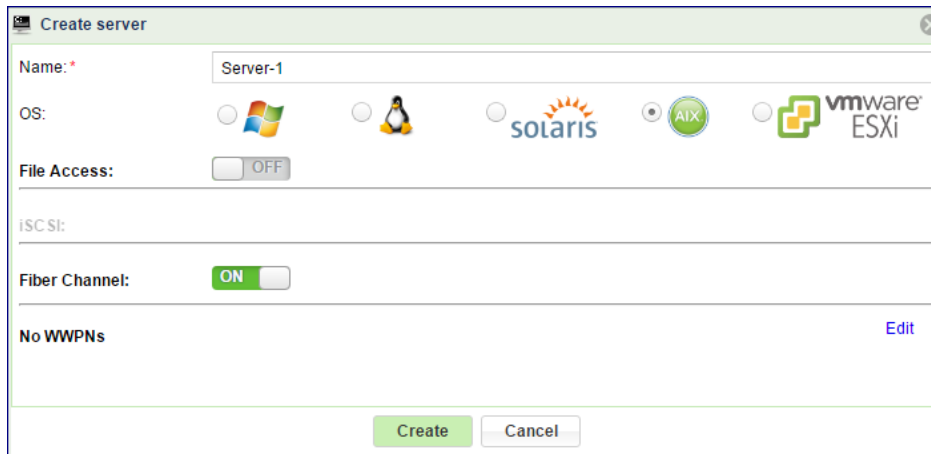
Establishing FC connection

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.con`:

```
#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.

6.1.4 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

Config server

Name: * server-215

File Access: ON

IP or CIDR Block: * 130.30.2.215

Enable IPsec:

iSCSI: ON

IQN: * iqn.1993-08.org.debian:01:51843c6c3a9a

Enable IPsec:

CHAP: **VPSA** **Host** [Edit](#)
 User: Liran_Large_Cache1 Mutual CHAP Disabled
 secret: CmhO7Lyzg1mg

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

Config server

Name: * SERVER-204

File Access: ON

IP or CIDR Block: * 180.80.2.204

Enable IPsec:

iSCSI:

Fiber Channel: ON

21000024ff7bc175 [Edit](#)
 21000024ff7bc174

6.2 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:	<input type="text" value="srv-00000001"/>
Name:	<input type="text" value="SERVER-204"/> ✎
VPSA CHAP User:	<input type="text" value="Dima_VPSA5"/>
VPSA CHAP Secret:	<input type="text" value="gWmfDExT3u2q"/>
Host CHAP User:	<input type="text"/>
Host CHAP Secret:	<input type="text"/>
IP or CIDR Block:	<input type="text" value="180.80.2.204"/>
iSCSI IQN:	<input type="text"/>
IPsec iSCSI:	<input type="text" value="Disabled"/>
IPsec NFS:	<input type="text" value="Disabled"/>
WWPN 1:	<input type="text" value="21000024ff7bc174"/>
WWPN 2:	<input type="text" value="21000024ff7bc175"/>
Registered:	<input type="text" value="yes"/>
OS:	<input type="text" value="Microsoft Windows Server 2012 R2 Datacenter 6.3.9600"/>
Added:	<input type="text" value="2016-03-30 16:48:54"/>
Modified:	<input type="text" value="2016-03-30 16:48:54"/>

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.

Volumes

A list of all the Volumes attached to this Server.

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.

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.

Logs

Displays all event logs associated with this Server.

Performance Alerts

A VPSA 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.

MANAGING VOLUMES, SNAPSHOTS AND CLONES

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.

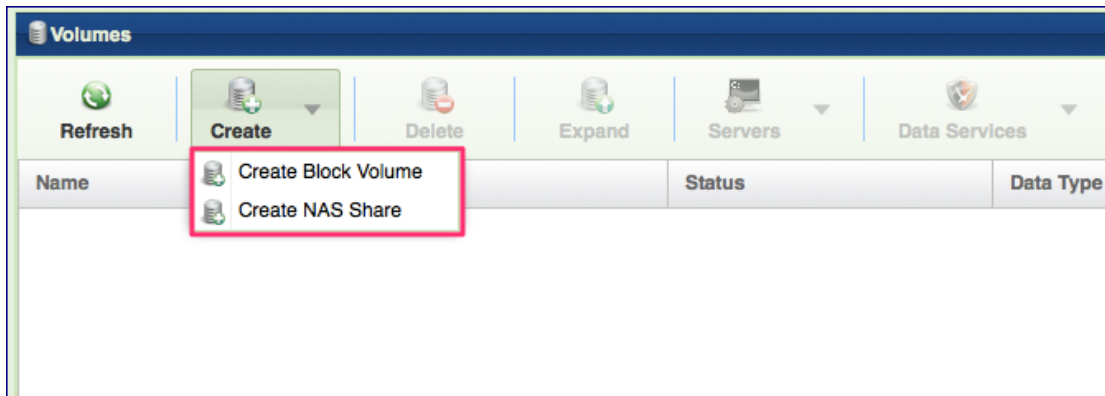
Snapshots are read-only representations of the Volume's data at a given point-in-time. They are thinly provisioned and share the same data chunks with their Volume as much as possible until you actually modify the chunk's data. This triggers a Redirect On Write (ROW) operation where a new chunk is provisioned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

7.1 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.



Creating a Block Volume

Create Block Volume

Name: *

Capacity (GB): *

Pool: *

Name	Status	Free Capacity
Pool1	normal	1.06 TB Free / 1.06 TB

Encrypted:

Attach Default Snapshot Policies:

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:** Object 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).
- **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 please check [Managing Encrypted Volumes](#)
- **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.

Creating a NAS Share

🏠 Create Share Volume
✕

Name: *

Capacity (GiB): *

Export Name:

Pool: *

Name	Status	Free Capacity
RAID-10-Pool-1	normal	10.80 TiB Free / 10.81 TiB

Encrypted:

atime Update:

Attach Default Snapshot Policies:

Read Ahead Size KiB: * 16KiB 64KiB 128KiB 256KiB 512KiB

User Quotas: * Off On

Group Quotas: * Off On

Project Quotas: * Off On

▲ Automatic Expansion

Automatic Expansion:

▲ SMB Options

SMB Only (optimized for performance):

Allow Guest Access:

Encryption mode: Off Desired Required

Enhanced Windows ACLs:

File Creation Mask:

Directory Creation Mask:

Map Archive:

Browseable:

Hidden files:

Hide Unreadable:

Hide Unwriteable:

Hide Dot files:

Store DOS Attributes:

Enable Oplocks:

SMB Serial small IO workload optimized:

▲ NFS Options

NFS root squash:

NFS all squash:

NFS anonymous GID:

NFS anonymous UID:

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 Expeort 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 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.
- **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 please 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. By default atime Update is disabled. Enabling it will impact performance.
- **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 a 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 a 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 here about Project Quotas: [Setting Project Quotas](#)

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

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 VPSA 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://blogs.msdn.microsoft.com/openspecification/2012/10/05/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.

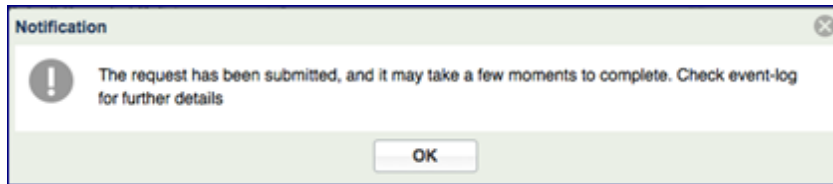
✓ Note:

- 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.

✓ **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.

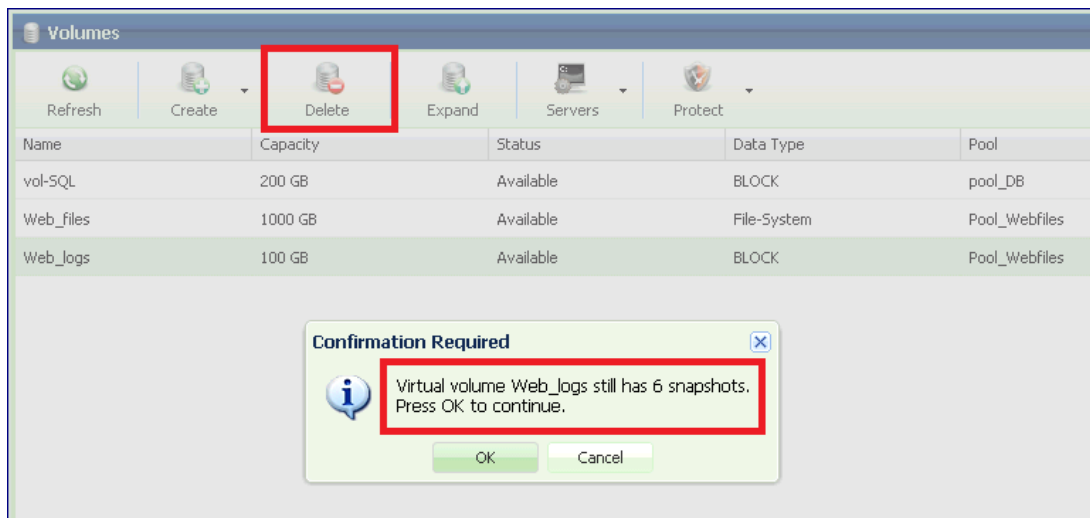


Deleting a Volume/Share

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.

7.2 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:

<input checked="" type="checkbox"/>	Name	IP or CIDR Block	iSCSI IQN	IPsec NFS	OS	Access Type
<input checked="" type="checkbox"/>	Server-213	130.30.2.213	iqn.1993-08.org.de...	Disabled	Ubuntu 14.0...	NFS or SMB

<input checked="" type="checkbox"/>	Name	IP or CIDR Block	iSCSI IQN	IPsec iSCSI	OS	Access Type
<input checked="" type="checkbox"/>	SERVER-204	180.80.2.204		Disabled	Microsoft W...	FC

- 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):

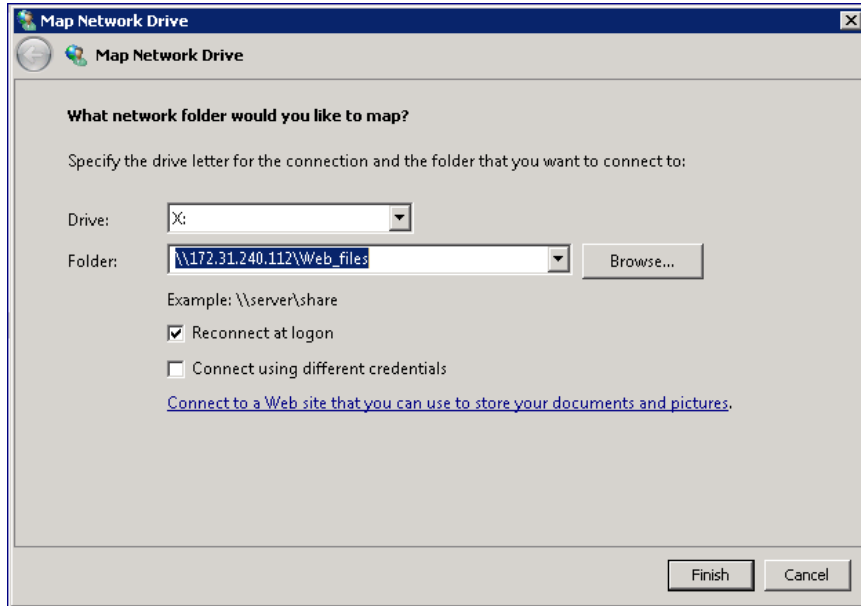
```
$mount -t nfs4 <NFS_Export_Path>/<mount point>
```

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](#) and 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.
2. The first time you connect from a Windows Server to a VPSA share you are requested to enter an SMB User name and Password. Please check [Creating SMB Users](#) for more details (or use SMB guest access).



Format a Volume

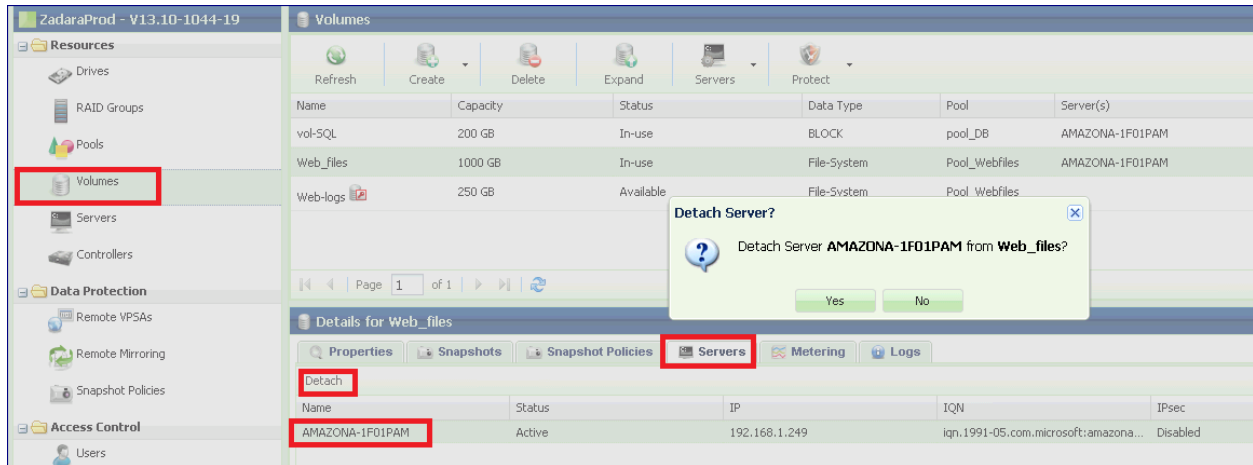
Once the Volume is attached to the Server and identified by the Operating System as a drive, use the specific OS tools to format the drive to the needs of the OS or file-system used. Allocation units of 512B to 64KB are supported.

To detach a Volume

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

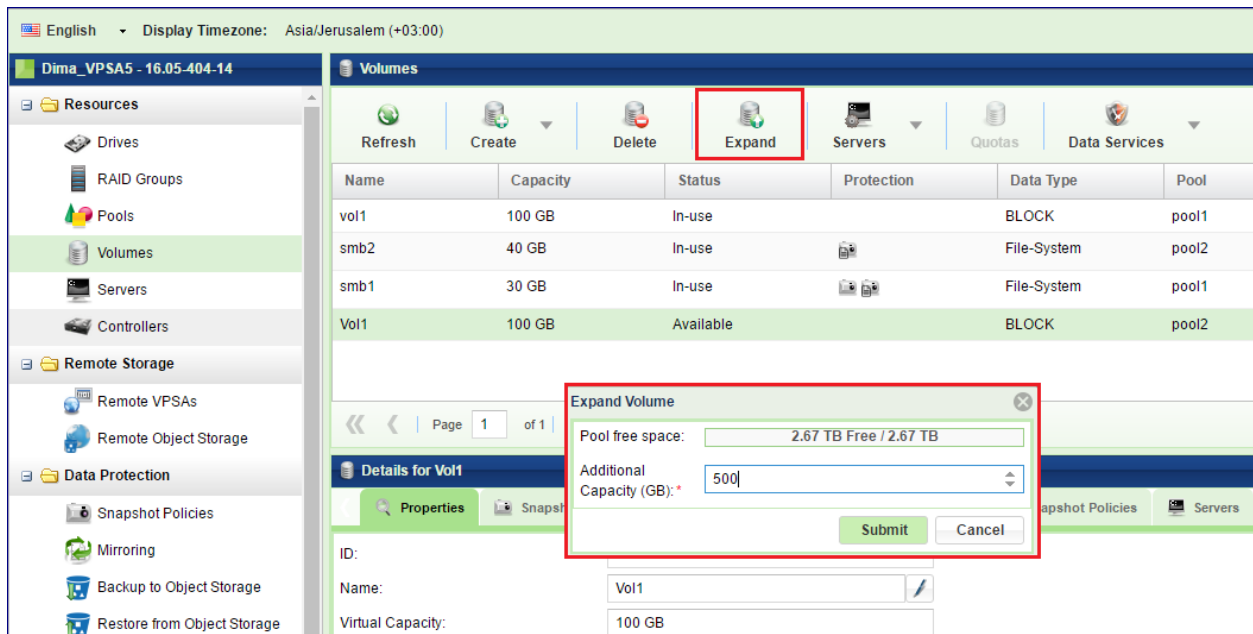
To detach a Volume from a Server, go to the [Volumes](#) Page and click the [Servers > Detach from Server\(s\)](#) button. You will be requested to select the Servers from which to detach this Volume.

Alternatively, you can view the attached Servers list in the Volume's South Panel, select the Server to detach from and click the Detach button on the top-left corner of the South Panel:



7.3 Expanding a Volume

You can expand a Volume anytime, regardless if the Volume has Snapshots, Clones or is being remotely mirrored. To expand a Volume go to the [Volumes](#) page, select the Volume and press the Expand button. Enter the amount of virtual capacity you'd like to expand the Volume by and press Submit.



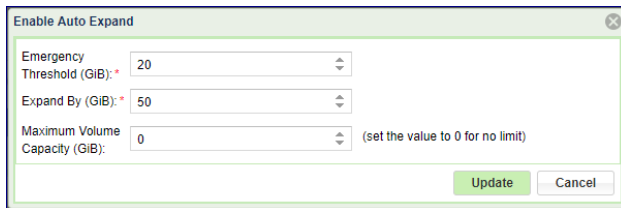
7.3.1 Volume Automatic Expansion

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

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

Auto 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)



The screenshot shows a dialog box titled "Enable Auto Expand". It contains three input fields, each with a spinner control:

- Emergency Threshold (GiB): 20
- Expand By (GiB): 50
- Maximum Volume Capacity (GiB): 0 (set the value to 0 for no limit)

At the bottom of the dialog, there are two buttons: "Update" and "Cancel".

By default all volumes are created with Auto Expansion disabled. To enable it check the Automatic Expansion checkbox on the share creation dialog, or enable it from the Capacity Properties page.

Create Share Volume

Capacity (GiB): *

Export Name:

Pool: *

Name	Status	Free Capacity
RAID-10-Pool-1	normal	800 GiB Free / 6.88 TiB

Encrypted:

atime Update:

Attach Default Snapshot Policies:

Read Ahead Size KiB: * 16KiB 64KiB 128KiB 256KiB 512KiB

User Quotas: * Off On

Group Quotas: * Off On

Project Quotas: * Off On

Automatic Expansion

Automatic Expansion:

Emergency Threshold (GiB): 10

Expand By (GiB): 50

Maximum Volume Capacity (GiB): 0 (set the value to 0 for no limit)

SMB Options

NFS Options

Submit Cancel

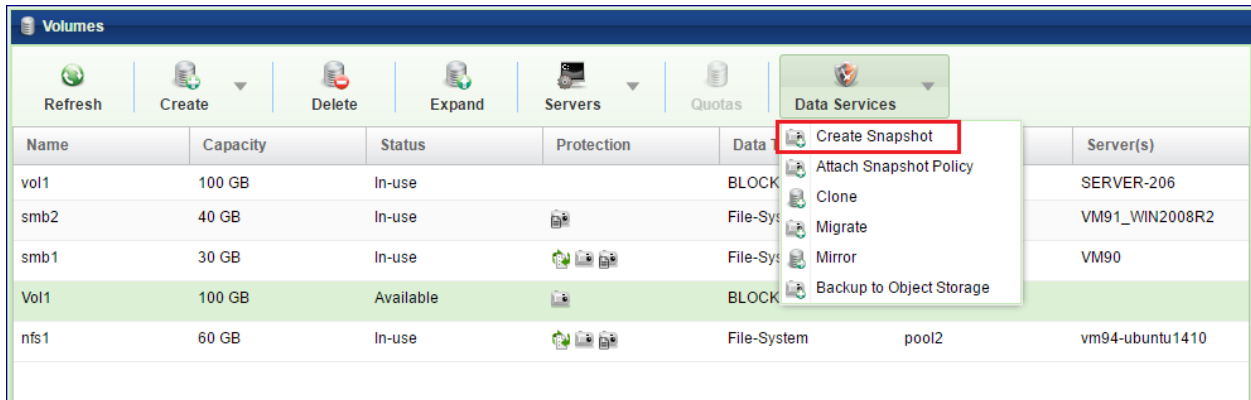
7.4 Managing Snapshots and Snapshot Policies

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.

7.4.1 Manual creation & deletion of Snapshots

To manually create a Snapshot:

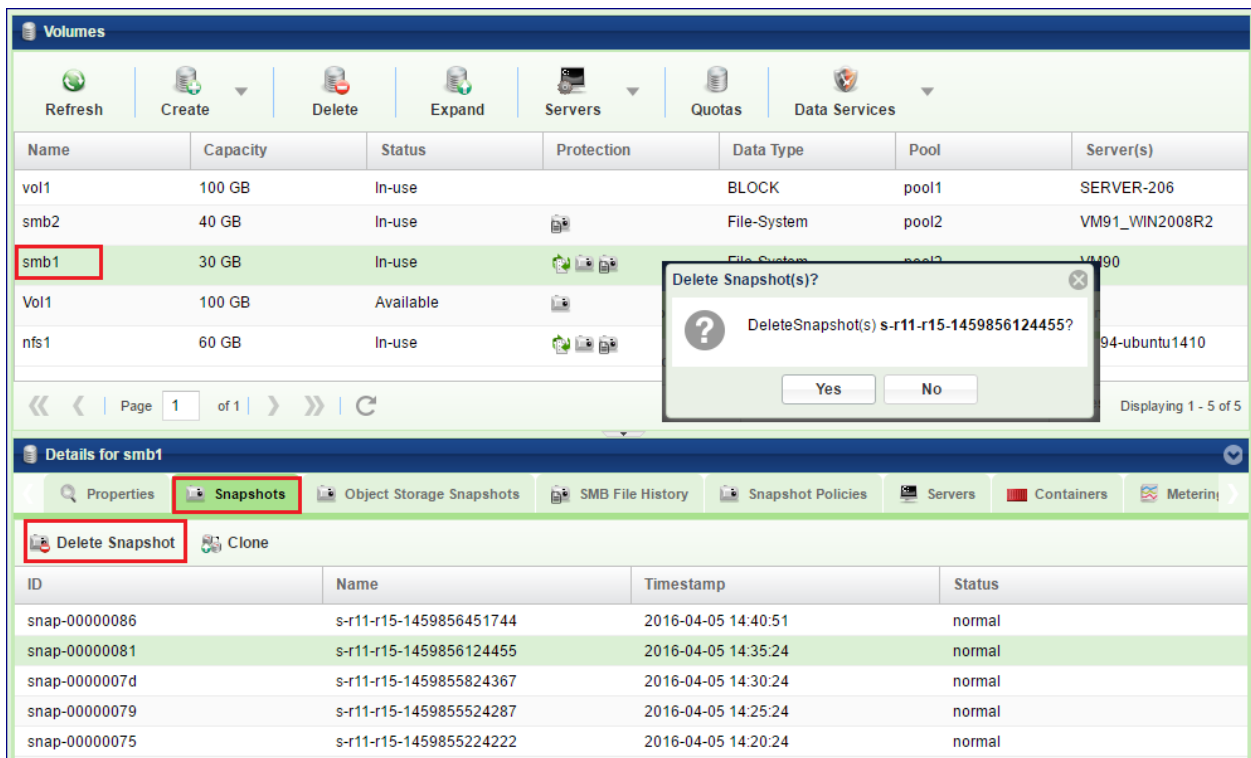
- Go to the Volumes page, press the Data Services button and select Create Snapshot.



- Enter a Unique Snapshot name 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 [Managing Remote Mirroring](#) for details.

7.4.2 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 VPSA.

