



Zadara Storage Cloud

VPSA User Guide

February 2014

Revision B

© 2011–2014 ZADARA Storage, Inc. All rights reserved.

Contents

1	Preface	5
1.1	Intended Audience	5
1.2	Document History	5
2	Introduction	5
2.1	VPSA Components	5
2.1.1	VPSA Management Console	5
2.1.2	Virtual Controller	5
2.1.3	Dedicated Drives	6
2.1.4	Zadara Storage Cloud Orchestrator	6
3	Getting Started	6
3.1	Registering a Zadara Account & Creating a VPSA	6
3.2	The VPSA GUI	9
3.3	Creating RAID Groups & Pool & Volumes	10
4	Managing VPSAs	10
4.1	Adding and removing Disk Drives	10
4.2	Managing Zadara Engines	11
4.3	Assigning Public IPs	12
4.4	Adjusting Cache	13
4.5	Managing Enterprise Suite	14
4.6	Hibernating you VPSA	15
5	Managing RAID Groups and Disks	15
5.1	Creating a RAID Group	15
5.1.1	Understanding RAID levels	16
5.2	Viewing RAID Group properties	17
5.3	Understanding Hot Spare Drives	19
5.4	Managing Sync Speed	20
5.5	Understanding Media Scans	20
5.6	Force Recovery	21
5.7	Replacing a Drive	21
5.8	Shredding a Drive	22
5.9	Viewing Drive properties	23
6	Configuring Storage Pools	24
6.1	Understanding Storage Pools	24

6.2	Creating and Managing Pools	24
6.2.1	Creating a Pool.....	24
6.2.2	Expanding Pool Capacity.....	26
6.2.3	Enabling Cache.....	26
6.3	Viewing Pool properties.....	26
6.4	Managing Pool Alerts	29
7	Managing Servers.....	30
7.1	Adding a Server	30
7.1.1	Adding a Server automatically.....	30
7.1.2	Adding a Network Range.....	33
7.1.3	Adding a Server manually.....	33
7.1.4	Configure Server Attributes	37
7.2	Viewing Servers Properties.....	38
8	Managing Volumes, Snapshots and Clones	40
8.1	Creating and Deleting a Volume.....	40
8.2	Attaching & detaching Volumes to Servers	43
8.3	Expanding a Volume.....	45
8.4	Managing Snapshots and Snapshot Policies	46
8.4.1	Manual creation & deletion of Snapshots	46
8.4.2	Managing Snapshot Policies.....	47
8.5	Cloning a Volume.....	49
8.6	Managing Encrypted Volumes	51
8.7	Viewing Volume Properties	52
8.8	Filtering Snapshots	55
9	Managing Access Control	56
9.1	Adding & Deleting Users.....	56
9.2	Managing User Passwords.....	57
9.3	Managing NAS Users Access Control.....	59
9.3.1	Creating NFS Users.....	59
9.3.2	Creating SMB Users.....	60
9.3.3	Editing SMB Users Password.....	61
9.4	Creating NFS Groups	61
9.5	Enabling Active Directory Authentication.....	62
9.5.1	Joining the VPSA to Active Directory	62
9.5.2	Leaving an Active Directory	63
10	Managing Remote Mirroring.....	64

10.1	Connect to a remote VPSA	64
10.2	Viewing remote VPSA Properties	65
10.3	Creating a Remote Mirror	66
10.4	Pause & Continue Remote Mirror	67
10.5	Managing Mirror Lifecycle.....	67
10.5.1	Clone Destination Volume for Dev & Test of Remote Mirror	68
10.5.2	Breaking a Mirror	68
10.5.3	Reconnecting a Mirror.....	69
10.6	Viewing Remote Mirror Properties	70
11	Managing TechSupport Tickets	73

1 Preface

This documentation presents information specific to Zadara Storage products. The information is for reference purposes and is subject to changes without prior notice.

1.1 Intended Audience

This document is intended for end users and storage administrators who wish to consume Enterprise Storage-as-a-Service via the Zadara Storage VPSA service.

1.2 Document History

Date	Revision	Description
November 2013	A	Initial revision for V2.0 Release Updated to version 13.10 content
February 2014	B	Updates for Release 14.01

2 Introduction

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.

2.1 VPSA Components

2.1.1 VPSA Management Console

The VPSA Management Console is the your gateway to the Zadara Storage ecosystem through which you can create, view, and modify your VPSA configurations on the many different Clouds that Zadara Storage offers.

2.1.2 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 clustering mode 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.

2.1.3 Dedicated Drives

The Zadara Storage Cloud Orchestrator assigns dedicated drives per VPSA. The drives are provisioned from different Storage Nodes (SNs) 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. This provides a complete isolation and privacy of your data drives. No bad neighbors can access your drives, impact your performance, or compromise your privacy and security. Hence, your drives' QoS are guaranteed.

2.1.4 Zadara Storage Cloud Orchestrator

This is the orchestrator software running inside the Zadara Storage Cloud, managing physical resources, provisioning VPSAs, and monitoring the Cloud. The VPSA Management Console communicates with the Zadara Storage Cloud Orchestrator to create, modify and update the user VPSAs.

Note: management of the VPSA internal objects (such as Raid Groups, Pools and Volumes) and the VPSA data is completely separated from the Zadara Cloud management.

3 Getting Started

3.1 Registering a Zadara Account & Creating a VPSA

- Go to <https://manage.zadarastorage.com/register/> and complete the form to register a Zadara Account.
- Go to your **VPSA Management Console** at <https://manage.zadarastorage.com>, Login using the registered username\email & password, and press "**Create VPSA**." The following dialog will appear:

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.
- **VPSA Description** – Give the VPSA a description.
- **Select Cloud Provider** – From your VPSA Management Console you can provision and manage VPSAs across many 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 an efficient iSCSI or NFS\CIFS connectivity.
- **Select the Zadara Engine** – The Zadara Engine defines the compute characteristics of your VPSA's Virtual Controllers (VCs). Each engine type defines the following characteristics:
 - Number of CPUs that are assigned to your VPSA's VCs.
 - Amount of RAM that is assigned to your VPSA's VCs.
 - Default size of protected SSD Cache.

The Zadara Engine type can be changed at any time throughout the lifetime of your VPSA depending on your applications workload needs.

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.

- **Drive Quantities** – Select the type and number of drives which are to be allocated for your VPSA.
 - The Zadara Cloud Orchestrator allocates full and 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 maximum of number of drives per Zadara Engine type. The stronger the Engine is, the more drives you can add. The following is the max drives per Engine type:

Engine Type	Maximum # of drives
Baby	5
Basic	10
Boost	20
Blazing	40

- **Enterprise Suite** – Enable this checkbox if you'd like to leverage the VPSA Enterprise Suite feature set, including:

- Snapshot
- Cloning
- Remote Mirroring
- Volume Encryption

Note: Enterprise Suite can be enabled\disabled at any time throughout the lifetime of your VPSA.

- Once you have completed selecting the above VPSA characteristics, press the **Submit** button to confirm the VPSA creation request. The requested VPSA will appear in the “Awaiting Approval” list.

1 VPSA(s) Awaiting Approval (click to hide)

Name	Status	Provider	Drives	Zadara Engine
Zadara_Demo	Awaiting Approval	AWS US East (N. Virginia) 1	3	Basic

- VPSA creation requires the approval of the 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:

VPSA Management Console

Create VPSA

My VPSAs

Name	Management Address	Drives	Provider	Status	Refresh	Delete
ZadaraProd		4	AWS US East (N. Virginia) 2	Launching...		

- 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:

VPSA Management Console

Create VPSA

My VPSAs

Name	Management Address	Drives	Provider	Status	Refresh	Delete
ZadaraProd	https://vsa-00000112-aws2.zadaraavpsa.com	4	AWS US East (N. Virginia) 2	Ready		

ZadaraProd (AWS US East (N. Virginia) 2)

Name: ZadaraProd

Description: N/A

Status: Ready

Zadara Engine: Basic

Time Created (GMT):

2013-11-16 21:29:50

Management Address:

https://vsa-00000112-aws2.zadaraavpsa.com

IP Address: 172.31.240.108

Public IP: None

Cache: 20GB (20GB from Engine)

Enterprise Suite: Enabled

Change Engine

Assign Public IP

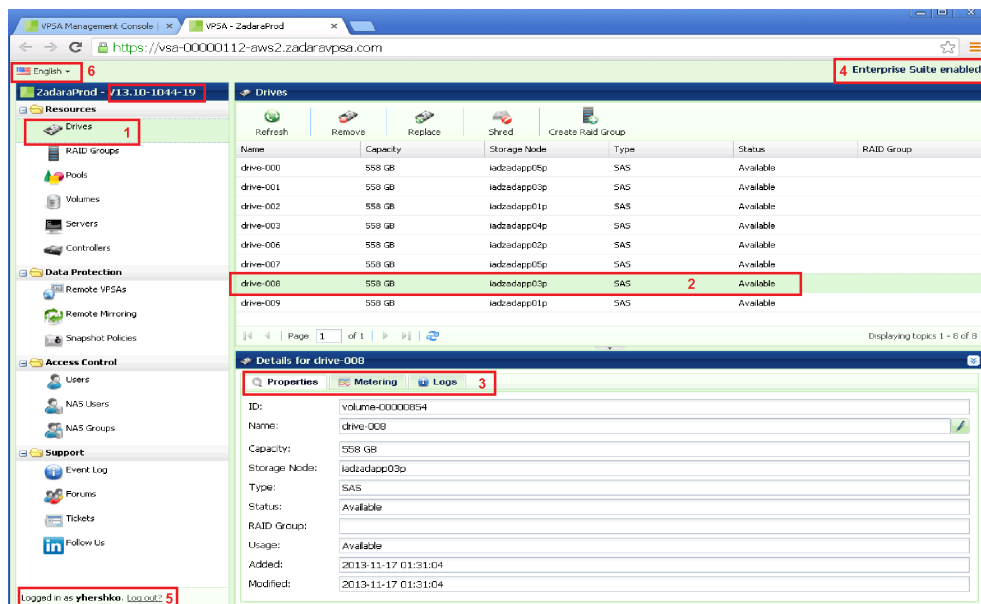
Add Drives

Adjust Cache

Disable Enterprise Suite

- 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.