Name	Create Policy
Hourly Snapshots for a Day	Every hour
Daily Snapshots for a Week	Once per day at midnight
Weekly Snapshots for a Year	Every Sunday at midnight
Daily Backup for a Year	Once per day at midnight

To create a new Snapshot Policy:

- Go to the [Snapshot Policies](#) page and press the Create button.
- **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.

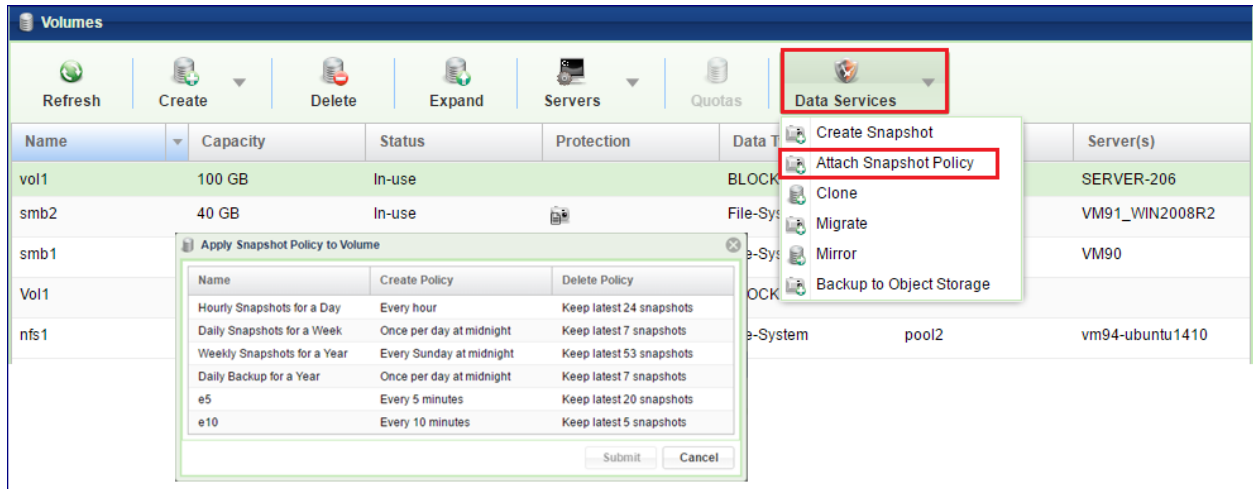
- **Set as default policy for newly created volumes** – Select this checkbox if you'd like all new Volumes to default to this Snapshot Policy. Select the appropriate Creation Policy from the drop down list.
- Define the number of Snapshots to retain in the deletion policy.
- Allows Empty Snapshot Creation – 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.

To Edit a Snapshot Policy

- Go to the [Snapshot Policies](#) page, select the Policy and press the Edit button.
- You can edit all of the Snapshot Policy's attributes: Name, Creation Policy, Deletion Policy Allow Empty Snapshots Creation and Set as Default Policy.
- 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.
- 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.

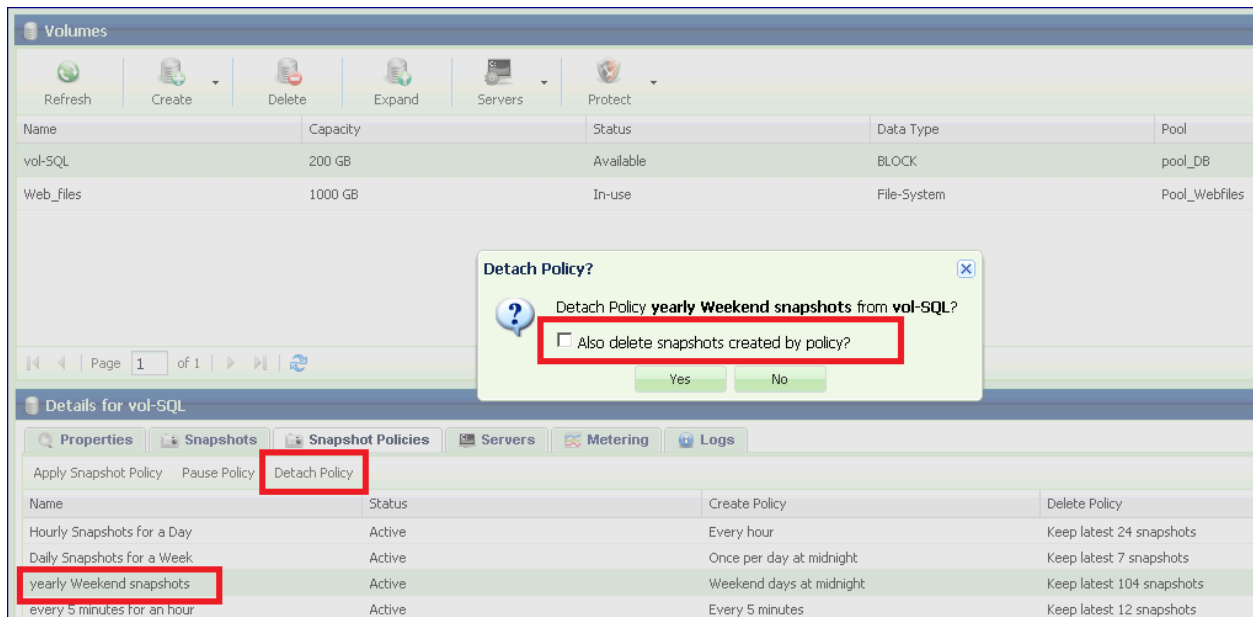
To Apply a Snapshot Policy on a Volume

- Go to the [Volumes](#) page, select the Volume and select [Data Services > Attach Snapshot Policy](#) from the menu.
- Select the Snapshot Policy to apply to the Volume and press the Submit button.



To detach a Snapshot Policy from a Volume

- Go to the [Volumes](#) page, select the Volume and press the Snapshot Policies south tab to view the Volume's applied Snapshot Policies.
- Select the Snapshot Policy to delete and press the Detach Policy button on 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.

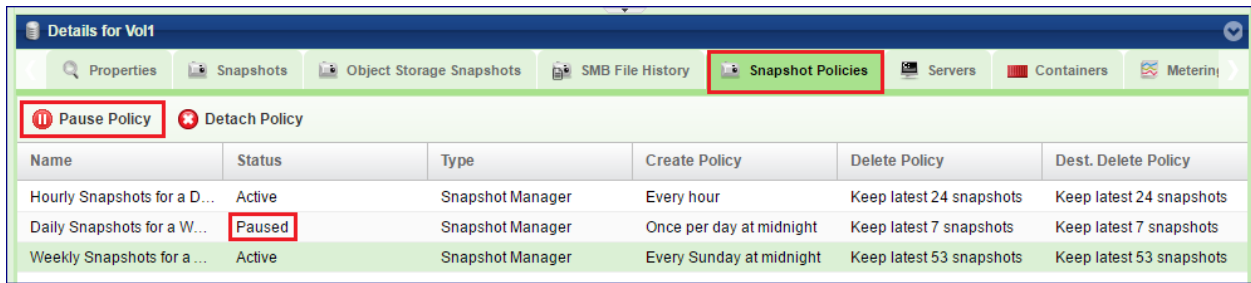


Pause/Resume a Snapshot Policy

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 have this Policy active as well.

- To pause a Snapshot Policy, 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.
- Select the Snapshot Policy and press the Pause Policy button on the top left corner of the South Panel.
- The Policy status will change to "Paused".

- To resume a Policy: The Pause / Resume button toggles according to the current Policy status. Select a Policy in a Paused state and press the Resume Policy button. The Policy Status will change to “Active”.

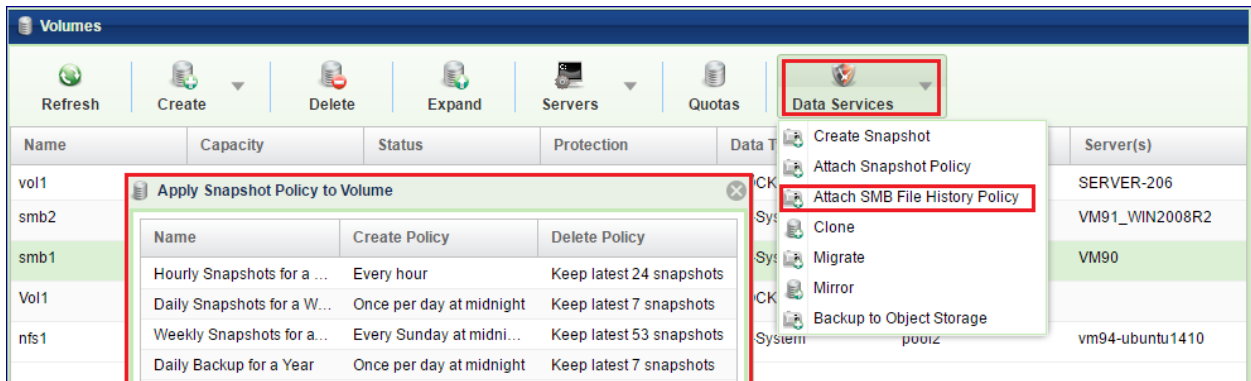


7.5 Managing SMB File History

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.

To Apply a SMB File History Policy on a Volume

- Go to the Volumes page, select the NAS Volume and select [Data Services-> Attach SMB File History Policy](#) from the menu.
- Select the Snapshot Policy to apply to the Volume and press the Submit button.

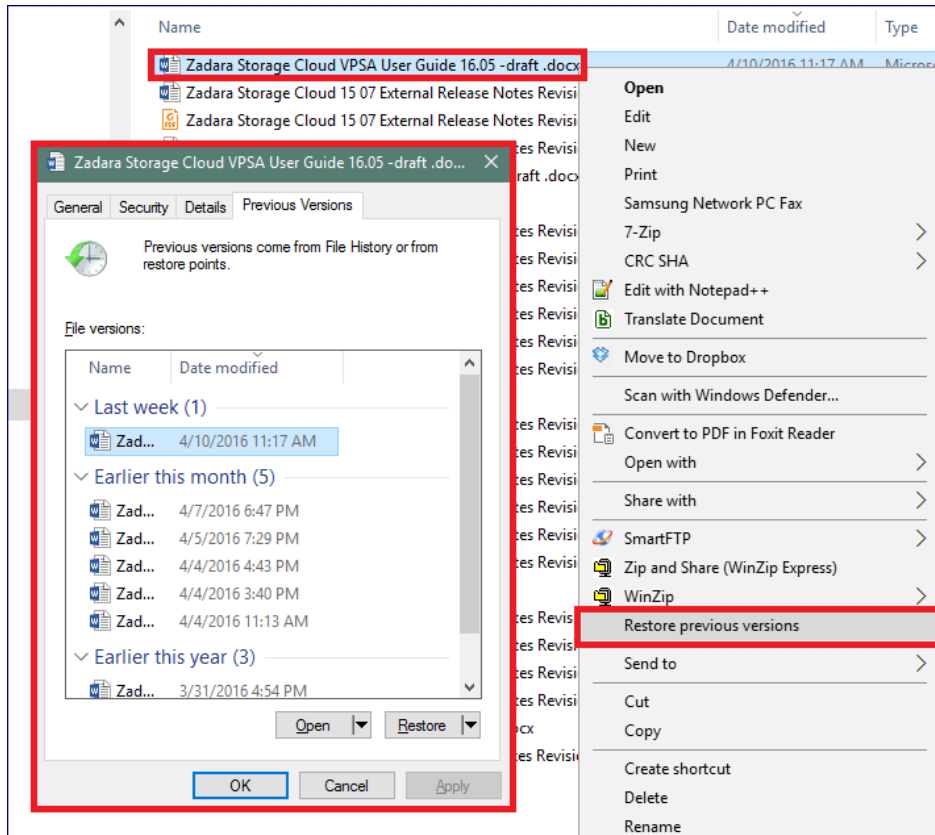


To detach a SMB File History Policy from a Volume

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

To restore files from SMB File History

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

✓ **Note:** 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

7.6 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

Go to The Volumes page, select the Volume to be cloned and press the Data Services > Clone button.

- **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.**
- press the Submit button to complete the operation.

The screenshot shows the 'Volumes' management page with a 'Data Services' dropdown menu open, highlighting the 'Clone' option. Below, the 'Clone Volume' dialog box is displayed. The 'Clone Name' field contains 'fc_vol1'. The 'Clone from' table lists three snapshots, with the second one, 's-r70-1460889203002', selected. The 'Details for fc_vol1' panel shows the volume's properties: Name: fc_vol1, Virtual Capacity: 60 GB, and Mapped Capacity: 20 GB.

Name	Capacity	Status	Protect
smb29	20 GB	Available	
smb30	20 GB	Available	
smb31	20 GB	In-use	
smb32	20 GB	Available	
smb34			

Clone Volume			
Clone Name: *			
Clone from:			
<input type="checkbox"/>	Name	Timestamp	Status
<input type="checkbox"/>	s-r70-1460889263164	2016-04-17 13:34:23	normal
<input checked="" type="checkbox"/>	s-r70-1460889203002	2016-04-17 13:33:23	normal
<input type="checkbox"/>	s-r70-1460730146743	2016-04-15 17:22:26	normal

Details for fc_vol1	
Properties	Value
ID:	Volume 00000000000000000000000000000000
Name:	fc_vol1
Virtual Capacity:	60 GB
Mapped Capacity:	20 GB

- Alternatively, you can go to the Volumes page, select the Volume to be cloned, press the Snapshots tab at the South Panel, select the desired point-in-time Snapshot and press the Clone button at the top left corner of the South Panel.
- Enter a name for the new cloned Volume.

The screenshot shows the 'Volumes' management page. At the top, there are buttons for Refresh, Create, Delete, Expand, Servers, and Protect. Below these is a table of volumes:

Name	Capacity	Status	Data Type	Pool
vol-SQL	200 GB	In-use	BLOCK	pool_DB
Web_files	1000 GB	In-use	File-System	Pool_Webfiles

The 'Snapshots' tab is selected, showing a list of snapshots:

ID	Name	TimeStamp	Status
snap-00000023	s-r7-1385040942847	2013-11-21 08:35:42	normal
snap-00000022	s-r7-1385040642810	2013-11-21 08:30:42	normal
snap-00000021	s-r7-1385040342769	2013-11-21 08:25:42	normal
snap-00000020	s-r7-1385040042733	2013-11-21 08:20:42	normal
snap-0000001f	s-r7-1385039742688	2013-11-21 08:15:42	normal
snap-0000001e	s-r7-1385039442645	2013-11-21 08:10:42	normal

A dialog box titled 'Clone from Snapshot' is open, displaying the following text:

You are about to create a clone of the Volume **vol-SQL** from the Snapshot **s-r7-1385040642810** which was created at: **2013-11-21 08:30:42**

Please enter a name for this new Volume:

vol-SQL-test

Buttons for OK and Cancel are visible at the bottom of the dialog.

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

The screenshot shows the 'Volumes' management page with the following table:

Name	Capacity	Status	Data Type	Pool
vol-SQL	200 GB	In-use	BLOCK	pool_DB
Web_files	1000 GB	In-use	File-System	Pool_Webfiles
vol-SQL-test	200 GB	Available	BLOCK	pool_DB

The 'vol-SQL' row is labeled 'Original Volume' and the 'vol-SQL-test' row is labeled 'Cloned Volume'.

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

7.7 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](#).

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 the Data Services option and then select Migrate.

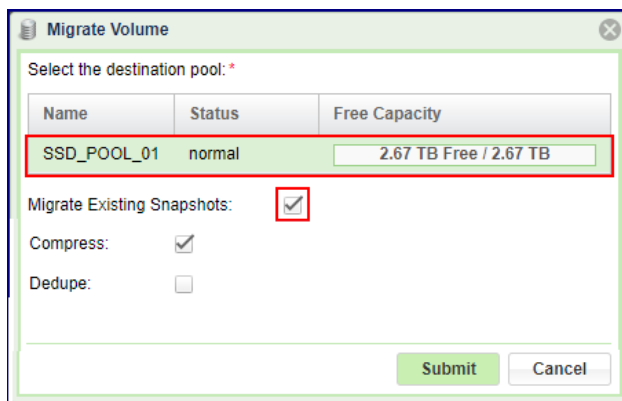


- **Destination pool** – Select the destination Pool to migrate to, from the list of available pools. Make sure to select a Pool with sufficient free capacity.
- **Migrate Existing Snapshots** – Check the checkbox if the migration of the volume should include the existing snapshots of the volume. In case “Migrate Existing Snapshots” is checked all snapshots will be migrated to the destination Pool. Note that in case the “Migrate Existing Snapshots” is not checked, the Volume snapshots will be deleted.

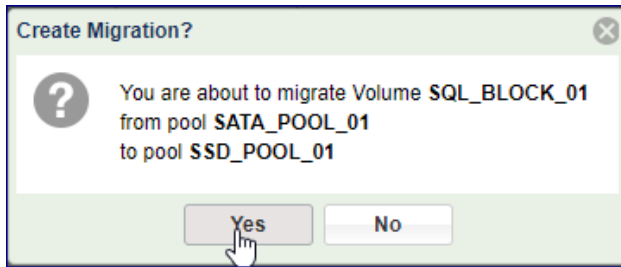
VPSA FLASH ARRAY

- **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 the Submit button to start the operation.



3. Review the details and confirm the Online Volume Migration operation.



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, a new tab is available - Migration Status. The Migration Status tab will provide real-time migration information while the migration is still running.

Name	Capacity	Status	Protection	Data Type	Pool	Server(s)
SQL_BLOCK_01	100 GB	Migrating		BLOCK	SATA_POOL_01	YOK-WIN-LAB-01
SQL_BLOCK_02	100 GB	In-use		BLOCK	SATA_POOL_01	YOK-WIN-LAB-01

Details for SQL_BLOCK_01

Properties | Snapshots | Object Storage Snapshots | Snapshot Policies | Servers | Containers | Metering | Logs | Performance Alerts | **Migration Status**

Abort | Pause | Continue

Name: mgrjob-00000001

Status: Syncing

Progress: 8751 MB done of 36472 MB (24.0%) 00:03:07 left

Transfer Rate (MB/s) vs Time (wrt-bandwidth)

1 sec | 10 sec | 1 min | 10 min | 1 hour | Refresh | Auto

4. The user has complete control on the migration task as it can be Paused or Aborted from the Migration Status tab.

- Upon completion, the Migration Status tab will be removed from the Volume south panel. A log entry will be added as an indication of a successful migration.

The screenshot displays the VPSA Storage Array management interface. The left sidebar shows a navigation tree with categories like Resources, Remote Storage, Data Protection, and Container Service. The main area is titled 'Volumes' and contains a table of storage volumes. Below the table is a 'Details for SQL_BLOCK_01' section with a 'Logs' tab selected, showing a list of events.

Name	Capacity	Status	Protection	Data Type	Pool	Server(s)
SQL_BLOCK_01	100 GB	In-use		BLOCK	SSD_POOL_01	YOK-WIN-LAB-01
SQL_BLOCK_02	100 GB	In-use		BLOCK	SATA_POOL_01	YOK-WIN-LAB-01

Title	Time
Migration of volume SQL_BLOCK_01 to storage pool SSD_POOL_01 completed successfully.	2017-09-25 10:45:14
Migration of volume SQL_BLOCK_01 to storage pool SSD_POOL_01 started successfully.	2017-09-25 10:41:21