3.2 The VPSA GUI



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

- Main navigation left panel** – Traverse through the various VPSA entities. The selected entity is highlighted.
- 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.
- The south pane** – Displays detailed information regarding the selected object. All objects have at least 3 tabs:
 - Properties** – Detailed properties of the object.
 - Metering** – Typically IO workload metering info.
 - Logs** – List of event-log messages related to that object.
- Enterprise Suite enable/disable** – Indicated at the top-right corner. You can enable/disable Enterprise Suite any time via the VPSA Management Console. You need to refresh the browser (f5) to see the updated status.

5. **Logged-in username** – Displayed at the bottom left corner.
6. **Selected Language** – Displayed at the top left corner .

3.3 Creating RAID Groups & Pool & Volumes

- Create a Raid Group. For more details check [here](#).
- Create a Pool. For more details check [here](#).
- Create an iSCSI\NFS\SMB Thin Provisioned Volume. For more details check [here](#).
- Attach the Volume to a Server. For more details check [here](#).

Congratulations! You have a new VPSA provisioned and ready.
The following sections describe in detail the various capabilities and services of your VPSA.

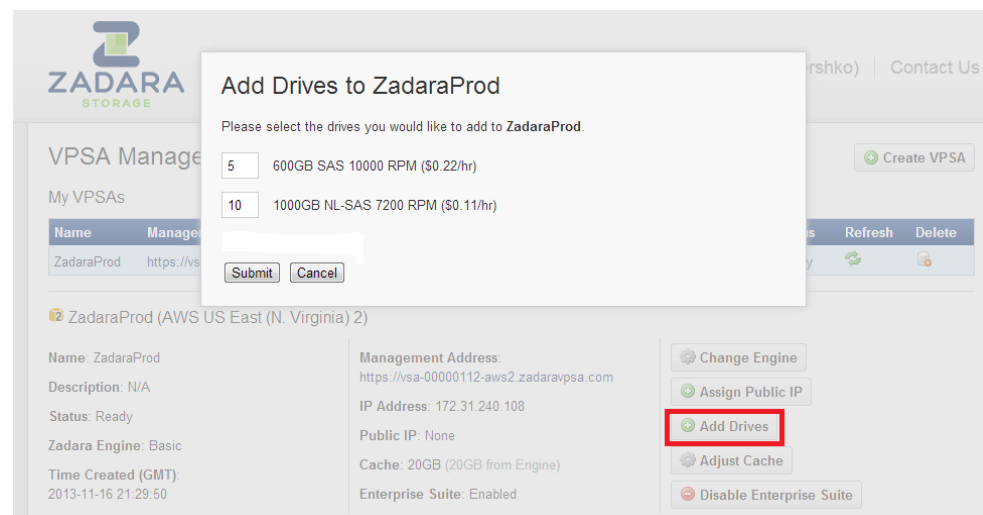
4 Managing VPSAs

You create, delete, and manage the resources composing your VPSAs via **the VPSA Management Console**.

This section describes the available operations at the Management Console (<https://manage.zadarastorage.com>).

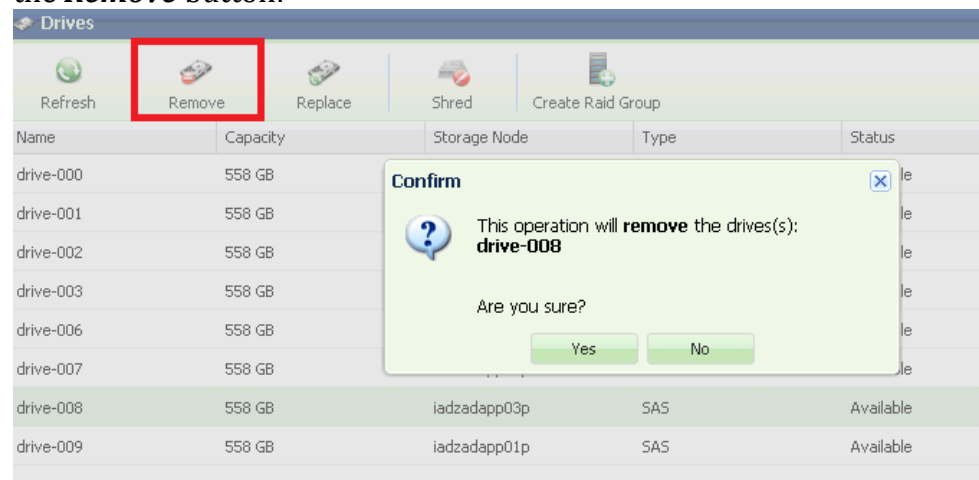
4.1 Adding and removing Disk Drives

To add Drives to your VPSA, go to the VPSA Management Console, select the VPSA, and then press the **Add drives** button.



- Select the number of drives from each available drive type you wish to add to your VPSA, and press **Submit**.
- This operation requires the approval of the Zadara Storage Cloud Admin. Once approved, you'll see the number of drives in the VPSA Management Console update accordingly, and if you refresh the Drives page in the VPSA GUI, the new drives will be displayed.

You remove unused drives (in status "Available") from within the VPSA. Go to the VPSA GUI -> Drive page, select the drive you wish to remove, and press the **Remove** button:



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

4.2 Managing Zadara Engines

The Zadara Engines define the compute resources that are dedicated to your VPSA. Those resources are not shared with any other VPSA or tenant within the Zadara Storage Cloud. The following Zadara Engines are defined:

Zadara Engine Type	Dedicated compute resources
Baby	1 CPU, 4 GB RAM
Basic	2 CPU, 8 GB RAM
Boost	4 CPU, 16 GB RAM
Blazing	8 CPU, 32 GB RAM


Each VPSA is provisioned with an SSD Cache partition to be used for both metadata, and user application data read and write caching. The SSD cache partition is protected under RAID-1, where each mirror copy resides on a different SN, thus ensuring Cache resilience to SN failure. Each Engine type comes with a default SSD Cache partition size. You can request additional SSD capacity for caching. For more details see ["Adjusting Cache."](#) The default SSD cache size is as follows:

Zadara Engine Type	SSD Cache Size
Baby	20GB
Basic	20GB
Boost	40GB
Blazing	80GB

The Zadara Engine type also defines the maximum number of drives that can be provisioned to a VPSA:

Engine Type	Maximum # of drives
Baby	5
Basic	10
Boost	20
Blazing	40



Note: Downgrading a VPSA engine type is prohibited if it exceeds the maximum number of drives limitation of the target Engine Type.

To change the Zadara Engine type, press the  **Change Engine** button in the VPSA Management Console:



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."

Name	Management Address	Drives	Provider	Status	Refresh	Delete
ZadaraProd	https://vsa-00000112-aws2.zadaraavpsa.com	9	 AWS US East (N. Virginia) 2	Upgrade Pending...		

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

4.3 Assigning Public IPs

For security & privacy reasons, you cannot access the VPSA by default from the public Internet. The VPSA FrontEnd IP address, used for VPSA management (via GUI and REST API) and for data IO workload (via iSCSI\NFS\SMB protocols), is allocated on the Zadara Storage Cloud "Front-End" network 10GbE interface which is routable from the Cloud Servers network. Servers outside of your Cloud Servers network cannot reach this IP address.

You are required to assign a Public IP address to your VPSA when you'd like to access it from the public Internet. A typical use case requiring Public IP addresses is when you're doing Asynchronous Remote Mirroring between two VPSAs in different regions, or even 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.

Assign Public IP

To assign a Public IP address to your VPSA, go to the VPSA Management Console and press the **Assign Public IP** button. You can see the assigned IP address in your VPSA details in the VPSA Management Console and in the VPSA GUI, under Controller->Public IPs.

Remove Public IP

To remove it, simply click the **Remove Public IP** button in the VPSA Management Console.

Notes:

- Access to the VPSA GUI is blocked through the Public IP for security reasons.
- NATed Servers are not supported for iSCSI, NFS, and SMB protocols over the Public IP.

4.4 Adjusting Cache

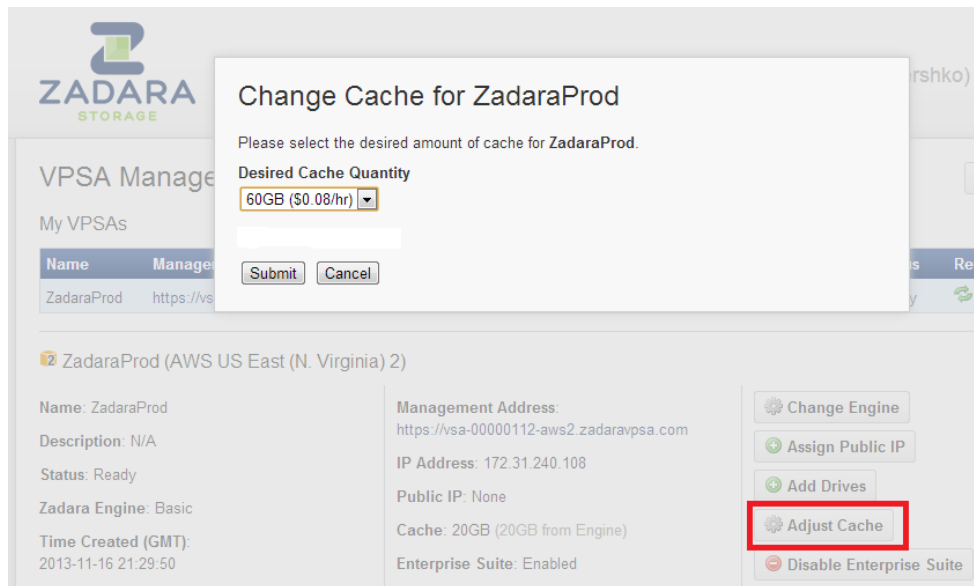
Each VPSA is provisioned with an SSD Cache partition, which is utilized by the VPSA for both metadata, and user application data read and write caching. The SSD cache partition is protected under RAID-1, where each mirror copy resides on a different SN thus ensuring Cache resilience to SN failure.

The VPSA SSD Cache size is elastic, meaning that you can increase or decrease the SSD Cache size based on your Application Servers workload. The initially assigned default SSD cache size is also the minimal cache size for a given Zadara Engine.

The default and maximum SSD Cache size depends on the Engine type as follows:

Zadara Engine	Default SSD Cache Size	Max SSD Cache Size
Baby	20GB	80GB
Basic	20GB	80GB
Boost	40GB	160GB
Blazing	80GB	320GB

To change the SSD Cache size for your VPSA, go to the VPSA Management Console and press the **Adjust Cache** button:

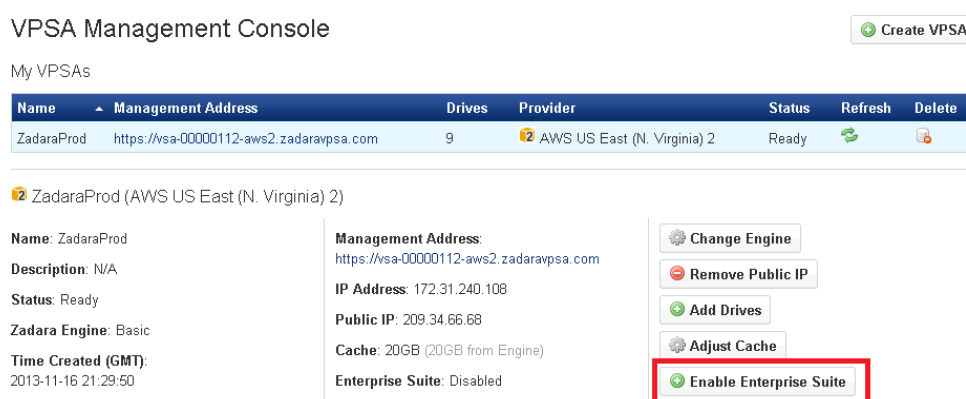


4.5 Managing Enterprise Suite

The VPSA Enterprise Suite extends your VPSA functionality to provide enterprise-level security and data management capabilities, including:

- **Data-at-rest Volume encryption** – Using user-controlled encryption keys.
- **Business continuity** – Zero-capacity, space-efficient, instantly available, unlimited number of read-only Snapshots.
- **Test & Dev** – Zero-capacity, space-efficient, instantly available read-write Clones of your Volume at any point-in-time snapshot.
- **Snapshot-based Asynchronous Remote Mirroring** – Crosses regions and crosses providers for Disaster Recovery.

To enable or disable Enterprise Suite, go to the VPSA Management Console and press the **Enable\Disable Enterprise Suite** button:



Note: Disabling the Enterprise Suite will have the following impact on existing objects:

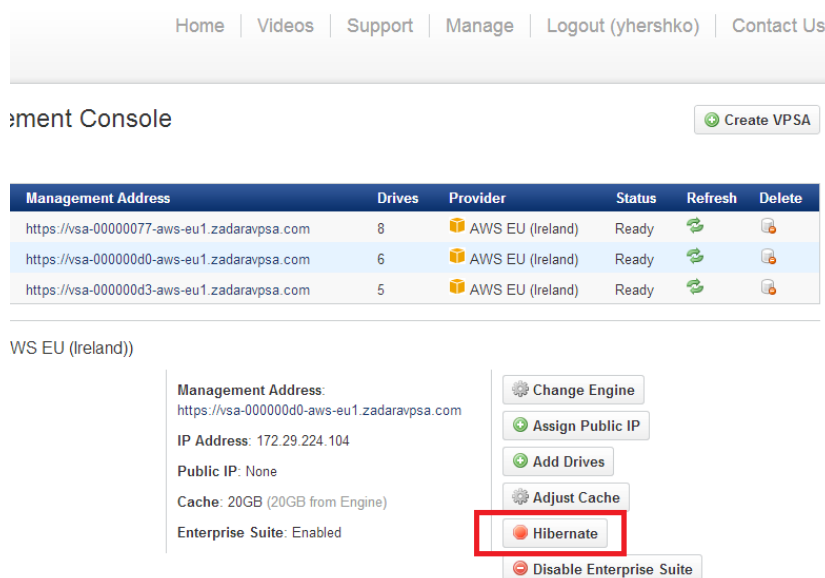
- Snapshot and Clones are not affected.
- Snapshot Policies are paused.
- Remote Mirrors are paused.

- Encrypted Volumes are not affected.
- New Encrypted Volumes cannot be created.

4.6 Hibernating you 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. Hibernating a VPSA involves the process of deleting its Virtual Controllers (the Zadara Engine) while sustaining its 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 NFA\SMB volumes. While the VPSA is in Hibernate state you will not be billed for the Zadara Engine. Resuming a Hibernated VPSA may take a few minutes.

To hibernate a VPSA, go to the VPSA Management Console and press the ***Hibernate*** button:



To resume access to the VPSA, go to the VPSA Management Console and press

the ***Restore*** button.

5 Managing RAID Groups and Disks

5.1 Creating a RAID Group



VPSA RAID Groups define the level of protection and performance of your data according to the selected RAID-level, number of drives, and type of drives. RAID

groups span across full-capacity drives from different Storage Nodes. Thus, a RAID Group is resilient to a single drive failure (RAID-6 allows for a 2 drive failure), as well as to a complete Storage Node failure.

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).
- Select **Protection Type**. See the below description for the various RAID levels.
- Select whether to allocate a drive as a Hot Spare for this RAID group. See more details about managing **Hot Spares** [here](#).
- Select the drives that participate in the RAID Group. As noted below, for RAID-1 a minimum of 2 drives is required, for RAID-5 a minimum of 3 drives, and for RAID-6 a minimum of 4 drives.
- For maximum redundancy, drives **MUST** be selected from different Storage Nodes (VPSA will prevent you from doing otherwise).
- It is possible but not recommended to mix drives of different types in a RAID Group.

Create Raid Group

Name: RAID-5-DB

Protection: ☐ RAID1 ☒ RAID5 ☐ RAID6

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

Hot Spare: ☐

Drives:

<input checked="" type="checkbox"/>	Name	Type	Capacity	Status	Storage Node
<input checked="" type="checkbox"/>	drive-020	SAS	558 GB	Available	iadzadapp05p
<input checked="" type="checkbox"/>	drive-021	SAS	558 GB	Available	iadzadapp03p
<input checked="" type="checkbox"/>	drive-022	SAS	558 GB	Available	iadzadapp01p
<input checked="" type="checkbox"/>	drive-023	SAS	558 GB	Available	iadzadapp04p
<input checked="" type="checkbox"/>	drive-024	SAS	558 GB	Available	iadzadapp02p

Create Cancel

5.1.1 Understanding RAID levels

RAID level	Description
RAID-1 – Mirroring	Offers a good combination of data protection and performance. RAID-1, or Drive Mirroring, creates fault tolerance by storing duplicate sets of data on a minimum of two hard drives. There must be 2 or 3 drives in a RAID-1. RAID-1 and RAID-10 are the most costly fault tolerance methods because they require 50 percent of the drive capacity to store the redundant data. RAID-1 mirrors the contents of one hard drive in the array onto another.

	If either hard drive fails, the other hard drive provides a backup copy of the files, and normal system operations are not interrupted.
RAID-5	Offers the best combination of data protection and usable capacity, while also improving performance over RAID-6. RAID-5 stores parity data across all the physical drives in the array and allows more simultaneous read operations and higher performance. If a drive fails, the controller uses the parity data and the data on the remaining drives to reconstruct data from the failed drive. The system continues operating with a slightly reduced performance until you replace the failed drive. RAID 5 can only withstand the loss of one drive without total array failure. It requires an array with a minimum of three physical drives. Usable capacity is $N-1$, where N is the number of physical drives in the logical array.
RAID-6	Offers the best data protection and is an extension of RAID-5. RAID-6 uses multiple parity sets to store data and can therefore tolerate up to 2 drive failures simultaneously. RAID-6 requires a minimum of 4 drives. Performance is slightly lower than RAID-5 due to parity data updating on multiple drives. Usable capacity is $N-2$ where N is the number of physical drives in the logical array.
RAID-10– Mirroring and Striping	Offers the best combination of data protection and performance. RAID-10, or drive mirroring, creates fault tolerance by storing duplicate sets of data on a minimum of four hard drives. RAID-10 is the most costly fault tolerance method because it requires 50 percent of the drive capacity to store the redundant data. 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 provides a backup copy of the files, 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. Note: RAID-10 is achieved in VPSA by creating RAID-1 RAID Groups, and striping them together at the Pool level.

5.2 Viewing RAID Group properties

The following properties and metering information are displayed in the RAID-Group details, found 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.
Protection	Selected RAID level—RAID-1, RAID-5 or RAID-6.
Capacity	Total protected and usable capacity of the RAID Group.
Available Capacity	RAID Group usable capacity which is not allocated to any Pool.
Stripe Size	Stripe size (per drive) for RAID-5 and RAID-6.
Mirror Number	Number of mirror copies for RAID-1.
Protection Width	Number of Drives participating in a RAID-5 and 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
- Location (Storage Node)
- Capacity (in GB)
- 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, 1 Minute, 10 Minutes, or 1 Hour. The **Auto** button lets you see continuously-updating live metering info (refreshed every 3 sec).

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 SCSI 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.