Volume SQL_BLOCK_01 was successfully attached to Server YOK-WIN-LAB-01, IQN [iqn.1991-05.com....	2017-09-25 09:21:11
Volume SQL_BLOCK_01, of size 100 GB and access mode BLOCK, was created successfully.	2017-09-25 09:13:40

7.8 Managing Encrypted Volumes

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.

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.

The encryption attribute of Volumes 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.

To create a Master Encryption Password, go to the [Settings](#) page, Security tab and press the Edit in the Encryption section. Read the instructions and warning. Type your Password and Save.

The screenshot shows the VPSA Settings page for 'VPSA1 - 16.05-404-39'. The left sidebar contains a navigation menu with categories: 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 Memory Pools, Containers), NAS Access Control, System (Settings, User Management, Logs, Support). The main content area is titled 'Settings' and has tabs for General, Security, NAS, Metering, and Container Service. The Security tab is active, showing the following configuration:

- Passwords policy:** Enforce password expiration: No, Password history: 8
- Global VPSA CHAP:** User: VPSA1, Secret: 83hCQ0bJKd5

A prominent warning message in red text states: **IMPORTANT: PLEASE PAY CLOSE ATTENTION TO THESE INSTRUCTIONS. IF YOU DO NOT, THERE IS A RISK YOU WILL PERMANENTLY LOSE ACCESS TO YOUR ENCRYPTED DATA!**

Below the warning, the text reads: "In order to protect the privacy of your data, we store your data encryption/decryption keys in a way which is inaccessible to us. The benefit to you is that neither we, nor your Cloud Provider, nor anyone else can decrypt your data. Because your keys are inaccessible to us, **YOU MUST** select AND **SAVE** a password with which our systems can access your encryption/decryption keys."

The interface prompts the user to "Please enter a password for protecting your encryption/decryption keys:" and provides two input fields: "New Encryption Password" and "New Encryption Password Again".

A warning message follows: "YOU MUST SAVE THIS PASSWORD FOR FUTURE USE. There is a small chance that our system will prompt you for this password in the future. If it does and you are unable to produce your password, YOU WILL PERMANENTLY LOSE ACCESS TO YOUR ENCRYPTED DATA and there would be NO WAY FOR US TO RECOVER YOUR DATA."

An unchecked checkbox contains the text: "By submitting this form you acknowledge that you understand that it is YOUR RESPONSIBILITY TO SAVE THIS PASSWORD and that failure to do so can result in PERMANENT LOSS OF ACCESS TO YOUR ENCRYPTED DATA. You also understand that there is a chance you will be required to enter the password to access to your encrypted data."

A "Save" button is located at the bottom of the form.

At the bottom of the page, the "Cloud Admin Access:" status is shown as "Enabled".

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

Settings

General **Security** NAS Metering Container Service

Security

Passwords policy: Enforce password expiration: No, Password history: 8

Dual Factor Authentication: Enforce Dual factor Authentication: No

Global VPSA CHAP: User: Dima_VPSA2
Secret: C9rU32NkgYaF

Some (or all) of your volumes are encrypted.
You can change your password below:

Current Encryption Password: *

New Encryption Password

New Encryption Password Again

Save Remove encryption password

Cloud Admin Access: Enabled

IPSec Key: 974B477DD17D49A69F75FAF762DE7942

affect the encrypted data.

Store your Master Encryption Password in a secure place

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



Encrypted Volumes are displayed with the icon.

7.9 Protecting Files Shares with Built-in Anti-Virus

The Zadara VPSA provides Anti-Virus protection to file shares with its integrated McAfee Anti-Virus engine. Generally speaking, all you need to do is to enable virus scanning at the volume creation time, or at any time later. You can either use the default scanning policy, or modify it to fit your specific needs.

Infected files will be either deleted or quarantine according to the policy. Virus signatures and Virus scan engine updates are completely automatic.

Understanding the Anti-Virus page The following screen appears when selecting [Antivirus](#) on the navigation left panel.

The screenshot displays the Antivirus configuration interface. The left sidebar contains a navigation menu with 'Antivirus' highlighted. The main content area is divided into four sections:

- 1. Navigation:** The 'Antivirus' link in the left sidebar.
- 2. Engine Properties:** Shows engine status (Enabled), version (5800.7501), DAT version (8275.0), update time (2018-06-24T08:21:17.000+00:00), and quarantine volume pool (pool1).
- 3. Update On-Demand Scan Parameters:** Configures scan schedule (Custom, 14:50), frequency (Every Day Of Month), scan subfolders (YES), scan archives (NO), primary action (Clean), secondary action (Delete), file types to scan (All), and file type exclusions/inclusions.
- 4. Affected Files:** A table listing infected files. The table has columns: Volume Id, Volume Name, Name, Path, and Virus Name. The row shows: volume-00000003, av_nas1, 1, /export/av_nas1, EICAR test file.

This screen is divided into the following sections:

1. **Navigation** - Click here for Anti-Virus
2. **Engine Properties** - Use it to verify that your engine is up-to-date, and to enable/disable AV engine.
3. **Policy** - Use it to modify the scanning policy and adjust it for your needs.
4. **Quarantine Viewer** - Use it to view and manipulate infected files that were moved into the quarantine.

Activating Virus protection at Volume creation

Create Share Volume

Name: *

Capacity (GB): *

Export Name:

Pool: *

Name	Status	Free Capacity
pool1	normal	1.42 TB Free / 1.42 TB

Encrypted:

atime Update:

Attach Default Snapshot Policies:

Read Ahead Size KB: * 16KB 64KB 128KB 256KB 512KB

User Quotas: * Off On

Group Quotas: * Off On

Project Quotas: * Off On

— ▼ SMB Options

— ▼ NFS Options

▲ Antivirus Options

Enable On Demand Scan:

File Types To Scan:

Exclude File types:

Include file types:

Exclude Path:

You can activate virus protection when creating NAS share. - Follow the instruction in Creating a NAS Share. - Expand the Antivirus Options sections - Check the Enable On Demand Scan - Select All file types to scan, or specify specific file types extensions - If all file types are selected you may specify a list of excluded file extensions - You may exclude specific folders on the share from being scanned by specifying their full path

✓ **Note:** The above parameters overwrite the default Anti-Virus policy described in the next section.

You can also activate/deactivate virus protection on an existing Volume.

Open the [Volumes](#) page, select the volume of interest and click Antivirus. Click Attach Policy or Detach Policy. Provide the same parameters as described above.

Volumes						
Name	Capacity	Status	Protection	Data Type	Pool	Antivirus
vol1	200 GB	In-use		File-System	pool1	VM-111
vol2	200 GB	In-use		File-System	pool1	VM-111
vol3	200 GB	In-use		File-System	pool1	VM-111
vol4	200 GB	In-use		File-System	pool1	VM-111
vol_iozone	10 GB	Available		File-System	pool1	

Adjusting Anti-Virus Protection Policy The default (that applies to all Shares) can be modified for your needs on the Anti-Virus page as follows:

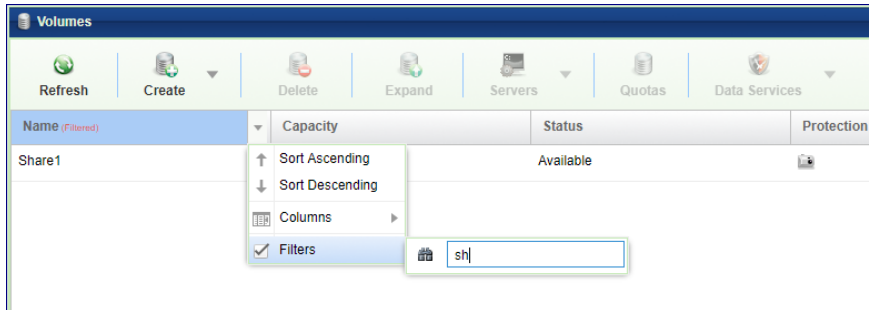
- Go to the [Antivirus](#) page
- **Scan Schedule** - Either select one of the pre-defined schedules (e.g. Everyday after midnight) or custom your own
- You can specify the day of the week, or the date of the month, and specify the time
- **Scan Subfolders** - Specify if you want the AV to scan files in subfolders
- **Scan Archives** - Specify if you want the AV to scan files within archive files such as ZIP, TAR, etc...
- **Primary & Secondary Actions** - Specify the actions you want the AV engine to take: (If the primary action failed, the AV engine tries the secondary action)
 - Clean - remove the virus from the infected file
 - Delete the infected file from the NAS share, and move it into the quarantine
 - Ignore the unaffected file and continue the scan
- **File Types to Scan** - Either All files, or specify specific file types extensions
- **Exclude File Types** - If all was selected, you have the options to specify file extensions to skip
- click Update to keep the changes made to the policy.

Managing Quarantined files The Infected Files pane lists all the files that were found infected and were removed to the quarantine. Review these files, select the rows in interest and take one the following actions: - Delete the files - Restore the files into their original folder

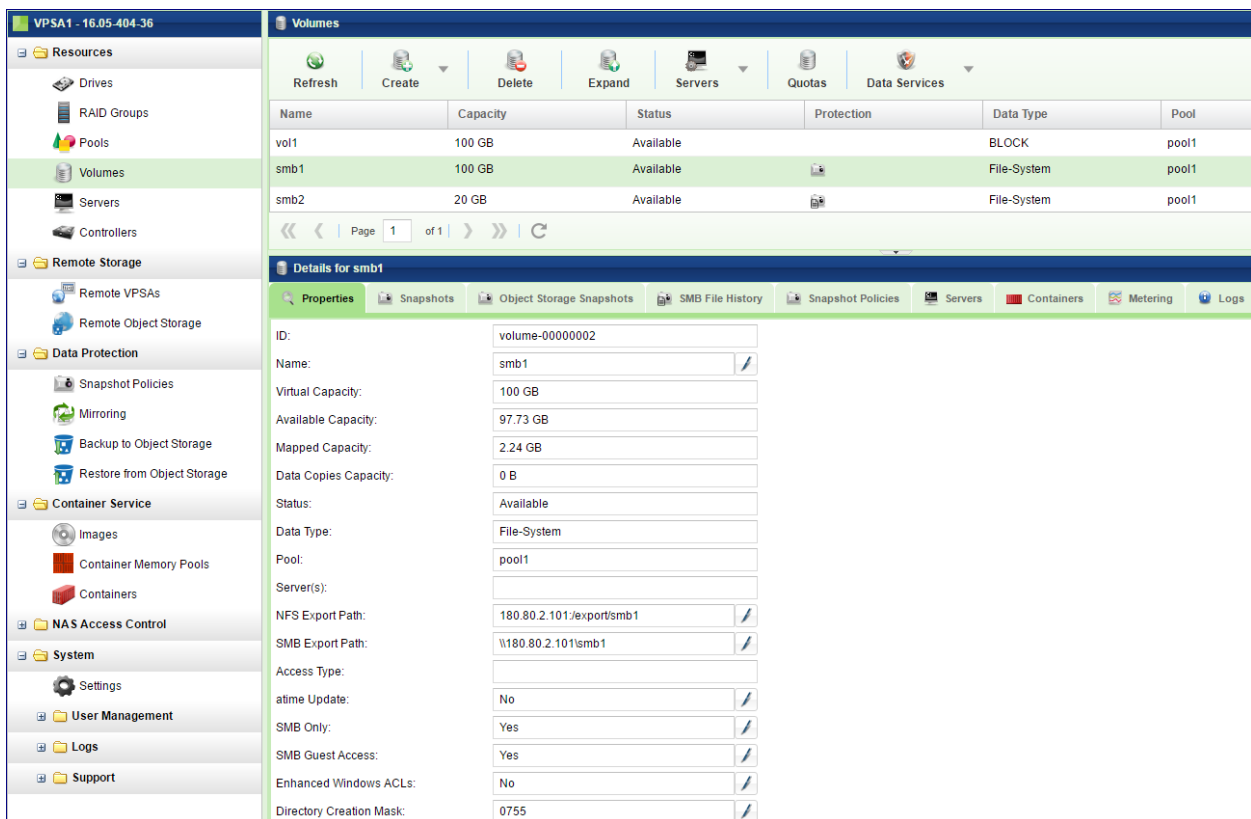
7.10 Viewing Volume Properties

Filtering Volumes