5.3 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 issuing “Add Spare” and “Remove Spare” on a selected RAID Group in the VPSA RAID Group GUI 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. (To prevent drive replacement and RAID rebuild process due to intermediate failures, like an SN reboot, the VPSA tries to recover the drive for 30 minutes before declaring it 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 not used in any other RAID Group within the VPSA. This process can be executed manually, or automated via the VPSA REST APIs. Simply identify the failed drive, issue the “Replace” command, and select the available drive to use for the replacement. For more details check [here](#).

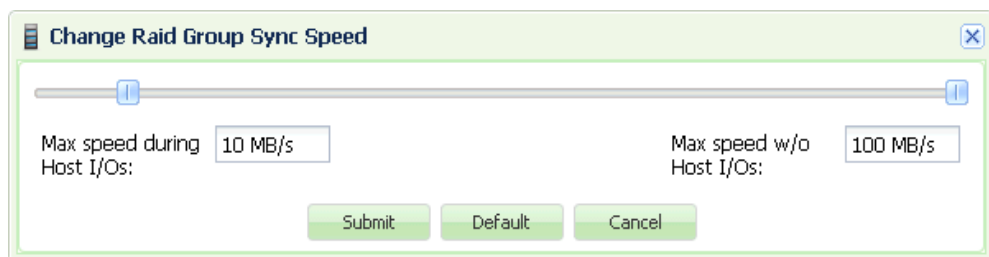
5.4 Managing Sync Speed



RAID Group Sync Speed controls the rate with which data is synchronized during a Rebuild process, either on a newly created RAID group or following a drive replacement.

Setting the Sync Speed is a tradeoff between the need to complete the RAID rebuild as fast as possible in order to return to full redundancy level, and the ability to supply good response time and throughput for Server IO workload. 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 IOs. Set it low if you want to prioritize the Servers IOs. Set it high if you’d like to prioritize the RAID rebuild process.
 - Default value: 10 MB/s
 - Range: 1 - 100 MB/s
- **“Max Speed W/O Host I/Os”** – Controls the sync speed when there are no Server IOs. Typically you would set it to max value (100 MB/s) unless it consumes too much of the VPSA resources (depending on the Engine type) impacting performance of other Raid Groups (which do have active Server IOs).



Sync Speed can be set and modified at any time and can vary between RAID groups. The Sync Speed also applies to Media Scan (see below).

5.5 Understanding Media Scans



Media Scan is the process of checking RAID-5 and 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 decide to trigger a media scan on a RAID Group, beyond this once-a-month scan, if there was an event that is suspected of risking data integrity, such as the failure of two or more drives in a RAID-5.

RAID status will change to “Repairing X%” during the media scan. At the end of the media scan, it saves the results of the scan in an event-log message.

Media scan cannot be aborted in the middle, but it can be paused. Press the ***Pause Media Scan*** button to pause the operation (the Media Scan button toggles to ***Pause Media Scan*** for RAID Groups which are in media scan).

RAID status will change to “Repairing Paused” when the media scan is paused.

5.6 Force Recovery

Force Recovery can be issued only on Failed RAID-5 and RAID-6 RAID Groups, after one or more of the failed drives have been recovered.

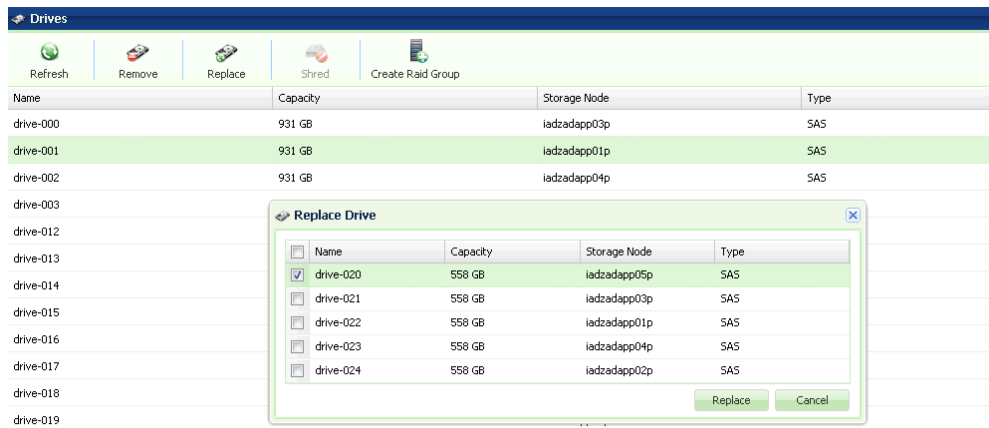
If all the drives were recovered, the VPSA will have enough information to determine how to recover the RAID automatically, but if one drive is permanently gone in a RAID-5 RAID Group or two drives are permanently gone in a RAID-6 RAID Group the VPSA is unable to determine if the available drives contain the most up-to-date data and hence cannot 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).

Note: 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.

5.7 Replacing a Drive

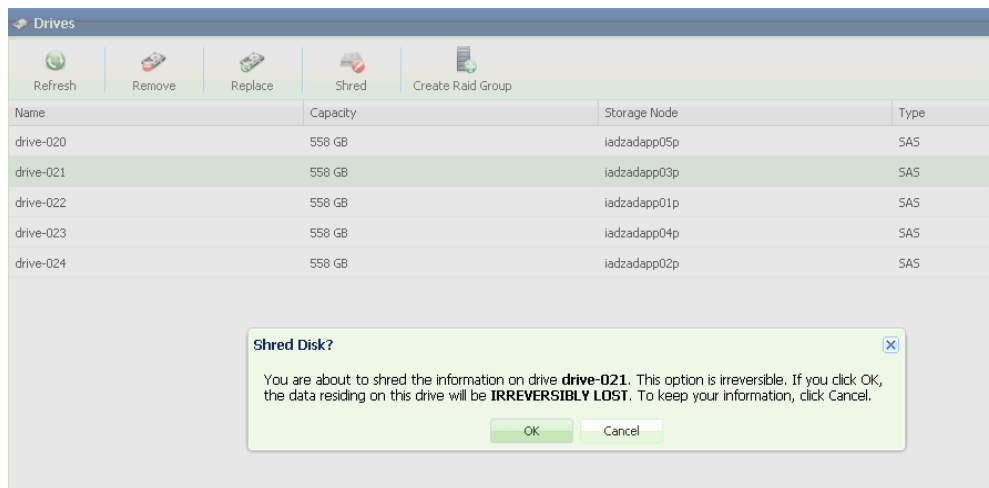


Press the **Replace** button to replace a drive. When selecting the replacement drive you must choose a drive that will not break the RAID Group redundancy (i.e., you cannot have two or more drives from the same Storage Node in a RAID Group). If you select a drive that has a different type, or larger 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 Resyncing state.

5.8 Shredding a Drive



Shredding is the process of erasing the data on a drive for security and privacy reasons by writing random data. Typically, you will shred a drive before returning it to the Zadara Cloud or before deleting your VPSSA.

Shredding is applicable only on drives in Available status (i.e., not in a RAID group).

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

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

Note: Shredding is irreversible!

5.9 Viewing Drive properties

You can view the following properties and metering information in the Drives Details south panel tabs:

Properties

Each drive includes the following properties:

Property	Description
Name	An internally assigned unique ID.
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 inside 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 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 containing this drive.
Usage	In-use.
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 one of the following: 1 Second, 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.

6 Configuring Storage Pools

6.1 Understanding Storage Pools

Storage Pools are virtual entities managing storage provisioning from 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. They share the same data chunks as the Volume. Anytime you actually modify the data in the Volume, or in one of the Clones, the data chunk is being copied-on-write (COW) in order to apply the new written data without affecting the data set of the other objects sharing the same data chunk.

Pools attributes define the way Volumes, Snapshots and Clones are provisioned.

6.2 Creating and Managing Pools

6.2.1 Creating a Pool

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. You will see the following dialog appear:

Create Pool

Name:

Raid Group(s):

<input checked="" type="checkbox"/>	Name	Protection	Status	Available
<input checked="" type="checkbox"/>	R11	RAID1	Normal	556 GB
<input checked="" type="checkbox"/>	R12	RAID1	Normal	556 GB

Capacity (GB):

Type: ☐ Transactional Workloads ☒ Repository Storage

Cached: ☒

Striped: ☒

Select the Pool attributes:

- **Display Name** – Can be modified anytime later.
- **Raid Group(s) selection** – Select one or more RAID Groups from which protected storage capacity will be allocated for this Pool.
- **Capacity** – The Pool's physical capacity 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 smaller capacity than is available in the selected RAID groups, the capacity will be evenly distributed between the RAID Group.
- **Type** – The VPSA supports two types of Pools: Transactional & Repository Pools. The difference is the chunk size used for the mapping of virtual LBAs to Physical Drive addresses. The following table describes the tradeoff of each type and the recommended use cases:

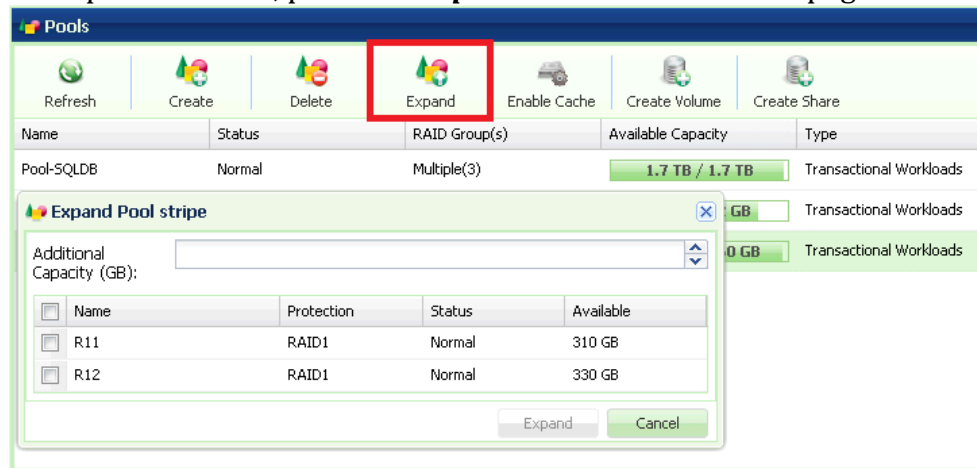
	Transactional Pool	Repository Pool
Chunk size	256KB	1MB
Pros	<ul style="list-style-type: none"> • Faster COW operation • Space efficiency on Random writes to Snapshots 	<ul style="list-style-type: none"> • Smaller metadata size • Sequential workload performance is similar to transactional Pools
Cons	Increased metadata size	<ul style="list-style-type: none"> • Slower COW operation. • Less space efficient
Use Case	Transactional Workload with Snapshots	<ul style="list-style-type: none"> • Repository type workload. • Large Pools • Many snapshots to keep

- **Cached** – Use SSD to Cache Servers Reads and Writes.
 - All Pools that are marked as "Cached" share the VPSA Cache.

- Write Cache is mostly efficient on random writes intensive workloads.
- 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, RAID-5 and RAID-6 creates RAID-10, RAID-50 and RAID-60 type of configuration. Use striping to improve performance of random workloads since the IOs will be distributed and all drives will share the workload.

6.2.2 Expanding Pool Capacity

To Expand the Pool, press the **Expand** button on the Pools page.



You can expand a Pool using capacity from any RAID Group. 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, which may impact the pool performance and protection QoS.

6.2.3 Enabling Cache



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

One use case to leverage this capability is to enable Caching only after the initial copy of the data into the VPSA. The initial copy typically generates 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 more random type of IO workload.

6.3 Viewing Pool properties

The Pools details are shown in the following south panel tabs:

Properties

Each Pool includes the following properties:

Property	Description
ID	An internally assigned unique ID.
Name	User assigned name. Can be modified anytime.
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.
Mode	<ul style="list-style-type: none">• Simple – There are one or more RAID Groups concatenated.• Striped – There are two or more RAID Groups which are striped.• Mixed – There are two or more RAID Groups which are concatenated <i>and</i> striped.
Free Capacity State	<ul style="list-style-type: none">• Normal• Alert• Protected
Type	<ul style="list-style-type: none">• Transactional Workloads• Repository Storage
Status	<ul style="list-style-type: none">• Normal• Creating• Deleting• Partial\Failed – At least one of the underlying RAID groups is failed, or the Pool metadata cannot be initialized at Start Of the Day.
Stripe Size	Applicable only for Pools of Mode Striped (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
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.
----------	--

Alerts

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

- **Alert Mode Threshold** - “Alert me when it is estimated that the Pool will be at full capacity in X Minutes.”
 - Default Value: 360 minutes
- **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 capacity in X Minutes.”
 - Default Value: 60 minutes
- **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
- **Emergency Mode Threshold** - “Delete snapshots, starting from the oldest, when there is less than the following capacity left in the Pool”
 - Default Value: 1 GB

Raid Groups

This tab lists the RAID Groups participating in the selected Pool. Each RAID Group includes the following information:

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

Volumes and Dest Volumes

These two tabs provide two lists for the provisioned Volumes and the Provisioned Remote Mirroring Destination Volumes. Please note that the Dest volumes are not displayed in the main Volume page since most operations are not applicable on them. Displaying the list of the Dest Volumes in the Pools south panel provides a complete picture of the Objects consuming capacity from the Pool. Each Volume includes the following information:

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

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 one of the following: 1 Second, 1 Minute, 10 Minutes, or 1 Hour. The **Auto** button lets you see continuously-updating live metering info (refreshed every 3 sec).

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 .

Logs

Displays all event logs associated with this Pool.

6.4 Managing Pool 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, and hence a Pool Protection Mechanism is required to alert and protect when free Pool space is low.

The VPSA Pool Protection Mechanism is mostly time 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 available capacity (which is 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 Pool 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 which are performed as part of the Pool Protection Mechanism:

- **Alert Threshold** – The estimated time (in minutes) before running out of free space in which an alert support ticket is to be submitted and an email to be 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 + email will be generated when the estimated time left is half of this threshold.
 - Default: 360 minutes (6 hours)
 - Minimum: 1 minute

- **Protection Threshold** – The estimated time (in minutes) before running out of free space in which the VPSA starts blocking the creation of new Volumes, Snapshots and Clones on that Pool. A support ticket & email are generated as well. When crossing this threshold, the Free Capacity State changes to “Protect” and the available capacity will be shown in Red.
 - Default: 60 minutes (1 hour)
 - Minimum: 1 minute
- **Capacity consumption rate calculation window size** - 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 it is impacted by intermediate changes in capacity allocations which can result from changes in workload characteristics and/or creation/deletion of new Snapshots and Clones.
 - Default: 60 minutes (1 hours)
 - Minimum: 1 minute
- **Emergency Threshold** – When the Pool’s free capacity 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 exist the Emergency threshold (i.e, when free capacity is greater than the threshold). A support ticket & email are generated as well. When crossing this threshold, the Free Capacity State changes to “Emergency” and the available capacity will be shown in Red.
 - Default: 1 GB
 - Minimum: 1 minute

7 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, NFS or SMB\CIFS protocols.

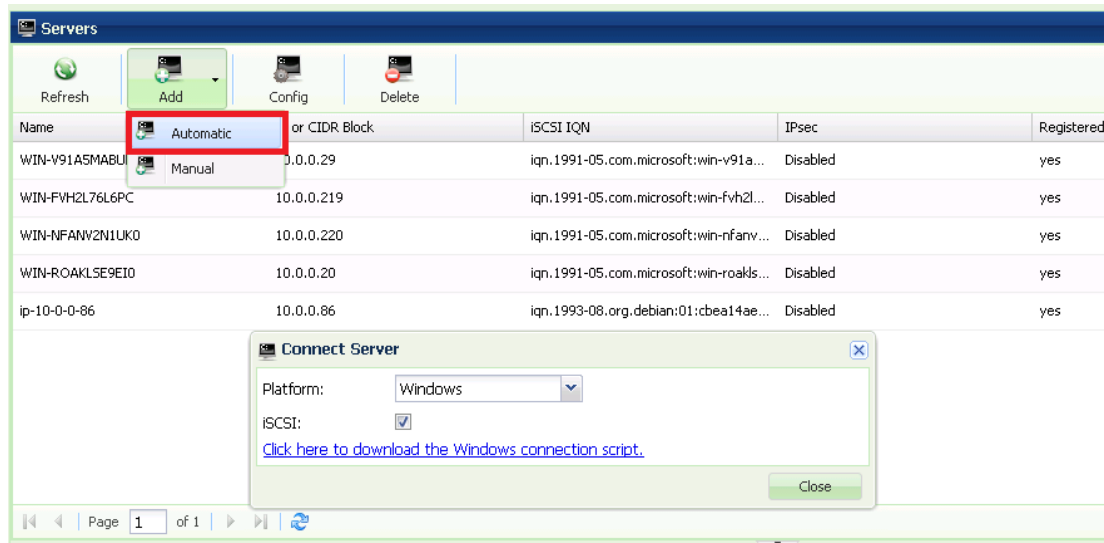
7.1 Adding a Server

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

- Creation of a Server Object in the VPSA database.
- Setting the Server IQN for iSCSI connectivity and/or the server IP address for NFS\SMB connectivity.
- Establishing CHAP authentication handshake between the Server and the VPSA for iSCSI.
- Register Server OS information (optional).

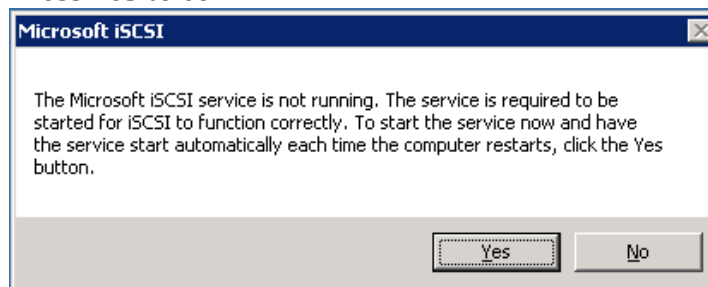
7.1.1 Adding a Server automatically

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

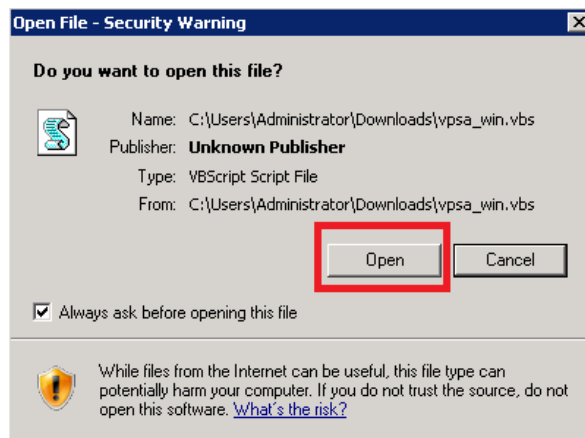
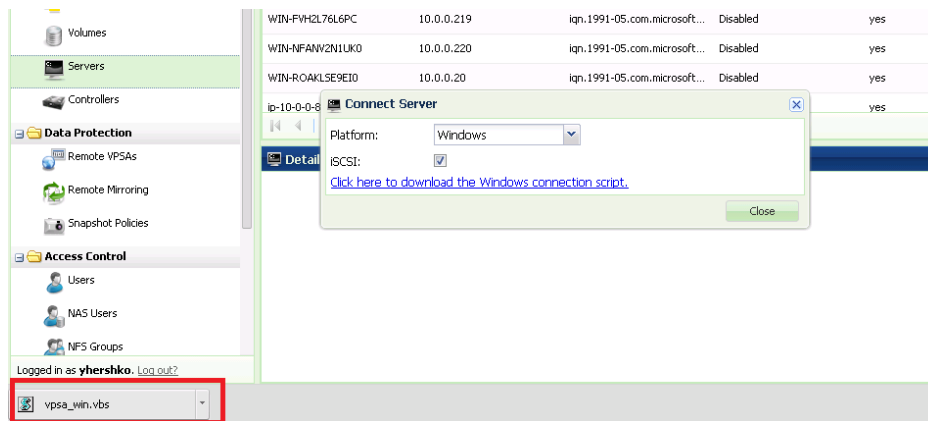


To Add a Windows Server:

- 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:



- 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.
- Click the download link. This will 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 Chrome browser.