In a VPSA with many volumes it might be difficult to locate a specific volume in the [Volumes](#) page. The following Filtering option may be useful. In the [Volumes](#) page click the little arrow on the Name column title, select Filter and start typing the name of the volume of interest. The table will be filtered accordingly.



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:



Properties

Each Volume includes the following properties:

Property	Description
ID	An internally assigned unique ID.

continues on next page

Table 1 – continued from previous page

Property	Description
Name	User assigned name. Can be modified anytime.
Comment	User free text comment. Can be used for labels, reminders or any other purpose
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
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.
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

continues on next page

Table 1 – continued from previous page

Property	Description
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.
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
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
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 multipathing is configured.
Encrypted	Yes/No
Created	Date & time when the Volume was created.
Modified	Date & time when the Volume was last modified.

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 [here](#).

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

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 here [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

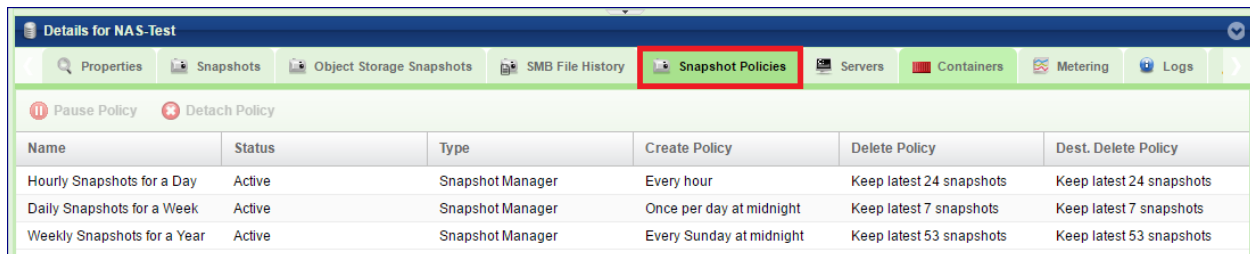
SMB File History (SMB Only)

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 [here](#).

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

Snapshot Policies



Name	Status	Type	Create Policy	Delete Policy	Dest. Delete Policy
Hourly Snapshots for a Day	Active	Snapshot Manager	Every hour	Keep latest 24 snapshots	Keep latest 24 snapshots
Daily Snapshots for a Week	Active	Snapshot Manager	Once per day at midnight	Keep latest 7 snapshots	Keep latest 7 snapshots
Weekly Snapshots for a Year	Active	Snapshot Manager	Every Sunday at midnight	Keep latest 53 snapshots	Keep latest 53 snapshots

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 [here](#).

Servers

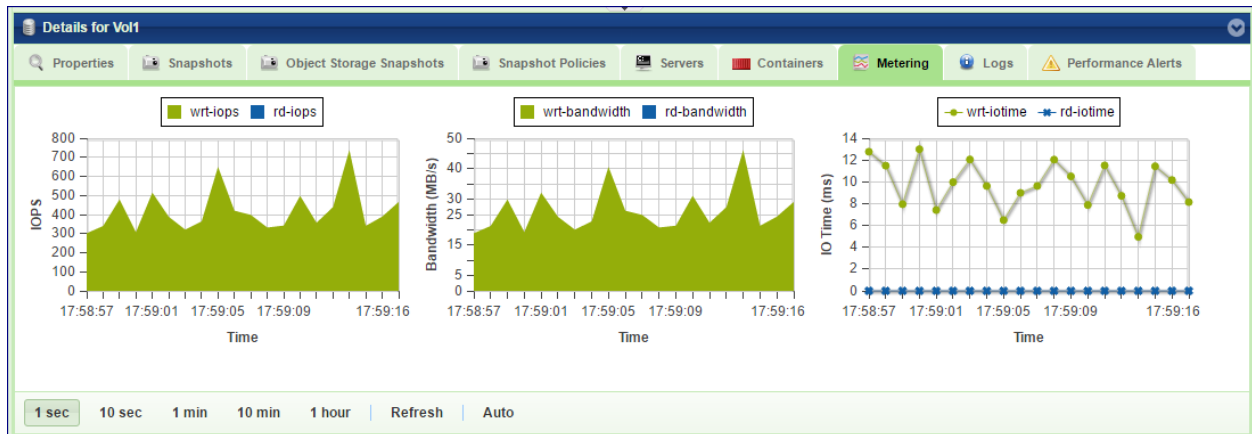
The Servers tab lists the Servers to which the Volume is attached. For Block Volumes the LUN Number associated with each Server is displayed. It also indicates if the server accesses the volume via iSCSI or FC.

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 [Managing Container Services](#)

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.

Logs

Displays all event logs associated with this Volume.

Title	Time
★ Snapshot creation on Volume fc_vol1 postponed because previous snapshot still deleting.	2016-04-13 23:11:08
★ Snapshot creation on Volume fc_vol1 postponed because previous snapshot still deleting.	2016-04-13 22:11:12
★ Snapshot creation on Volume fc_vol1 postponed because previous snapshot still deleting.	2016-04-13 21:11:11
✔ Snapshot Policy e1 was successfully attached to Volume fc_vol1.	2016-04-13 18:36:54
✔ Volume fc_vol1 was successfully attached to Server SERVER-204, WWPN [21000024ff7bc174.21000024ff7bc175], IP [180...	2016-04-13 18:33:59
✔ Volume fc_vol1, of size 60 GB and access mode BLOCK, was created successfully.	2016-04-13 18:32:57

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.

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

7.11 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

Go to The Volumes page, select a Volume and display the Snapshots tab in the South Panel. Press the Filter button at the bottom of the page. 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.

The screenshot displays the VPSA Storage Array management interface. At the top, there is a 'Volumes' section with a table listing various volumes. Below this, the 'Details for vol-SQL' section is visible, showing a list of snapshots. A 'Filter Snapshots' dialog box is open, allowing users to filter snapshots based on date and time. The dialog box includes fields for 'From Date/Time' (11/21/2013 07:00) and 'To Date/Time' (11/21/2013 10:00), an 'Origin' dropdown set to 'User', and a 'Snapshot Policy' table. The 'Filter Snapshots' dialog box is highlighted with a red border. At the bottom of the interface, a 'Filter' button is also highlighted with a red border.

Name	Capacity	Status	Data Type	Pool	Server(s)
vol-SQL	200 GB	In-use	BLOCK	pool_DB	AMAZONA-1F01PAM
Web_files	1000 GB	In-use	File-System	Pool_Webfiles	AMAZONA-1F01PAM
vol-SQL-test	200 GB	Available	BLOCK	pool_DB	
Web-logs	250 GB	Available	File-System	Pool_Webfiles	
test-share	1 GB	In-use			

ID	Name	Status
snap-00000036	s-r3-r7-1385046583732	normal
snap-00000035	s-r7-1385045743663	normal
snap-00000034	s-r7-1385045443624	normal
snap-00000033	s-r7-1385045143578	normal
snap-00000032	s-r7-1385044843532	normal
snap-00000030	s-r7-1385044543486	normal
snap-0000002f	s-r7-1385044243433	normal
snap-0000002e	s-r7-1385043943378	normal
snap-0000002d	s-r7-1385043643304	normal
snap-0000002c	s-r7-1385043343259	normal
snap-0000002b	s-r7-1385043043218	normal
snap-0000002a	s-r7-1385042743184	normal

Name	Create Policy	Delete Policy
Hourly Snapshots for a Day	Every hour	Keep latest 24 snapshots
Daily Snapshots for a Week	Once per day at midnight	Keep latest 7 snapshots
yearly Weekend snapshots	Weekend days at midnight	Keep latest 104 snapshots
every 5 minutes for an hour	Every 5 minutes	Keep latest 12 snapshots

MONITORING PERFORMANCE

8.1 Understanding Performance Monitoring

This chapter contains instructions to monitor 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 the 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 operates properly. The VPSA performance Monitor is a tool for pinpointing a storage performance bottlenecks. The following metrics are of interest to measure 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 means the amount of read or write operations done in one seconds 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.

8.2 The Performance Monitor

To open the VPSA Performance Monitor click the [VPSA GUI > 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

VPSA FLASH ARRAY

For the All 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, etc...

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:** Toggles between number of supported layout with different number of tails.
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 examples: If 1 minute interval was selected 60 points are displayed, each one is the average value for that specific minute. In total the last 1 hour is displayed.

The interval selection affects all tiles.

8.3 Customizing the Performance Monitor

8.3.1 Customizing the Layout

- Go to [VPSA GUI > Performance](#) and click the Layout selector




- Select the layout of your choice. Note that if the selected layout has fewer tiles than the original the other tiles will be lost, and should be set again.
- Drag the object of interest from the resources tree, and drop it into a tile. Do the same for all tiles.

8.3.2 Customizing a Tile

Each tile represent a single resource, and provide 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 the Charts button on the top right corner of the tile, and select the metric of interest.



- To display a table click the Tables button on the top left corner of the tile, and select the table of interest. The table provide performance information as well as other parameters such as **data reduction ratio**.

The screenshot shows a performance monitoring tile titled "Pool Summary". It features a "Tables" dropdown on the top left and a "Charts" dropdown on the top right. The main area is a table with the following data:

Ratio		
	Compression Ratio	Total
Performance Summary		
Provisioned Capacity		
Garbage Collection Capacity	7.4 : 1	7.4 : 1
Effective and Data Reduction Capacity		
Data Reduction Ratios		

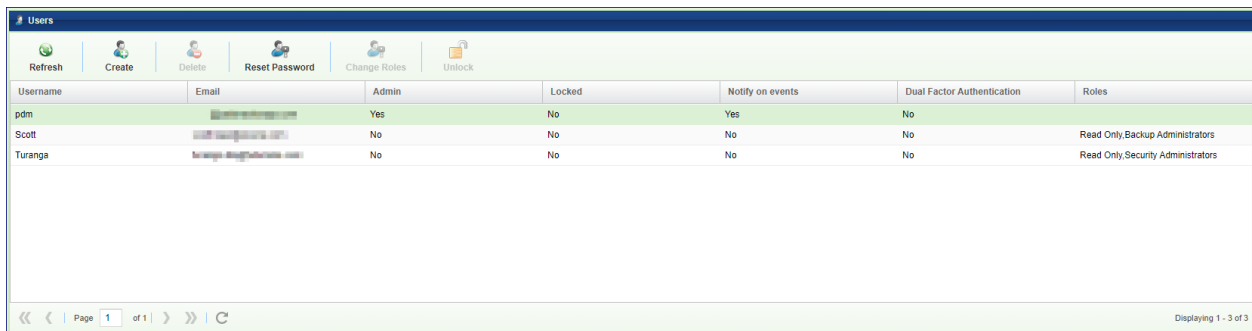
✓ **Note:** Some of the performance metering charts and table are for Zadara support use only.

MANAGING ACCESS CONTROL

9.1 Adding & Deleting 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 its own Password and Access Key.



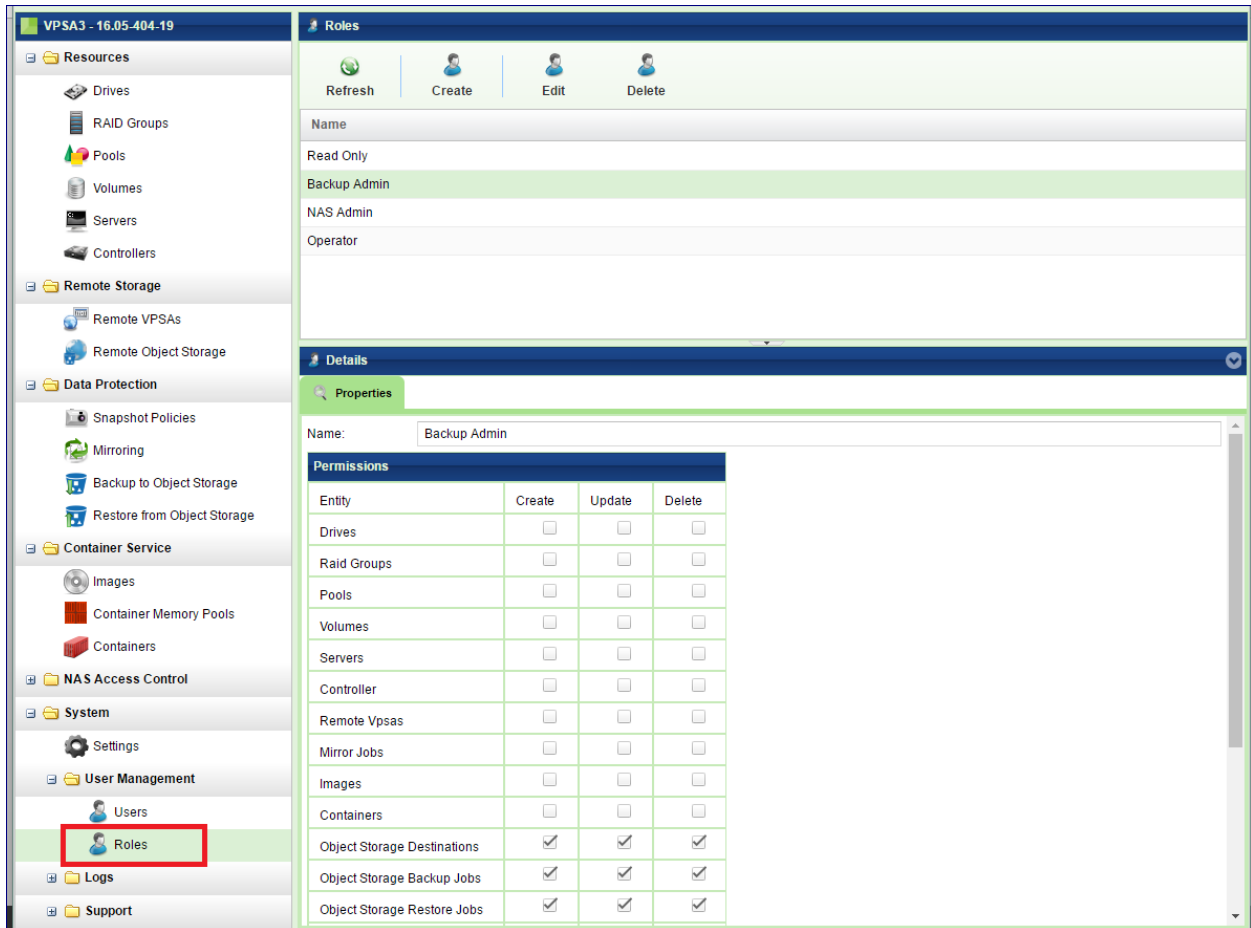
The screenshot shows the 'Users' management interface. At the top, there are navigation buttons: Refresh, Create, Delete, Reset Password, Change Roles, and Unlock. Below these is a table with the following columns: Username, Email, Admin, Locked, Notify on events, Dual Factor Authentication, and Roles. The table contains three rows of user data.

Username	Email	Admin	Locked	Notify on events	Dual Factor Authentication	Roles
pdm	[redacted]	Yes	No	Yes	No	
Scott	[redacted]	No	No	No	No	Read Only, Backup Administrators
Turanga	[redacted]	No	No	No	No	Read Only, Security Administrators

At the bottom of the table, there is a pagination control showing 'Page 1 of 1' and a refresh icon. The status 'Displaying 1 - 3 of 3' is visible in the bottom right corner.

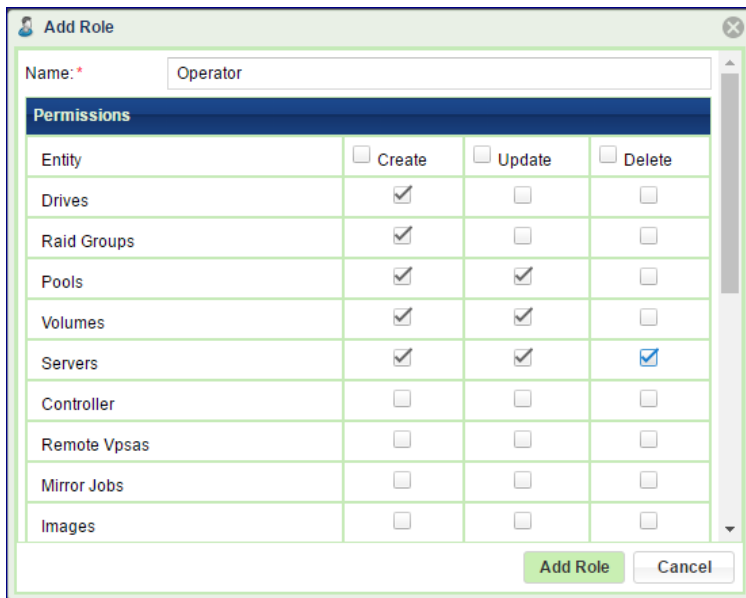
User Roles

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.



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 the Add Role button.



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 the Add User button.

Enter the Username and Email address and specify if this new User will be assigned the 'Admin' privilege (full control), or select specific Role. 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.

Changing a User 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 Role button.

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.

9.2 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.

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 the Change Password button.



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.

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 the Reset Password button.

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 button.

9.3 Managing Password Policy

The VPSA Admin can control the VPSA Password Policy. For details, see VPSA settings [Security](#).

9.4 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.

9.4.1 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.

User Information	
<div style="display: flex; justify-content: space-around;"> Change password Reset Access Key </div>	
Username:	dima
Email:	dima@zadarastorage.com
Access key:	QIFH7NFWPZ01E2XJ1SJT
Dual Factor Authentication:	<input type="text" value="Activated"/> Deactivate

Click Activate or Deactivate. Close the properties dialog, and logout.

The first time you login again, the following screen will pop up.

Confirmation

To use dual factor authentication, please download an authentication application (i.g "Google Authenticator") to your mobile device and scan the following QR code:



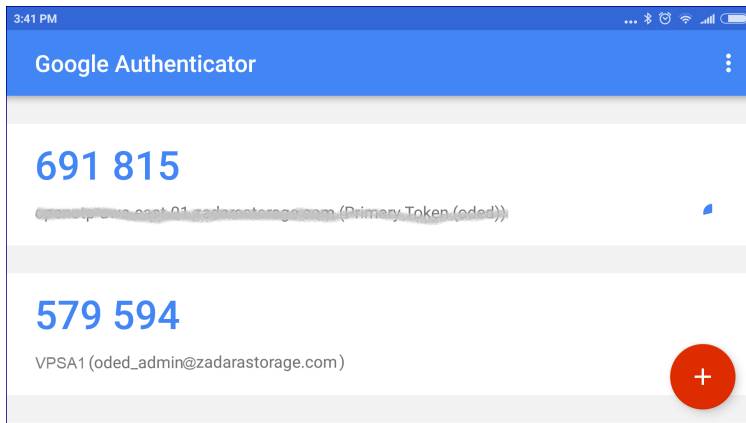
Alternatively, you may instead manually insert the following data into your application:

qa@zadarastorage.com
rel5vlkemqcsd7v4omhwwmf

Dual Factor Authentication
Token: *

Confirm
Cancel

Install 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. In case the device was lost or replaced, the user must ask the VPSA admin to reset her DFA settings. VPSA admin must contact Zadara support for resetting the DFA.

9.4.2 Enforcing Dual Factor Authentication

VPSA administrator can force DFA for all users. In setting/Security click Edit on the Dual Factor Authentication, check the checkbox and Save. This setting change does not have immediate effect. Next time each user will login, she will be required to set her mobile device Authenticator app as described above.

The screenshot shows the 'Settings' window with the 'Security' tab selected. The 'Enforce dual factor for all users' checkbox is checked and highlighted in yellow. Below it, a warning message states: 'Warning: Dual factor authentication will only be activated upon the next user log in. Please verify that all users who require dual factor authentication log in to their accounts.' A 'Save' button is located below the warning. Other settings include 'Passwords policy', 'Global VPSA CHAP', 'Encryption', 'Cloud Admin Access', and 'IPSec Key'.

Setting	Value	Action
Passwords policy:	Enforce password expiration: No, Password history: 8	Edit
Enforce dual factor for all users:	<input checked="" type="checkbox"/>	Close
Global VPSA CHAP:	User: eyalvpsa Secret: PIR3zD9CU5Yr	Edit
Encryption:	No password set	Edit
Cloud Admin Access:	Enabled	Edit
IPSec Key:	6D0E6A36B5724A4C83C79D93EA55278A	

✓ **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 her settings in the user properties as described above.

9.4.3 Creating NAS Users

9.5 Managing NAS Users Access Control

9.5.1 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:

Settings

General Security **NAS** Metering Container Service

NAS

Domain: Close

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 [VPSA 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.

Create NAS User

Username:

Authentication: NFS SMB

NFS UID:

9.5.2 Creating SMB Users

- Go to [VPSA 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

9.5.3 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:

- 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.

9.6 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.

- 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.

9.7 Enabling 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: Using Active Directory cannot be used while the VPSA configured to use LDAP service, the transition from LDAP to Active Directory based authentication should be handled carefully, as existing NAS permissions may be affected. In case you are considering such transition, contact Zadara support team for additional information.

9.7.1 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.

The screenshot shows a dialog box titled "Join Active Directory Server" with the following fields and options:

- Domain Name: * DOMAIN.COM
- Domain NetBIOS Name: * DOMAIN.COM
- Administrator Name: * Administrator
- Administrator Password: * *****
- DNS IP #1: * 192.168.1.100
- DNS IP #2 (optional):
- DNS IP #3 (optional):
- Advanced**
 - Active Directory UID Mapping:
 - Allow Trusted Domains: OFF
 - DNS Lookup realm: OFF
 - DNS Lookup kdc: ON
- Buttons: Submit, Cancel

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, gidNumber). 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, 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 VPSA shares, prior to joining the AD Domain.

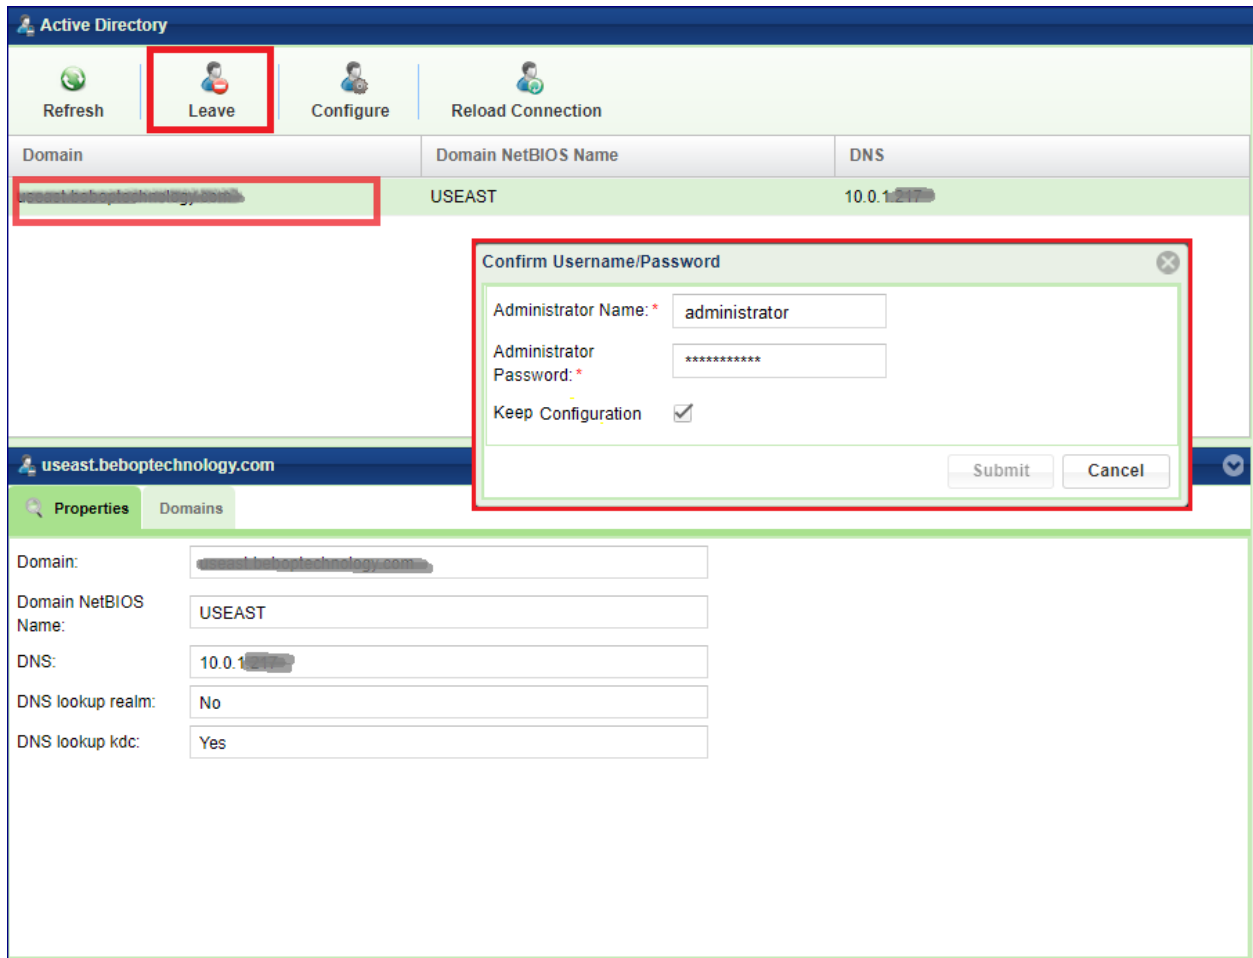
 **Note:** The joining of the VPSA to the Active Directory may fail if the time on the VPSA and the Active Directory Domain Controller is out of sync by more than a few minutes. Sync the time and try again. Different time zones are not an issue.

9.7.2 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 [VPSA 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).

9.7.3 Leaving an Active Directory

To leave the Active Directory, Go to [VPSA 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.

9.8 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.

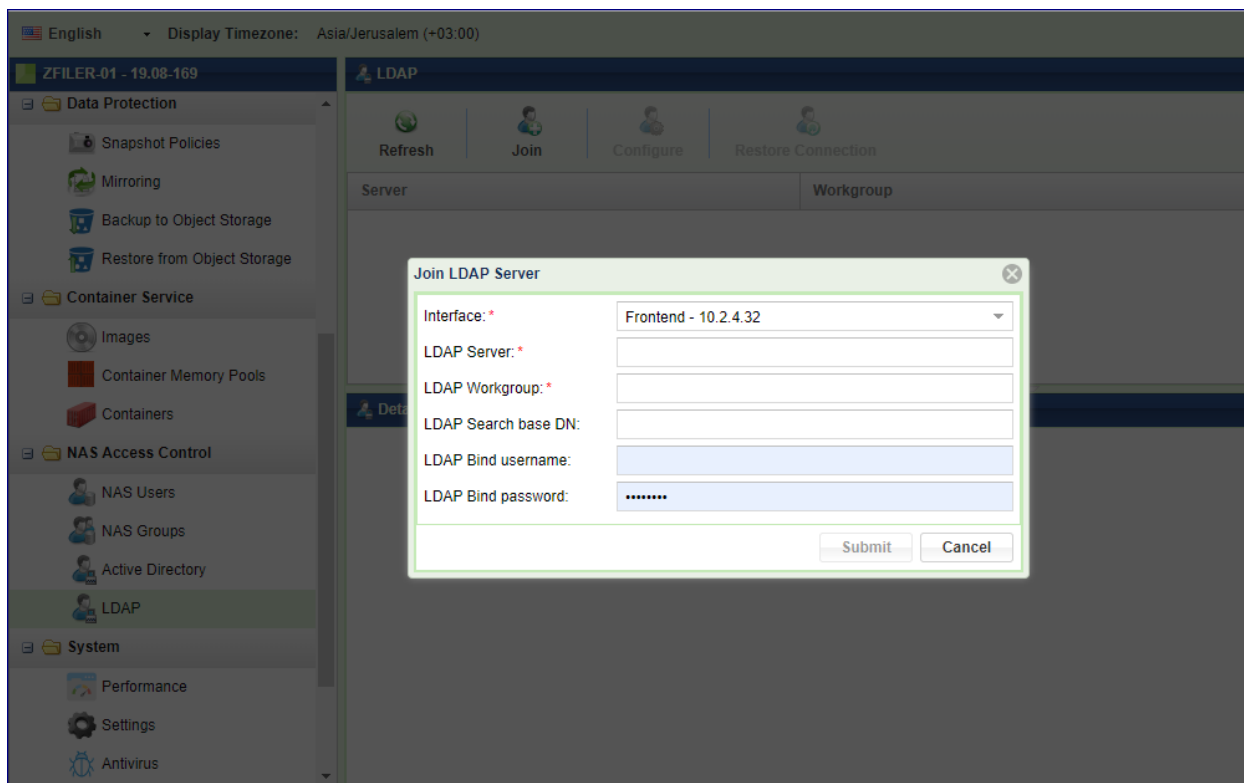
Starting from VPSA version 19.08, the VPSA SMB service can be configured to authenticate users against LDAP service (JumpCloud or similar).

✓ **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. In case you are considering such transition, contact Zadara support team for additional information.

9.8.1 Configuring LDAP service for NAS authentication

To enable the LDAP service navigate to [VPSA GUI > NAS Access Control > LDAP](#) and click the Join button.



Enter the following information:

1. **Interface** - the VPSA network interface that will be used for LDAP connectivity. In case of a 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)

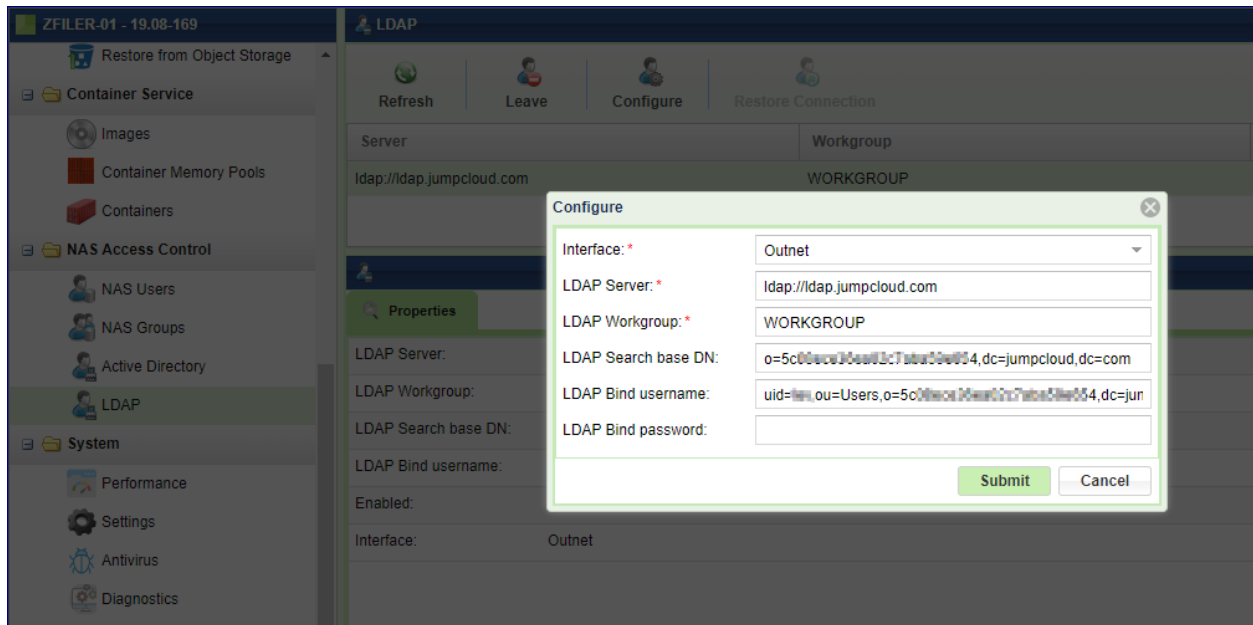
Note: In case of JumpCloud integration, Samba authentication should be enabled in the target directory. See <https://support.zadara.com/hc/en-us/articles/360036369912> for a KB article covering JumpCloud specific integration.

Click the Submit.

9.8.2 Updating LDAP configuration

In case the existing configuration needs to be changed, the directory parameters can be updated directly from the VPSA GUI.

Navigate to [VPSA GUI > NAS Access Control >LDAP](#) and click the Configure button.

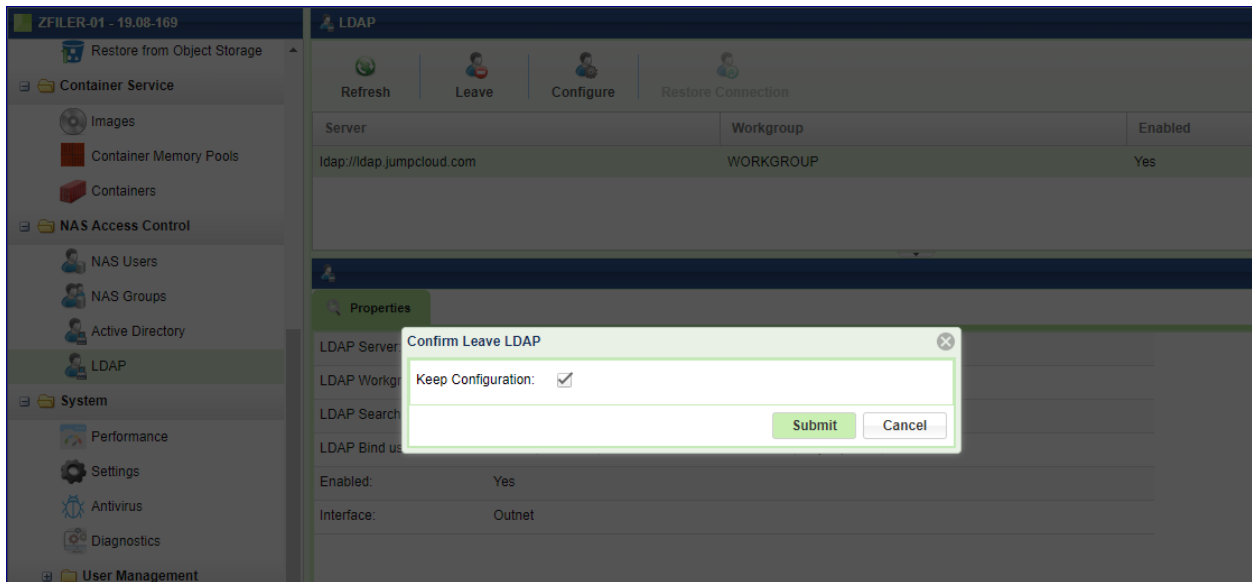


Once the configuration is submitted, the file services will be restarted in order to apply the new configuration.

9.8.3 Disable LDAP service SMB authentication

In case LDAP authentication is no longer needed, the LDAP authentication can be disabled from the VPSA GUI.

Navigate to [VPSA GUI > NAS Access Control >LDAP](#) and click the Leave button.



In case 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 [VPSA GUI > NAS Access Control > LDAP](#) and click the Restore button.

9.9 Managing NAS Quotas

9.9.1 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 tab in the VPSA Storage Array web 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 in green and has a red box around it. Below the table is a pagination control showing Page 1 of 1. The Details for test_nas_volume1 section is visible, showing various properties like Data Copies Capacity, Status, Data Type, Pool, Server(s), NFS Export Path, SMB Export Path, Access Type, atime Update, SMB Only, SMB Guest Access, Enhanced Windows ACLs, Directory Creation Mask, File Creation Mask, Map Archive, User Quotas, and Group Quotas. The User Quotas property is set to Off and is highlighted with a red box. A Quotas dialog box is open over the User Quotas property, showing the User Quotas section with radio buttons for Off (selected), On, and Account only. The dialog box has Submit and Cancel buttons.

Name	Capacity	Status	Protection
RecoPH	5 GB	In-use	
NFS-RecoPH	5 GB	Available	
test_nas_volume1	2 GB	Available	

Page 1 of 1

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 dialog box:
 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
```