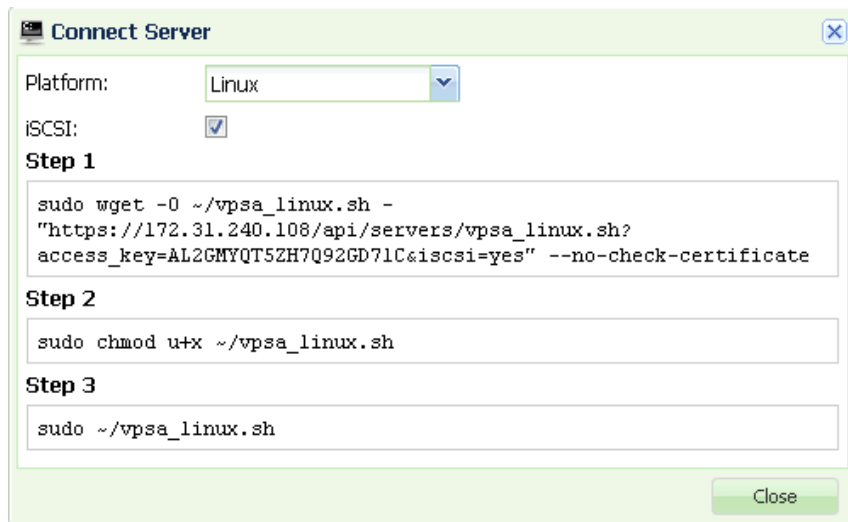
- Once the connect script successfully completes, the new connected Server will be listed in the VPSA Servers page with status = “Active,” Registered = “Yes,” and the correct OS details.

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.



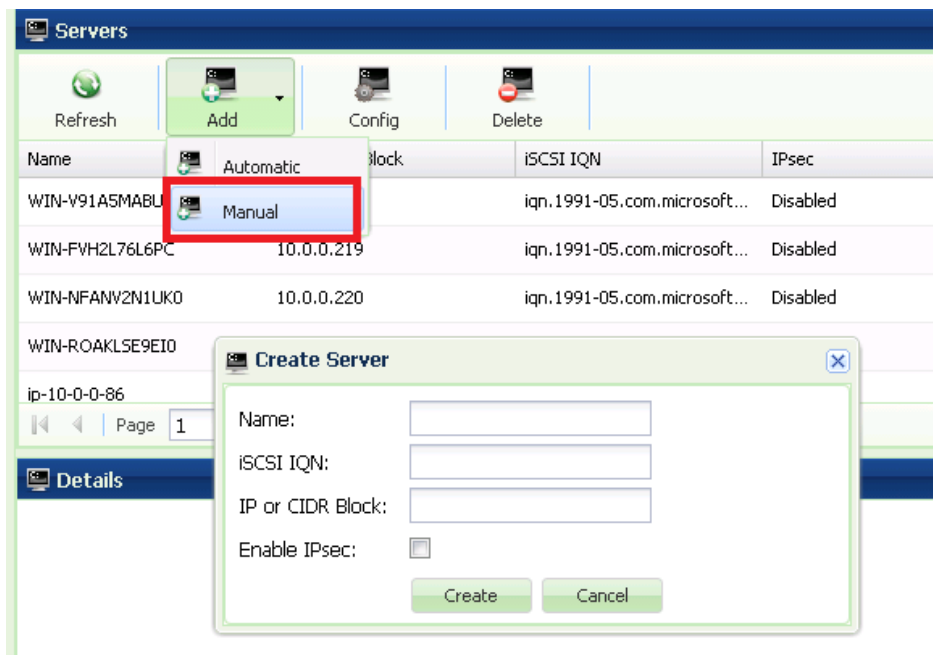
- Once the connect script completes successfully, the new connected Server will be listed in the VPSA Servers page with status = "Active," Registered = "Yes," and the correct OS details.

7.1.2 Adding a Network Range

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 following manual procedure to add this type of Server while specifying the IP range in CIDR notation (e.g. 192.168.1.1/24).

7.1.3 Adding a Server manually

If for some reason adding a server automatically doesn't work, or if you wish to add an IP Network Range, follow these steps to add the server manually. Go to Servers->Add and select **Manual**:



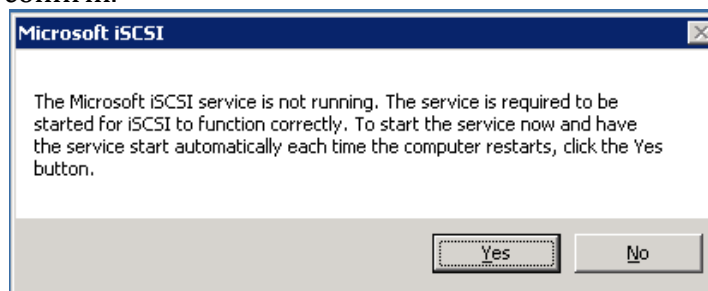
- Enter the Server Name.
- A Server Object must have at least one of the following attributes defined:
 - Server IP or CIDR – For NFS\SMB access.
 - IQN – For iSCSI access.
- Check the “Enable IPsec” checkbox if you wish to secure iSCSI traffic between the Server and the VPSA. Please note that your Server must be properly configured to utilize IPsec, and that performance is impacted.

7.1.3.1 Establishing an iSCSI connection

After adding a Server manually, 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 of Volumes.

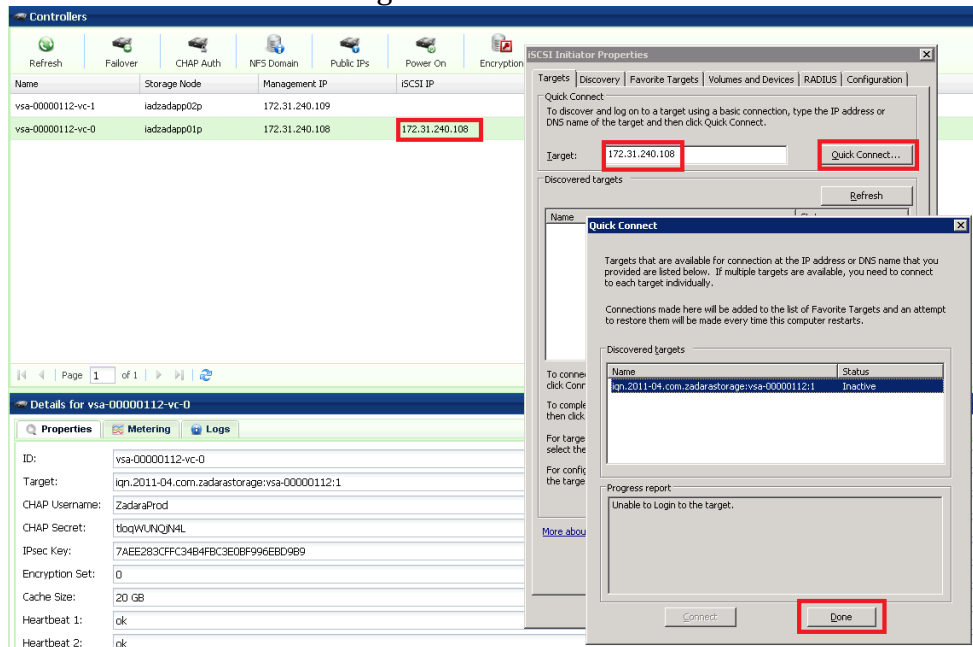
On Windows Servers:

- 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.

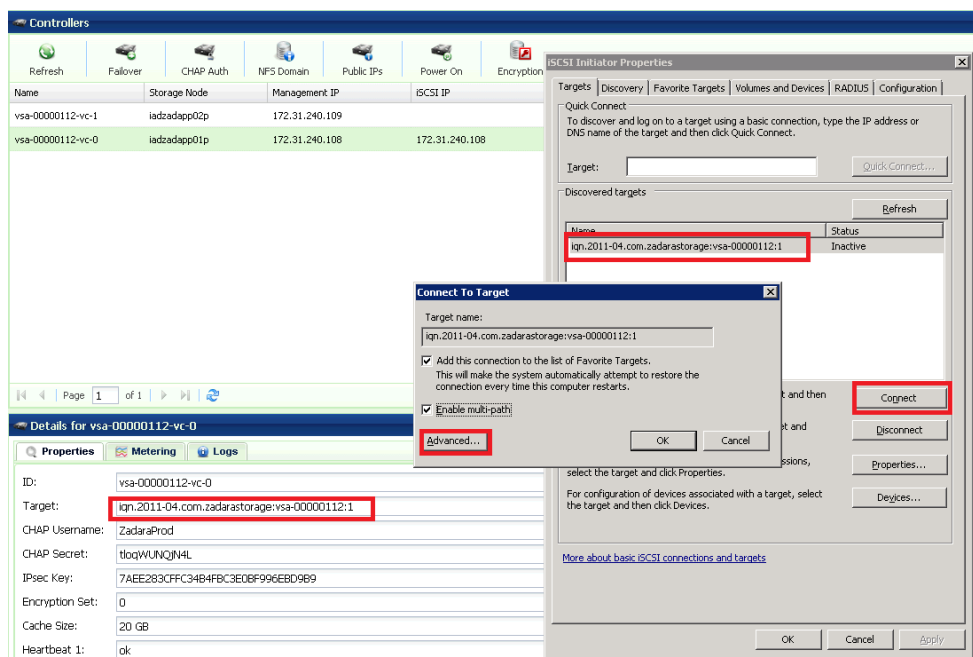


- The **Microsoft iSCSI Initiator Properties** dialog box will open, and the **Targets** tab will be displayed.