9.9.2 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 displays the VPSA GUI interface. At the top, the 'Volumes' section is active, with a table listing volumes. The 'NAS1' volume is selected, showing a capacity of 5 GB and a status of 'Available'. The 'Quotas' button is highlighted with a red box. A modal dialog titled 'Quotas' is open, showing the 'Users' tab. The dialog includes a 'Default Limit' field set to 0, an 'Update defaults' button, and a table with columns for 'Type', 'Id', 'Usage (MB)', and 'Limit (MB)'. The table is currently empty, and the status 'No data to display' is shown at the bottom. The dialog also features a 'Refresh' button and an 'Add Records' button.

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

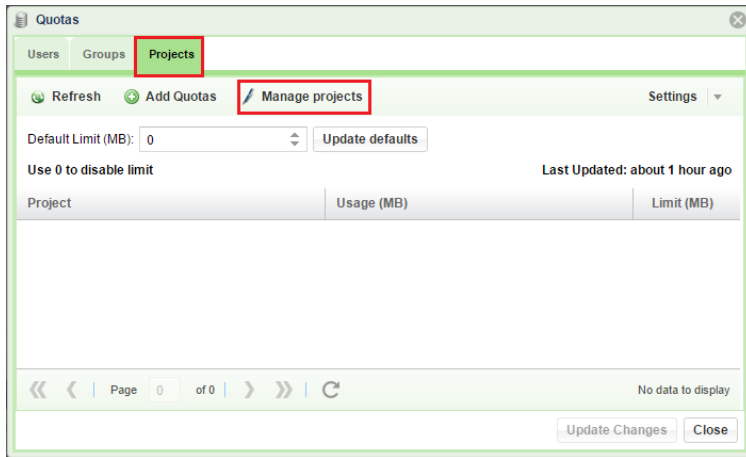
After making any additions or changes to Quota Limits, on the Quotas dialogue box press 'Refresh' to update the figures displayed.

9.9.3 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.

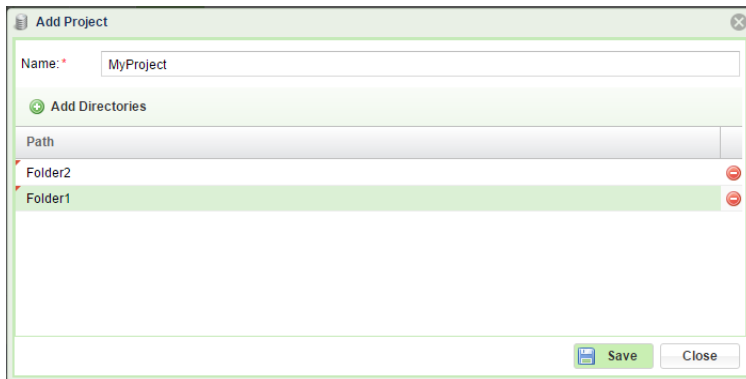
Defining Projects

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.

MANAGING REMOTE MIRRORING

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.

The VPSA 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 VPSA manages checkpoints to track the sync progress within a Volume/Snapshot. In case of a transport failure (line failure, VPSA failure etc.) the VPSA has a clear checkpoint from where to resume the sync.

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).

10.1 Creating a Local Mirror (on the same VPSA)

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 VPSA as the Destination VPSA 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 VPSA: *

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: *

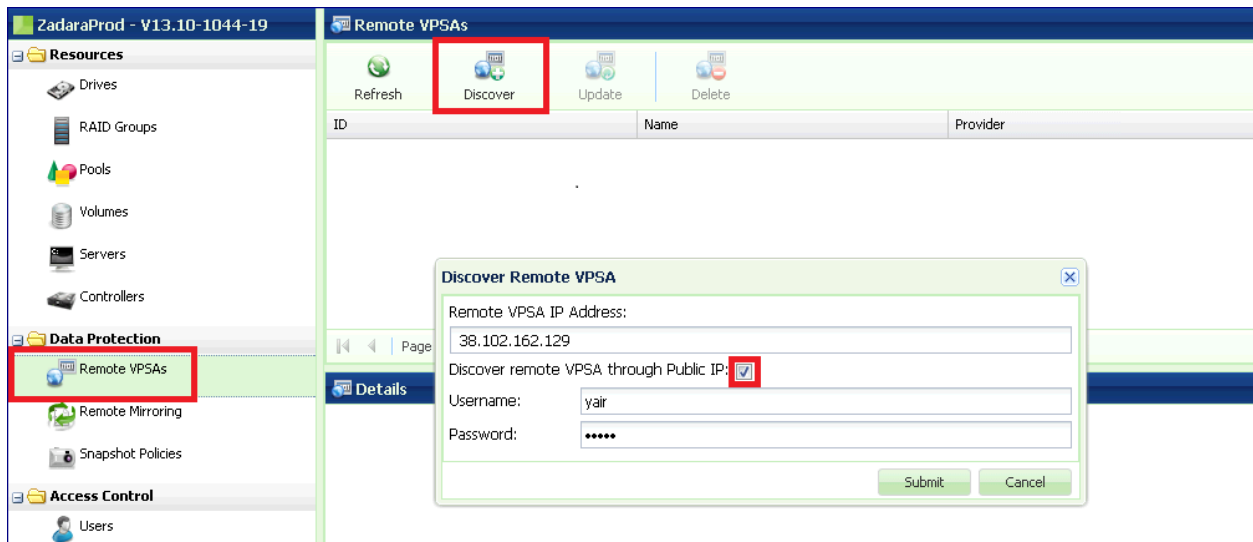
<input type="checkbox"/>	Name	Create Policy	Delete Policy	Dst. Delete Policy
<input checked="" type="checkbox"/>	Hourly Snapsh...	Every hour	Keep latest 24 ...	Keep latest 24 ...
<input checked="" type="checkbox"/>	Daily Snapshot...	Once per day a...	Keep latest 7 s...	Keep latest 7 s...
<input type="checkbox"/>	Weekly Snapsh...	Every Sunday ...	Keep latest 53 ...	Keep latest 53 ...

10.2 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.

Go 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.

10.3 Viewing remote VPSA Properties

The screenshot shows the ZADARA VPSA management interface. On the left, the 'Data Protection' section is expanded to show 'Remote VPSAs'. The main area displays a table of Remote VPSAs with the following data:

ID	Name	Provider	IP
vsa-000000de	ZadaraDR	aws3	38.102.162.128

Below the table, the 'Details for ZadaraDR' section is visible, showing a table of Pools:

Name	Total Capacity	Free Capacity	Mode	Stripe Size
pool_DB_DR	1.1 TB	1.1 TB	Stripe	64 KB
pool_webfiles_DR	1.1 TB	1.1 TB	Stripe	64 KB

You can view 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

10.4 Creating a Remote Mirror

You can create the Remote Mirror from the [Remote Mirroring](#) page by clicking Create. You will see a similar dialog:

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

Name the Mirror and the destination Volume, select the source Volume and click Next.

Create Mirror

Destination VPSA: *

WAN Optimization

I/O Performance Optimization

Destination Pool: *

Name	Total Capacity	Free Capacity	Version
RAID-10-Pool-1	3.47 TB	3.47 TB	3

Compress:

Dedupe:

Choose the Destination VPSA, WAN Optimization or I/O Performance Optimization (see below) and the destination Pool. Click Next.

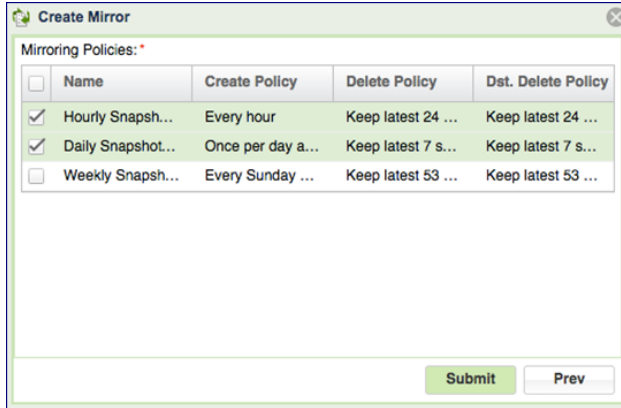
✓ Note: Mirroring a volume from a VPSA in a software version which is 19.08 and above to a target VSPA with a software version lower than 19.08 is not supported.

WAN Optimization v. I/O Performance Optimization

The VPSA supports selection between “I/O” and “WAN” data synchronization optimization. When “I/O” is selected, the VPSA synchronizes modified Pool chunks of 25KB or 1MB, depending on the Pool type. When “WAN” is selected, the VPSA 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.

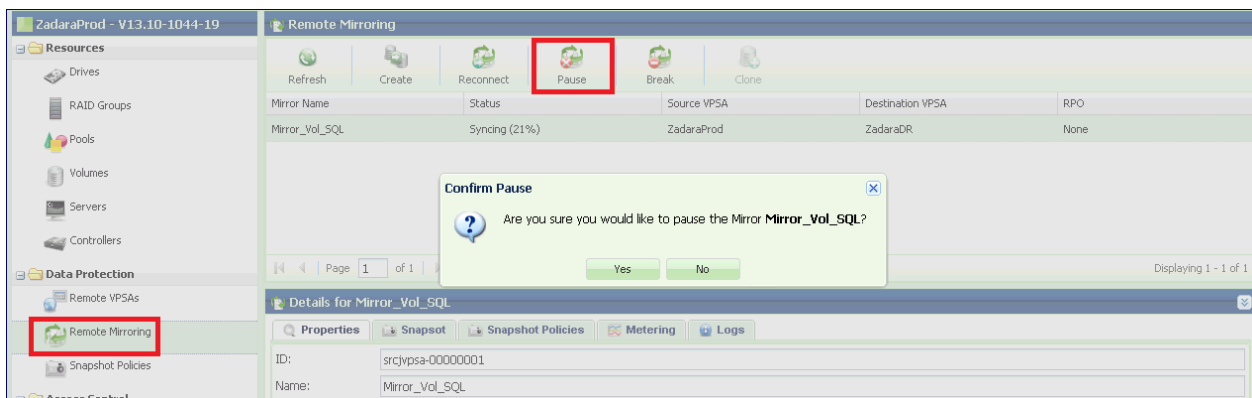
10.5 Replicate the same Volume to multiple destinations

It is possible to replicate the same Volume to multiple destinations. Just repeat the above steps, selecting a different destination each time.


10.6 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.



10.7 Managing Mirror Lifecycle

The Mirror controls the Remote Volume on the destination VPSA. As long as the Mirror is active it 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.

10.7.1 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.

10.7.2 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.

10.7.3 Reconnecting a Mirror

As previously described, the VPSA 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 Calchup	Paused	Remote Mirroring (NAS_SATA_VOL2_MIRR...	Once per day at midnight	Keep latest 45 snapshots	Keep latest 10 snapshots

This metadata allows the VPSA to identify a remote Volume, on a Remote VPSA, 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 VPSA 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				
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

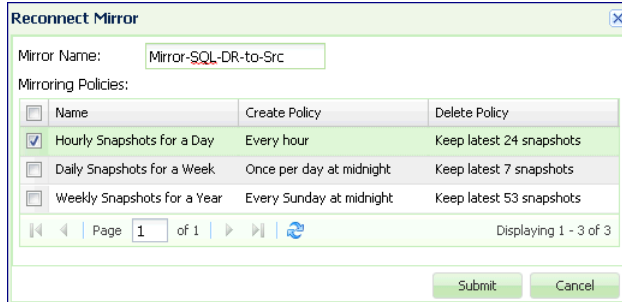
Page 1 of 1 Displaying 1 - 1 of 1

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.



✓ **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).

10.8 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:

The screenshot displays the VPSA Storage Array management interface. On the left, a navigation sidebar lists various system components, with 'Mirroring' highlighted in red. The main content area is titled 'Mirroring' and features a toolbar with icons for Refresh, Create, Reconnect, Pause, Break, and Clone. Below the toolbar is a table listing mirrors:

Mirror Name	Status	Source VPSA	Destination VPSA	RPO
VOLUME10_MIRROR	Idle	VPSA2	VPSA1	2016-04-2
SRM_MIRROR_cg-0000005b	Idle	VPSA2	VPSA1	2016-04-3
SRM_MIRROR_cg-0000005a	Idle	VPSA2	VPSA1	2016-04-2

Below the table is a pagination control showing 'Page 1 of 2'. The 'Details for SRM_MIRROR_cg-0000005b' section is visible, with tabs for Properties, Snapshots, Snapshot Policies, Metering, and Logs. The 'Properties' tab is active, showing the following details:

- 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

The 'Source VPSA' section includes:

- Provider: zadaraqa6
- VPSA: VPSA2
- Pool: pool1
- Volume: VOLUME3

The 'Destination VPSA' section includes:

- Provider: zadaraqa6
- VPSA: VPSA1
- Pool: pool1
- Volume: VOLUME3_MIRROR

Properties

Each Mirror includes the following properties:

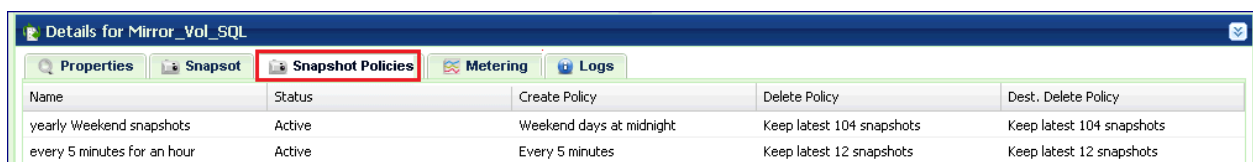
Property	Description
ID	An internally assigned unique ID.
Name	User assigned name. Can be modified anytime.
Comment	User free text comment. Can be used for labels, reminders or any other purpose
Status	<ul style="list-style-type: none"> • Idle – Mirror has nothing to Sync. • Failed • Paused • Syncing (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 tranfer 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.



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 VPSA.

The Source VPSA updates the Destination VPSA 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 VPSA**:

- **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 VPSA**:

- **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 VPSA.
- **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 VPSA.

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.

MANAGING REMOTE CLONES

Remote Clone feature makes a given snapshot of a source volume, instantly available (before data is copied) as a volume on another VPSA, in the same cloud or in a different cloud over any distance. Unlike Mirroring that might take a long time to replicate the data (depending on capacity and the link bandwidth), the cloned new volume is available immediately. Blocks of data are retrieved from the source volume on demand. During that time the clone behaves like any other volume, but it is dependant on its source. Once all the data was copied over, the relation between the volumes breaks, and the clone volume becomes a regular volume. In addition there is an option to have the system retrieve all the data in the background. IN this mode once all the data was retrieved, the connction to the original snapshot breaks, and the volume becomes independant.

Most common use cases are:

Instant Mobility - Rapid Migration of volumes between VPSAs with minimal downtime for the migrating application.

Useful for:

- Volume migration from a VPSA that runs out of capacity into a VPSA that has free capacity
- Volumes migration between Gen2 and Geb3
- Volumes migration between sites
- Volumes migration from Private cloud to the public cloud

How to:

- Stop the application
- Take a snapshot
- Create remote clone
- Attach the remote clone to a host
- Start the application on the new host

Offline Processing - Create an instant clone of a volume on another VPSA (local or remote) for offline processing without affecting the original production volume.

Useful for:

- Dev and Test
- Analytics
- Reporting

How To:

- Take a snapshot
- Create remote clone

- Attach the remote clone to a host
- Start the offline processing on the new host

11.1 Connect to a Remote VPSA

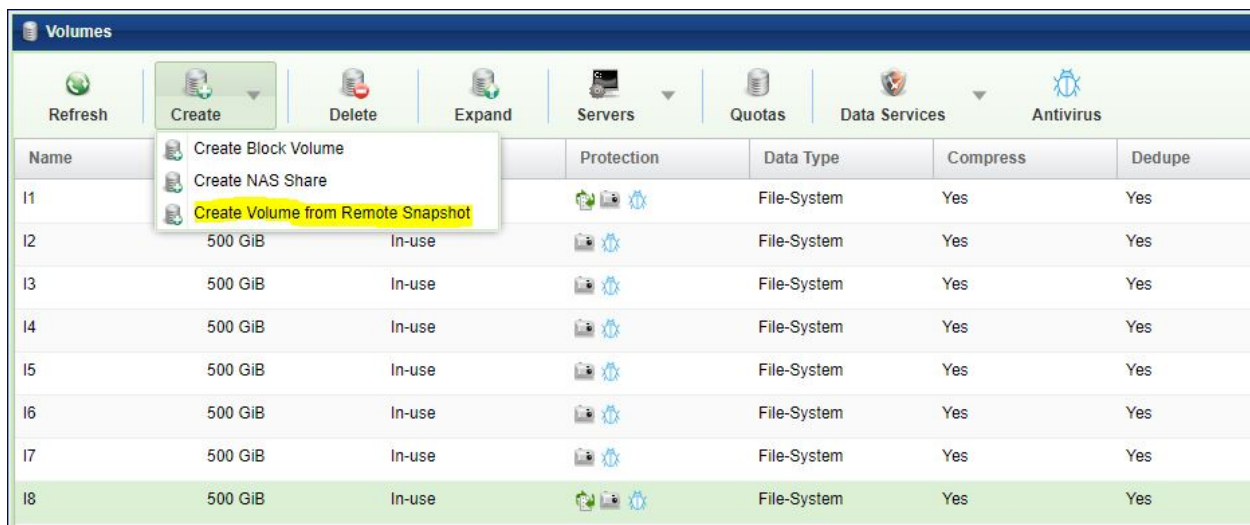
Clone can only be created on a pair of VPSA's known to each other. To establish connection between VPSAs, Remote VPSAs are discovered and defined the same way as done for remote mirroring.

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 here [Connect to a remote VPSA](#) to discover the remote VPSA and establish connection.

11.2 Create a Remote Clone

Remote Clone volume creation is done on the destination VPSA as a new type of volume. You can create the Remote clone from the [Volumes](#) page by clicking Create, and selecting [Create Volume from Remote Snapshot](#).



Volumes								
Refresh		Create	Delete	Expand	Servers	Quotas	Data Services	Antivirus
Name				Protection	Data Type	Compress	Dedupe	
I1					File-System	Yes	Yes	
I2	500 GiB	In-use			File-System	Yes	Yes	
I3	500 GiB	In-use			File-System	Yes	Yes	
I4	500 GiB	In-use			File-System	Yes	Yes	
I5	500 GiB	In-use			File-System	Yes	Yes	
I6	500 GiB	In-use			File-System	Yes	Yes	
I7	500 GiB	In-use			File-System	Yes	Yes	
I8	500 GiB	In-use			File-System	Yes	Yes	

The following dialog will open:

Source VPSA: * e_CHTHONIC

Source Volume: * smb1

Source Snapshot:

ID	Name	Timestamp
snap-00000530	s-r72-r73-r74-r76-156728...	2019-09-01 00:30:34
snap-00001250	s-r72-r73-r74-r76-156789...	2019-09-08 00:00:15
snap-00001778	s-r72-r73-r74-1568497260...	2019-09-15 00:41:00
snap-00001a3d	s-r72-r73-r76-r80-156892...	2019-09-20 00:20:28
snap-00001c8f	s-r72-r76-r80-1569012684...	2019-09-20 23:51:32

Page 1 of 8 | Filter | Displaying 1 - 5 of 36

Next Cancel

- Select the remote VPSA to clone from
- Select the volume to clone
- Select the Snapshot to use from the list of snaps of the selected source volume
- Press Next

New Volume Name: * MyNewClone

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:

Attach Default Snapshot Policies:

! 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

- Give the new volume a name
- Select the pool of the new volume
- Check if the new volume should be encrypted or not. The clone volume can be encrypted (even if the source is not).
- Check if the new volume should be deduped/compressed or not. The clone can be deduped/compressed (even if the source is not)
- Select the clone mode: (On-Demand: only the needed data is retrieved, Background: the rest of the data is retrieved in the background)

- Press Create

11.3 Monitoring Remote Clone

- Clone volumes are listed in the volumes table with a special identifier.
- As long as in retrieval mode each remote clone has a status: connected or disconnected
- If the connectivity to the source snapshot was lost during retrieval the clone becomes unavailable.
- Clone volume properties are identical to a regular volume.
- Capacity of a clone shows the virtual capacity of the original volume, and the physical actual capacity of the clone volume.
- The status of data retrieval can be seen on the south panel Remote Snapshot Status as long as the data is being copied between the original volume to its clone.
- You can pause/resume the data transfer in case there are load problems on the system.
- You can break the connection between the clone and its origin.
- Once all the data was retrieved, and the relationship between the source and the clone has been broken, the clone becomes a regular volume.

Volumes

Refresh Create Delete Expand Servers Quotas Data Services Antivirus

Name	Capacity	Status	Protection	Data Type	Compress
BV27	1 TiB	Available		BLOCK	Yes
BV28	1 TiB	Available		BLOCK	Yes
BV29	1 TiB	Available		BLOCK	Yes
BV30	1 TiB	Available		BLOCK	Yes
smb1	1.99 TiB	Available		File-System	Yes
nas	250 GiB	In-use		File-System	Yes
Clone1	1.00 TiB	creating		File-System	Yes

Page 1 of 1

Details for Clone1

Properties Remote Snapshot Status Snapshots Object Storage Snapshots SMB File History Snapshot Policy

Break Pause Continue

General		Source Information	
Name:	dstrclone-00000001	VPSA Name:	e_R1
Status:	Syncing	Volume ID:	
Mode:	Background	Volume Name:	nas1
Progress:	1013 MiB done of 1026 MiB (98.7%) 0 seconds left	Snapshot Name:	s-r3f-r4f-1568838629975
		Created At:	2019-09-26 16:47:47
		Modified At:	2019-09-26 16:47:48

11.4 Attaching a Clone

Remote Clone is attached to servers on the target VPSA as if it is a regular volume. Similarly, Detaching a volume from a server is done the same way as if it is a regular volume. Follow the instructions here: [Attaching & detaching Volumes to Servers](#)

11.5 Data Services

The following data services are not available on clone volumes while the data is retrieved:

- Snapshots
- Clones
- Mirrors
- B2OS
- Anti Virus

Once all the data was retrieved, and the relationship between the source and the clone has been broken, all data services become available.

11.6 Deleting a Clone

Remote Clone Volume deletion is done on the destination VP SA, and it breaks the relations with the source Volume. you can delete a Remote Clone the same way you delete any other volume. Follow the instructions here: [Creating and Deleting a Volume](#).

BACKUP TO OBJECT STORAGE

Zadara VPSA provides built in backup and restore capabilities to AWS S3, Google Cloud Storage, Zadara VPSA Object Storage or any other S3 compatible object storage. The backup process involves transporting VPSA Snapshots to the remote Object Storage for safe keeping.

12.1 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

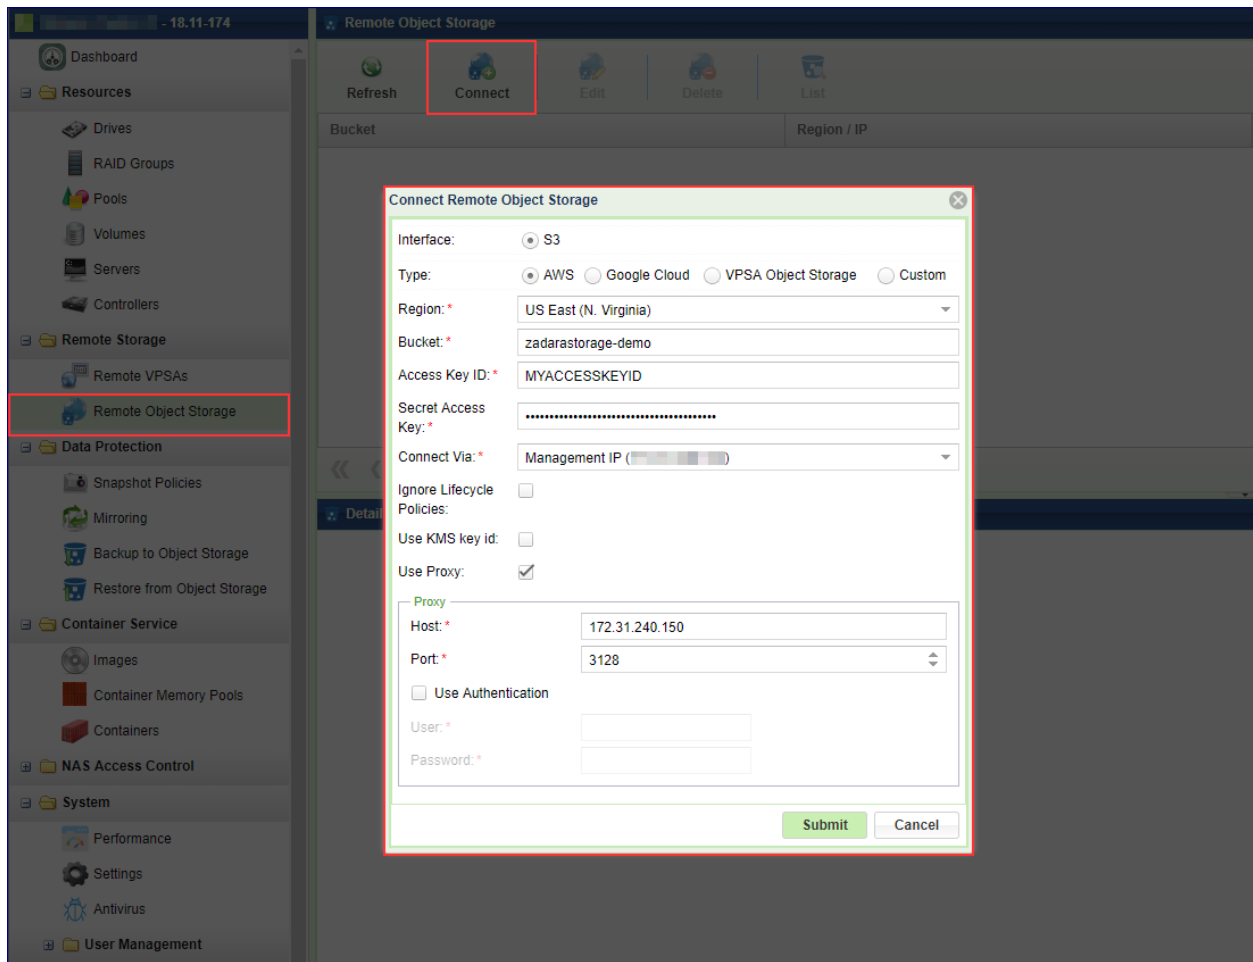
✓ **Note:**

- 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:



- Open the [VPSA GUI > Remote Object Storage](#) and click the Connect button.
- Select between AWS S3, Google Cloud Storage, VPSA Object Storage or Custom (S3 Compatible Object Storage).
- Enter the bucket/container name, access key and secret key.
- Select the connection method – via public IP, or the local management network.
- 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](#)

- In case 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)
- Press Submit

12.2 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, VPSA 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

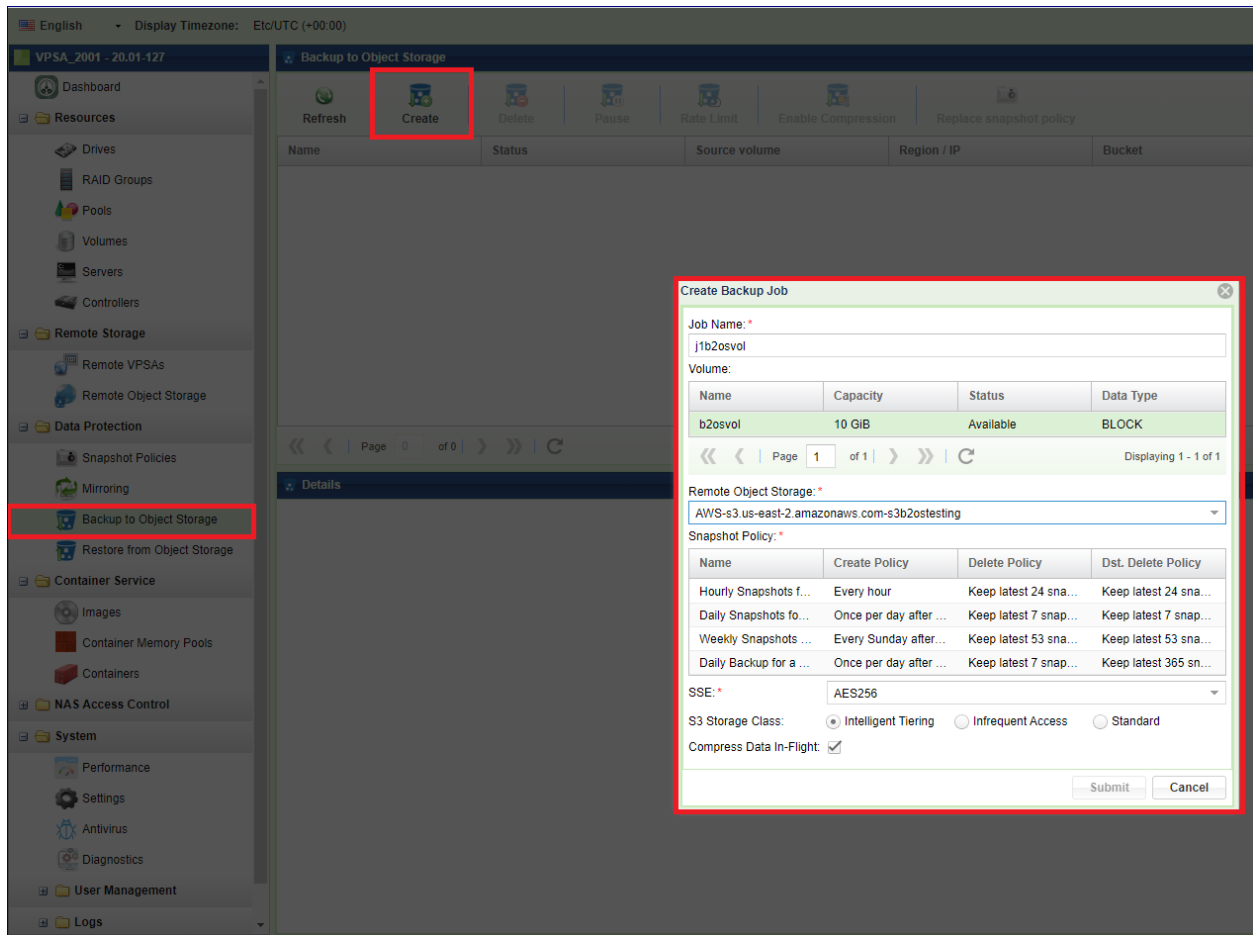
Restore Jobs Tab - List of all restore jobs using the selected Remote Object Storage

Logs Tab - List of event log messages related to that Remote Object Storage

12.3 Creating New Backups

In order to create a Backup for a given Volume, you must first have the Remote Object Storage connected as explained here [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
- 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:** 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.

- Check the Compress Data box if you want to compress the data in flight. This may save on the traffic fees
- Press Submit

12.4 Monitoring Backups

Remote Object Storage Backups can be managed and monitored from the VPSA GUI.

Open the [VPSA GUI > 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):

- Delete the Backup Job
- Pause / Resume
- Enable / Disable compression
- Rate Limit - Limit the backup job bandwidth (MB/s)
- Change the Snapshot Policy of the Backup Job
- Add a comment to a backup job
- Change a backup job target S3 storage class (AWS S3 Only)

✓ **Note:** If target S3 Storage class settings is modified for a specific backup job the new class will be applied on backups taken after this changed was performed. Previously created backups copies will not be modified.

Name	Status	Source volume	Region / IP	Bucket	SSE	RPO	Snapshot Policy
j2b2zevol	Idle	h2zevol	s3 us-east-2.amazonaws.com	s3b2zotesting	AES256	None	Hourly Snapshots for a Day

General	Source	Destination
ID: b1qjots-00000002	Volume: h2zevol	Type: AWS_S3
Name: j2b2zevol		Account:
Comment:		Endpoint: s3 us-east-2.amazonaws.com
Status: Idle		Bucket: s3b2zotesting
SSE Type: AES256		
Storage Class: Intelligent Tiering		
KMS Key ID:		
Snapshot Policy: Hourly Snapshots for a Day		
RPO: None		
Rate Limit (MB/s): 0		
Compression: YES		
Created: 2020-01-08 13:52:09		
Modified: 2020-01-08 13:52:10		

The Backup Job details are shown in the following South Panel tabs:

General	Source	Destination
ID: b1qjots-00000002	Volume: h2zevol	Type: AWS_S3
Name: j2b2zevol		Account:
Comment:		Endpoint: s3 us-east-2.amazonaws.com
Status: Idle		Bucket: s3b2zotesting
SSE Type: AES256		
Storage Class: Intelligent Tiering		
KMS Key ID:		
Snapshot Policy: Hourly Snapshots for a Day		
RPO: None		
Rate Limit (MB/s): 0		
Compression: YES		
Created: 2020-01-08 13:52:09		
Modified: 2020-01-08 13:52:10		

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.

12.5 Restore

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.

The screenshot displays the VPSA GUI interface. On the left sidebar, the 'Restore from Object Storage' option is highlighted with a red box. The main window shows the 'Restore from Object Storage' dialog with the 'Create' button highlighted in red. Below the dialog, a 'Create Restore Job' dialog is shown, also highlighted in red. This dialog includes a 'Remote Object Storage' dropdown menu set to 'AWS-s3.amazonaws.com-...', a 'Directory Listing' section showing a tree structure of folders and files, and 'Next' and 'Cancel' buttons at the bottom right.

✓ **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>

Create Restore Job

Remote Object Storage: *
s3.amazonaws.com-undefined

Browse destination
 Enter path manually

cloud_name.cloud_uid/vpsa_name.vpsa_id/volume_name.volume_id/object_snapshot_name
aws2.10C332C3DBFF42B7AE92E5C0FB856388/Test_VPSA_vsa-0000380/nas1.volume-0000003/20

Next Cancel

The Restore Job creates a new Volume from the selected Snapshot. Restore supports three modes of operation:

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.

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.

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.

Create Restore Job

Volume name: restored_vol

Mode: *

Restore - Full-capacity clone. Volume can be immediately attached to Servers. Data is retrieved from Object Storage on-demand and in a background process.