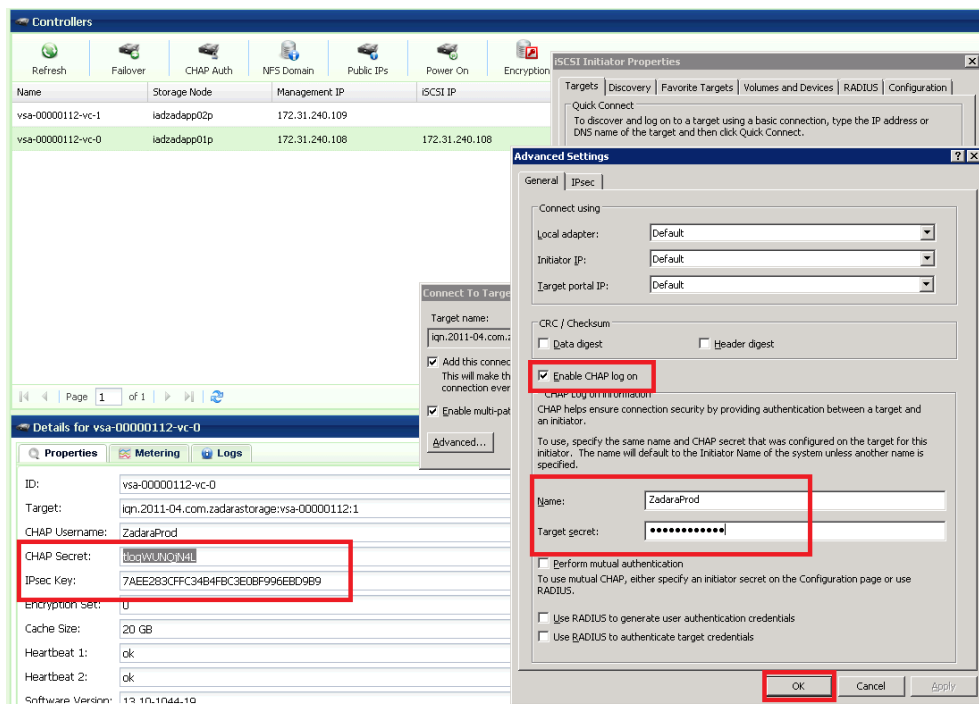
- On the **Targets** tab, type the iSCSI IP address of the VPSA (which is displayed in the VPSA GUI Controllers page) in the **Quick Connect** target text box, and then click the **Quick Connect...** button.
- The **Quick Connect** dialog box will be displayed, with the VPSA discovered iSCSI target in 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.

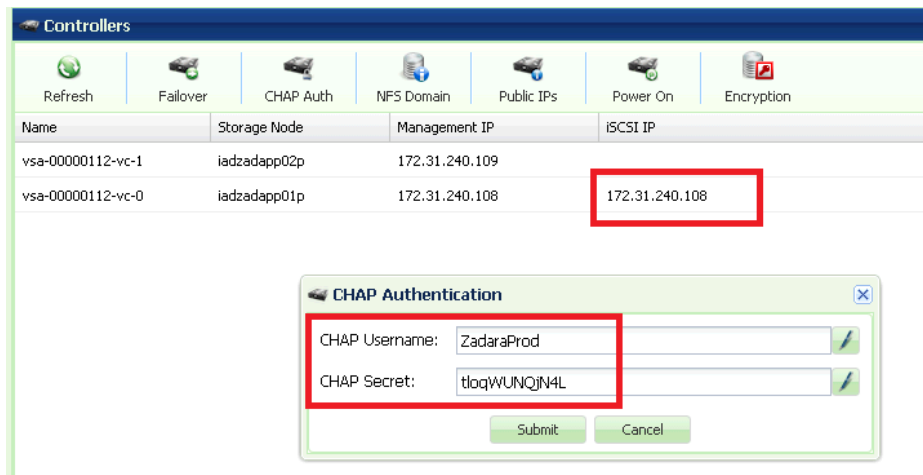


- In the **Targets** tab you'll see that the VPSA iSCSI target has moved from "Inactive" to "Connected" status. A new Server is created automatically in the VPSA and is displayed in the Servers GUI page. The name of the server is its iSCSI initiator IQN. You may change the Server Display Name.

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 Controllers 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 the 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.

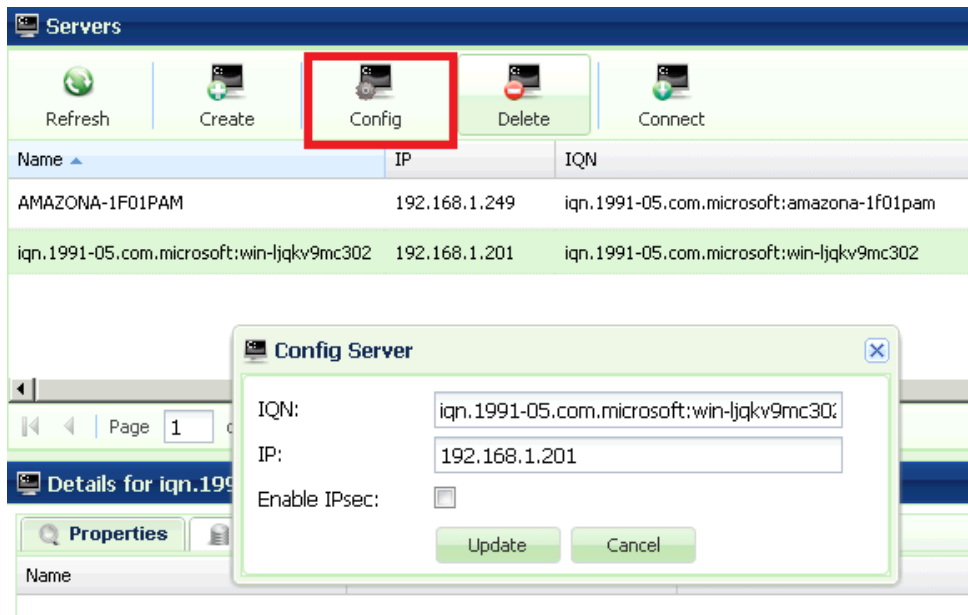
Note: To ensure that your Server automatically login to the VPSA after each reboot (or iscsid restart), run the following command on your Linux Server:

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

7.1.4 Configure Server Attributes

You can change the following Server Attributes using the Config Server dialog:

- Server IQN
- Server IP address
- Enable\Disable IPSec



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.

7.2 Viewing Servers Properties

The Servers Page displays a list of available Server objects, providing the following detailed information in the Servers details south panel tabs:

Name	IP	IQN	OS	Registered	Status	IPsec
AMAZONA-1F01PAM	192.168.1.249	iqn.1991-05.com.microsoft:amazona-1f01pam	Microsoft Windows Server 2008 R2 Datacenter 6.1.7601	yes	Active	Disabled
iqn.1991-05.com.microsoft:win-ljqkv9mc302	192.168.1.201	iqn.1991-05.com.microsoft:win-ljqkv9mc302		no	Active	Disabled

ID:	srv-00000002
Name:	iqn.1991-05.com.microsoft:win-ljqkv9mc302
IP:	192.168.1.201
IQN:	iqn.1991-05.com.microsoft:win-ljqkv9mc302
IPsec:	Disabled
Registered:	no
OS:	
Status:	Active
Added:	2013-11-19 06:39:05
Modified:	2013-11-19 06:39:05

Properties

Each Pool includes 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.
IP	IP Address of the Server.
IQN	Unique “iSCSI Qualified Name” of the Server.
IPSec	Enabled\Disabled
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.

Metering

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

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, 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 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.

8 Managing Volumes, Snapshots and Clones

VP SA virtual Volumes are thinly provisioned, utilizing an efficient and sophisticated block-level mapping layer. The Volume's virtual address space is chunked into virtual contiguous blocks (a.k.a. "Chunks"). When you create a Volume, it consumes zero Pool capacity. The first write to each Chunk triggers the provision of Pool capacity and mapping update of the virtual to physical addresses, thus Pool capacity is provisioned on demand.

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

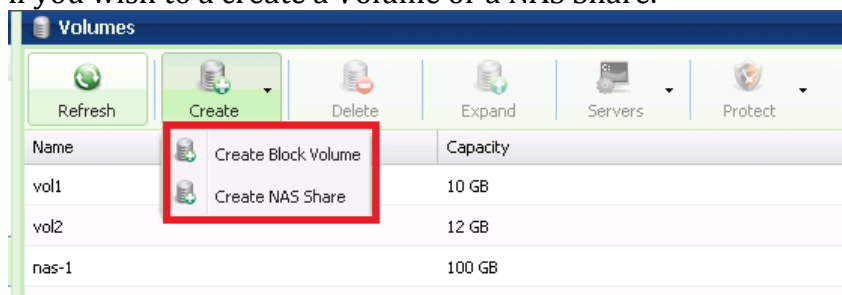
Snapshots represent read-only representations of the Volume's data at a given point-in-time. They are thinly provisioned as well, sharing data Chunks with the Volume as much as possible until you actually modify the Chunk's data. This triggers a Copy On Write (COW) operation where a new chunk is provisioned 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, and COW is triggered whenever you modify a chunk in the Volume or in the Clones.

Volumes can be block Volumes (exposed via an iSCSI protocol) or NAS Shares (exposed via NFS or SMB protocols).

8.1 Creating and Deleting a Volume

To Create a Volume, go to the Volumes Page and press the **Create** button. Select if you wish to create a Volume or a NAS Share.