Clone - Zero-capacity clone. Volume can be immediately attached to Servers. Data is retrieved from Object Storage only on-demand when accessed by the attached Servers.

Import Seed - Full-capacity clone, including snapshot time-stamping. Volume can be attached to Servers only after the Volume's data was fully retrieved from Object Storage. Use this mode to import initial data seed for Remote Mirroring.

Encrypted:

Destination Pool: *

Name	Free Capacity	Status
pool1	0 B Free / 5.36 TB	normal

Page 1 of 1 | Displaying 1 - 1 of 1

Submit Cancel

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.

MANAGING CONTAINER SERVICES

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.

13.1 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 here [Registering a Zadara Account & 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) ▼

00

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

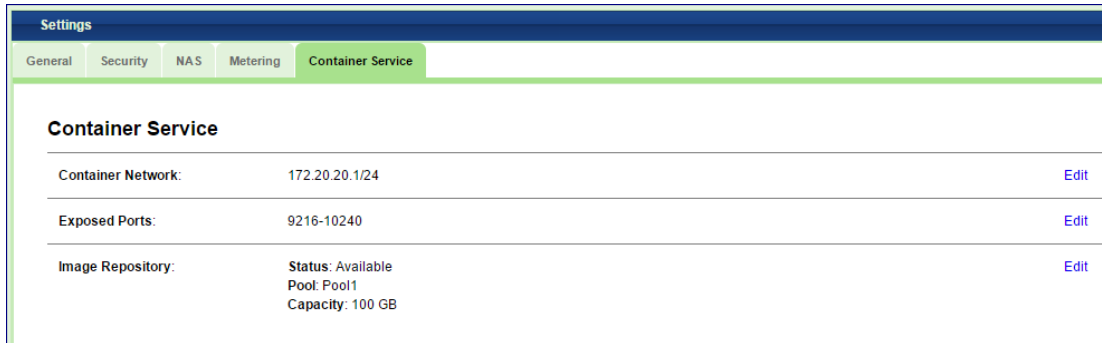
Adjust Cache

Hibernate

13.2 Creating 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.

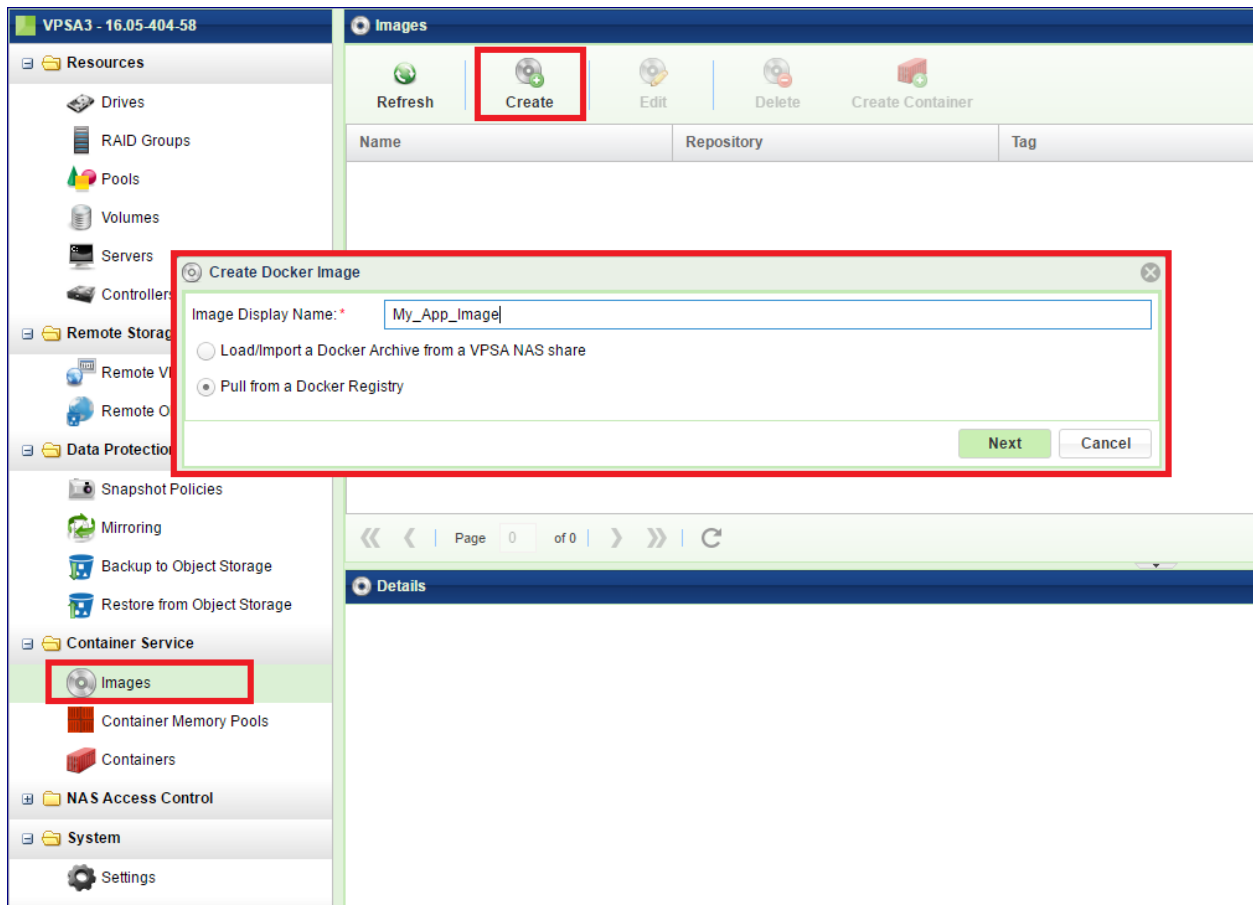
To create an Image Repository, open the [VPSA GUI > Settings > Container Service](#) tab and click Edit on the Image Repository section. Select the Pool that will host the Image Repository.




13.3 Creating 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.

Name	Repository	Tag	Status	Capacity
My_first_image	zadaraq/zubnutu	latest	normal	490 MB

13.4 Creating a Container Memory Pool

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.

To create a Container Memory Pool open the [VPSA GUI > Container Memory Pools](#) and click Create .

The screenshot displays the VPSA GUI interface for managing Container Memory Pools. On the left, a navigation sidebar lists various system components, with 'Container Memory Pools' highlighted in red. The main panel shows the 'Container Memory Pools' section with 'Refresh', 'Create', and 'Delete' buttons. The 'Create' button is highlighted in red. A modal dialog titled 'Create Containers Memory Pool' is open, containing two input fields: 'Name: *' with the value 'my_pool' and 'Memory limit (MB):' with the value '100'. 'Create' and 'Cancel' buttons are at the bottom of the dialog. Below the dialog, a pagination bar shows 'Page 0 of 0' and a refresh icon. At the bottom, a 'Details' section is visible but empty.

Give the new Memory Pool a name and select the amount of memory to allocate to this Container Memory Pool. The combined total of all of the Pools' limits must be less than or equal to the memory size of the ZCS engine. Click Create .

13.5 Creating 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 VPSA is utilizing this technology to allow user applications to run within VPSA 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 VPSA, 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 [VPSA GUI > Containers](#) and click Create .

The screenshot displays the VPSA GUI interface for managing containers. The left sidebar contains a navigation menu with categories like Resources, Remote Storage, Data Protection, Container Service, NAS Access Control, System, and User Management. The 'Containers' option under 'Container Service' is highlighted with a red box. The main panel shows a 'Containers' section with a toolbar containing 'Refresh', 'Create' (highlighted with a red box), 'Start', 'Stop', 'Delete', and 'CPU Utilization' buttons. Below the toolbar is a table listing containers:

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

Below the table, there is a pagination control showing 'Page 1 of 1' and 'Displaying 1 - 3 of 3'. The 'Details for c3' panel is expanded, showing various configuration fields:

- 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	

13.6 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.

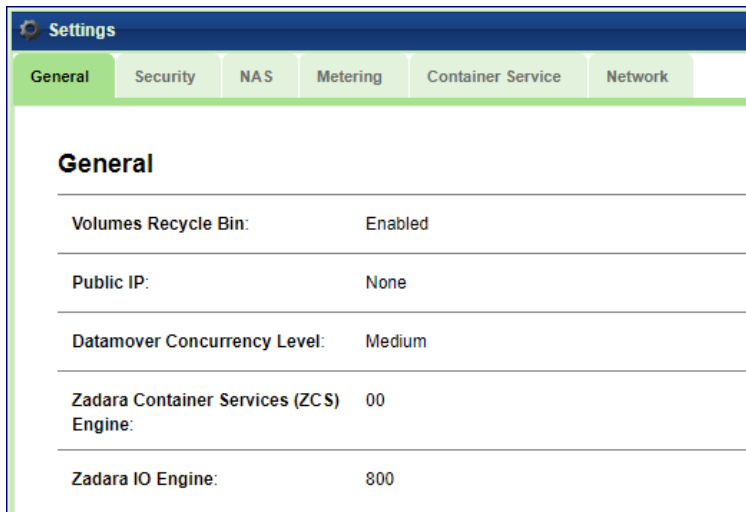
Metering

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.

SETTINGS

14.1 General



Settings					
General	Security	NAS	Metering	Container Service	Network
General					
Volumes Recycle Bin:	Enabled				
Public IP:	None				
Datamover Concurrency Level:	Medium				
Zadara Container Services (ZCS) Engine:	00				
Zadara IO Engine:	800				

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.

Public IPs

This displays any Public IPs assigned to the Controllers. A Public IP allows host connectivity from outside of the VPSA VPN.

Datamover Concurrency Level

You can control the load allowed for datamovers such as mirroring, cloning, etc... by setting the concurrency level. Default is Medium

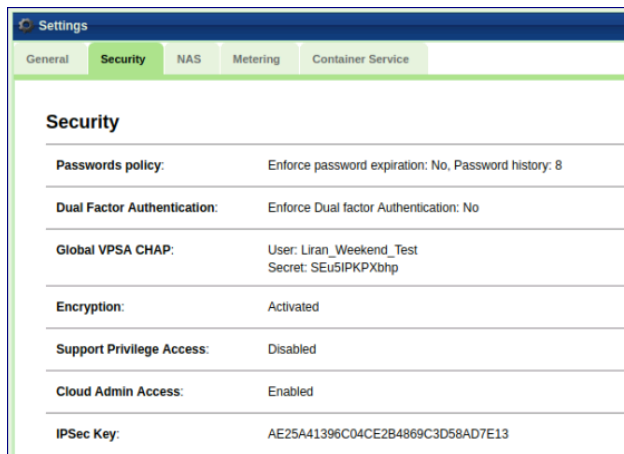
Zadara Container Services Engine

This displays which, if any, ZCS Engine has been configured via the Provisioning Portal.

Zadara IO Engine

This displays which VPSA IO Engine (Model) has been configured via the Provisioning Portal.

14.2 Security



Security	
Passwords policy:	Enforce password expiration: No, Password history: 8
Dual Factor Authentication:	Enforce Dual factor Authentication: No
Global VPSA CHAP:	User: Liran_Weekend_Test Secret: SEu5IPKPXbhp
Encryption:	Activated
Support Privilege Access:	Disabled
Cloud Admin Access:	Enabled
IPSec Key:	AE25A41396C04CE2B4869C3D58AD7E13

Password Policy

The VPSA Admin can control the VPSA Password expiration policy and password history policy.

Dual Factor Authentication The VPSA Admin can force all users to login to the VPSA GUI using dual factor authentication. For details see: [Dual Factor Authentication](#)

Global VPSA CHAP

This gives you a uniform username and password to use when you create Servers.

Encryption

This sets the Volume encryption password for the VPSA's data-at-rest encryption.

For more information on managing encrypted volumes see [Managing Encrypted Volumes](#).

Support Privilege Access

This controls the ability of Zadara support engineers to access the VPSA virtual controllers with privileged rights. Only the VPSA Admin can change this setting. If enabled, the VPSA Admin gets notification every time the privileged access is used.

Cloud Admin Access

This sets the cloud admin's VPSA GUI access (via the Command Center) to Enabled/Disabled status.

IPsec Key

This displays the key to be used when configuring IPsec tunneling for secured host connections.

14.3 NAS

Settings					
General	Security	NAS	Metering	Container Service	Network
NAS					
NFS Domain (NFS4 only):		localdomain			
NFS4 ID Mapping:		Disabled			
SMB NetBios name:		vsa-00000027			
SMB Character Set:		Unix: UTF-8 DOS: CP850			
Defragmentation:		Autorun: Enabled Status: Standby			
File System Trim:		Autorun: Enabled Status: Standby			

NFS Domain

This sets the domain name for NFS shares. This defaults to localdomain. (NFS4 Only)

NFS ID mapping

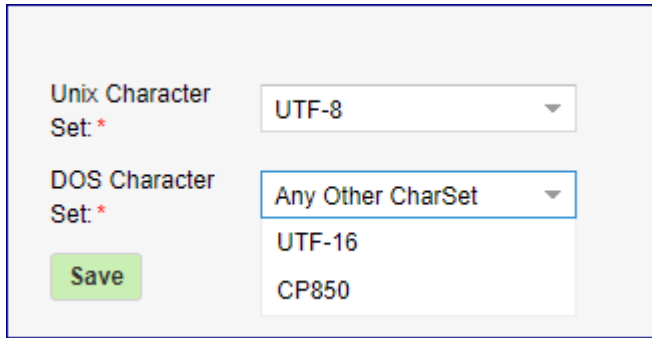
If enabled, each UNIX (Linux) User must be defined as a NAS User in the VPSA. If disabled, UNIX Users are authenticated on the UNIX host side.

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.

SMB Character Set

This gives you the default Character Sets used by the SMB service for SMB/CIFS 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.



Unix Character Set: * UTF-8

DOS Character Set: * Any Other CharSet

Save

UTF-16

CP850

Changing this value while clients are connected will cause them to temporarily lose access to all SMB shares.

Defragmentation

You can enable/disable background file system defragmentation. This also allows on-demand defragmentation.

File System Trim

This

- Periodic fstrim is triggered every weekend (On Saturday 00:00)
- Can be manually started and stopped via the settings page.

14.4 Metering

The VPSA 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.

14.5 Container Service

The screenshot shows the 'Settings' page with the 'Container Service' tab selected. The page displays the following configuration details:

Container Service	
Container Network:	172.20.20.1/24
Exposed Ports:	9216-10240
Image Repository:	Not Created

Container Network

This displays the internal IP range of the ZCS and is accessible only by the host VPSA.

Exposed Ports

This displays the Ports ranges that are exposed for host access.

Image Repository

This displays the Status, Pool and Capacity of the Image Repository that stores all of the ZCS containers and images.

14.6 Network

VPSA support Jumbo Frames. MTU size can be set to values from 1500 to 9000 for both the Front End (data) network, and the public network & VNIs (Public & VNI network MTU). Default is 1500 for both.

The screenshot shows the 'Settings' page with the 'Network' tab selected. The page displays the following configuration details:

Network	
FE MTU Size:	1500
Public & Virtual Networks MTU Size:	1500

Available MTU size options: 1500, 2048, 4096, 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.

DIAGNOSTICS

15.1 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 Diagnostic allows you to check connectivity over the selected network to any server:

- Go to the [Diagnostics](#) page.
- Select the VPSA network interface
- Enter the target IP address of the server in question.
- select if you want to ping the server, traceroute or both
- Click [Run](#).

It take a minute or so until the results show up in the output box.

English Display Timezone: Asia/Jerusalem (+03:00) [zadara_cloud_admin](#)

Dima_VPSA1 - 19.08.169 **Diagnostics**

- 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 Memory Pools
 - Containers
- NAS Access Control
- System
 - Performance
 - Settings
 - Antivirus
 - Diagnostics**
- User Management
- Logs
- Support

Network Diagnostics

Interface: Frontend - 10.2.9.28

Target Address: * 10.2.9.99

Ping:

Count:

Traceroute:

TTL:

```


Output
PING 10.2.9.99 (10.2.9.99) from 10.2.9.28 : 56(84) bytes of data.
From 10.2.9.28 icmp_seq=1 Destination Host Unreachable
From 10.2.9.28 icmp_seq=2 Destination Host Unreachable
From 10.2.9.28 icmp_seq=3 Destination Host Unreachable

--- 10.2.9.99 ping statistics ---
3 packets transmitted, 0 received, +3 errors, 100% packet loss, time 2041ms
pipe 3

-----
traceroute to 10.2.9.99 (10.2.9.99), 30 hops max, 60 byte packets
 1 ActiveFEIP (10.2.9.28) 3070.171 ms !H 3070.147 ms !H 3070.142 ms !H
    
```

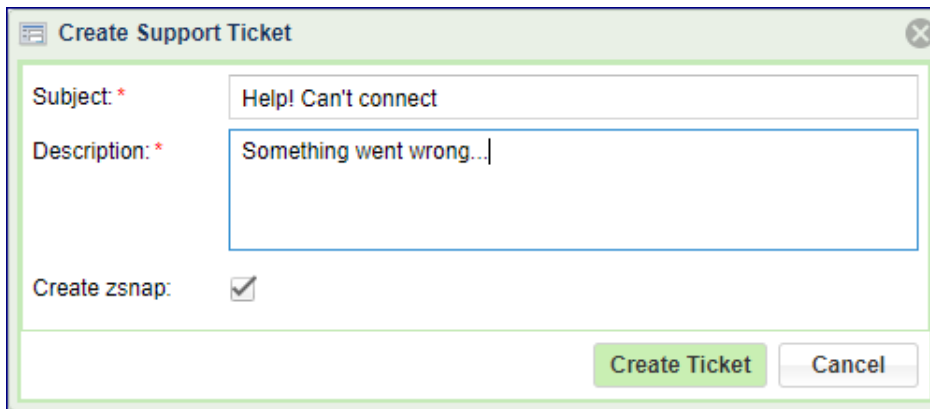
MANAGING TECH SUPPORT TICKETS

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.



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.