Creating a Block Volume

Create Thin Volume

Name: Vol-SQL

Capacity (GB): 200

Pool:

Name	Status	Free Capacity
Pool_Webfiles	normal	1.6 TB / 1.6 TB
pool_DB	normal	1.1 TB / 1.1 TB

Apply Snapshot Policy: ☒

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
Weekly Snapshots for a Year	Every Sunday at midnight	Keep latest 53 snapshots

Encrypted: ☒

Submit Cancel

Define the following Volume attributes in the “Create Thin Volume” dialog:

- **Name** – the Volume’s display name. Needs to be unique, and can be modified throughout the Volume’s lifetime.
- **Capacity** – Virtual Capacity of the Volume in GB. As all Volumes are thinly provisioned, no actual capacity is allocated when the Volume is created; hence the Virtual capacity is not bounded by the Pool capacity. You are permitted to over-provision a Pool, but you need to handle it carefully using Pool Protection Mechanism (check [here](#) for more details).
- **Pool** – Select the Pool that is most appropriate for your Volume’s QoS requirements (available capacity, caching, RAID protection, drive types etc).
- **Apply snapshot Policy** – Check [here](#) for detailed explanation regarding snapshot policies. You can apply and remove snapshot policies from a Volume at any time. If you select this check-box, you are requested to select one of the existing snapshot policies.
- **Encrypted** – Select the check-box if you wish to encrypt the volume’s data on the drives. Please note that you must first define an encryption password via the Controller Page. For more details about Volume encryption please check [here](#).

Creating a NAS Share

Create NAS Share

Name:

Capacity (GB):

Export Name:

atime Update: ☐

SMB Only (optimized for performance): ☐

SMB - Allow Guest Access: ☐

Pool:

Name	Status	Free Capacity
Pool-backup	normal	1.0 TB / 1.1 TB
Pool-webfiles	normal	2.7 TB / 2.7 TB
Pool-DB	normal	98 GB / 183 GB
Pool-T	normal	2.6 TB / 2.7 TB

Apply Snapshot Policy: ☒

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
Weekly Snapshots for a Year	Every Sunday at midnight	Keep latest 53 snapshots
every minute for an hour	Every 1 minute	Keep latest 60 snapshots
every 5 min for a day	Every 5 minutes	Keep latest 288 snapshots

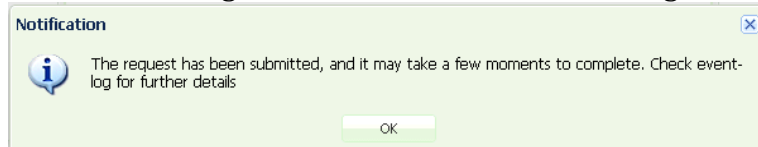
Encrypted: ☐

Define the following Volume attributes in the “Create Share” dialog:

- **Name** – the Share’s display name. Needs to be unique, and can be modified throughout the Share’s lifetime
- **Capacity** – Virtual Capacity of the Share in GBs. As all Shares are thinly provisioned, no actual capacity is allocated when the Share is created; hence the Virtual capacity is not bounded by the Pool capacity. You are permitted to over-provision a Pool, but you need to handle it carefully using Pool Protection Mechanism (check [here](#) for more details)
- **Export Name** –The name of the NFS\SMB mount point as seen by the Server. Must be Unique. By default it is identical to the Share name.
- **atime Update** – indicates whether to update access time of files and directories on every access, including read-access. By default atime Update is disabled. Enabling it will impact performance.
- **SMB Only** – Set this checkbox if you know that this NAS share will be attached to Servers only via the SMB protocol. In this case, VPSA can do some locking optimization to enhance performance.
- **SMB – Allow Guest Access** - Set this checkbox if you want to enable connection and access to the NAS share to anonymous users without requiring a password.
- **Pool** – Select the Pool that is most appropriate for your Share’s QoS requirements (available capacity, caching, RAID protection etc.).
- **Apply snapshot Policy** – See [Managing Snapshot Policies](#) for detailed explanation regarding snapshot policies. You can apply and remove snapshot policies from a Share at any time. If you select this check-box, you are requested to select one of the existing snapshot policies.

- **Encrypted** – Select the check-box if you wish to encrypt the Share’s data on the drives. Please note that you must first define an encryption password via the Controller Page. For more details about Volume encryption please check [here](#).

Note: Share creation involves the process of initializing a file system. It may take a few minutes depending on the Virtual capacity of the Share. During that time the share will be in “Creating” state. When initialization is completed, the Share’s status will change to “Available” and an event-log message will be saved.

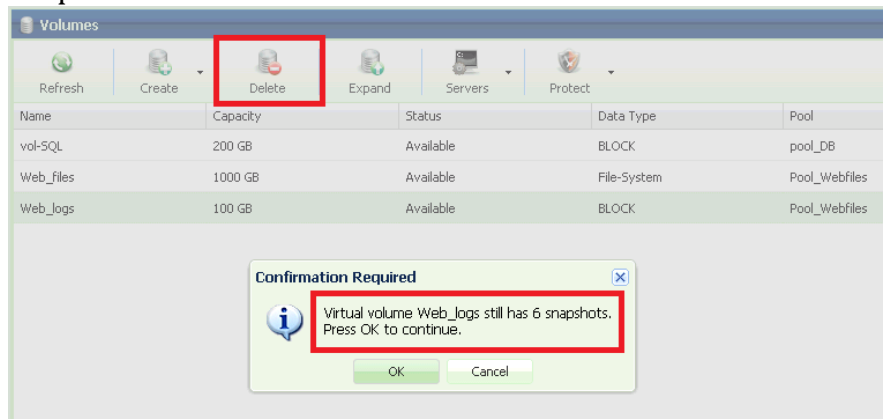


Deleting a Volume\Share

You can delete a Volume if it is not attached to any Server.

In the Volumes page, select the Volume and press the **Delete** button. It will move the Volume to “Deleting” status immediately. 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 needs to update chunk mapping and references accordingly.) When the deletion process completes, the Volume will disappear from the Volume page and an event-log message will be saved.

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



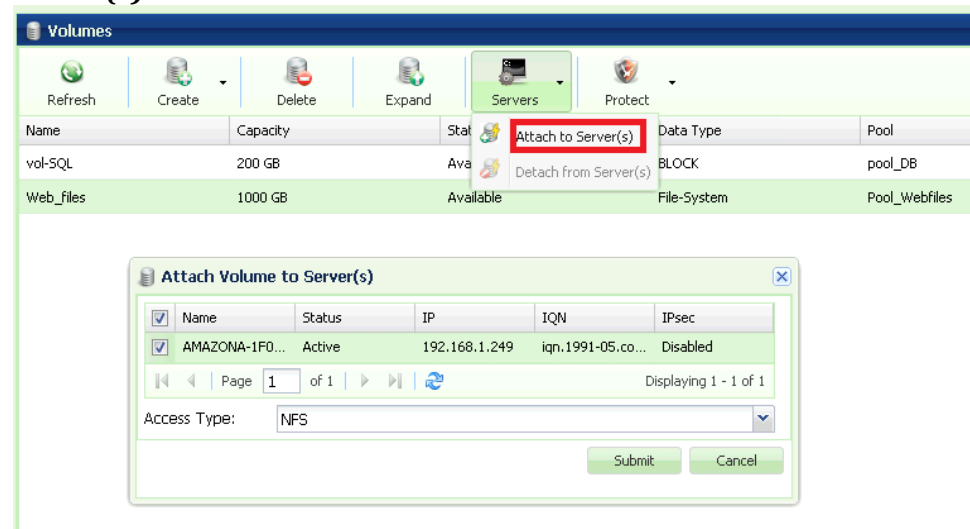
Clones of the deleted Volume are not affected by the deletion of the Volume.

8.2 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:



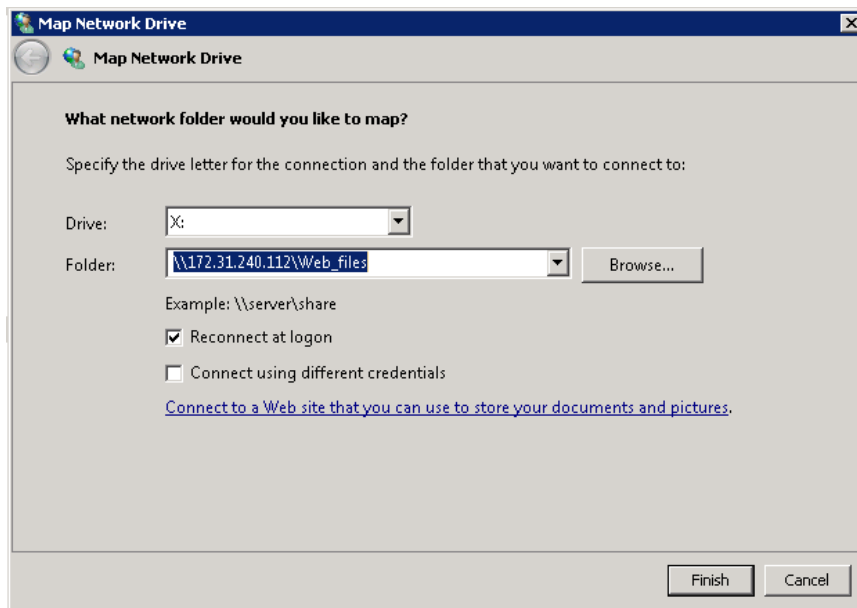
- Select the Server(s) that you'd like to have access to the Volume.
- For NAS Shares, select the access type: NFS or SMB.
- Press **Submit** to confirm.

Mounting an NFS Share on a Linux machine

1. Install NFS client:
 - a. On Ubuntu Servers do: 'apt-get install nfs-common';
 - b. On Redhat/CenOS Servers do: 'yum install nfs-utils'
2. Create a mount point:
 - a. `$ mkdir /mnt/nfs_share`
3. Run the following command as the superuser (or with sudo):
 - a. `$ mount -t nfs4 <NFS_Export_Path>/<mount point>`
 - b. You can find the **NFS_Export_Path** in the VPSA GUI, Volume page->Properties tab
4. Follow the step in [Managing NFS Users Access Control](#) to setup basic NFS authentication.

Mounting a 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, Volume page->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).

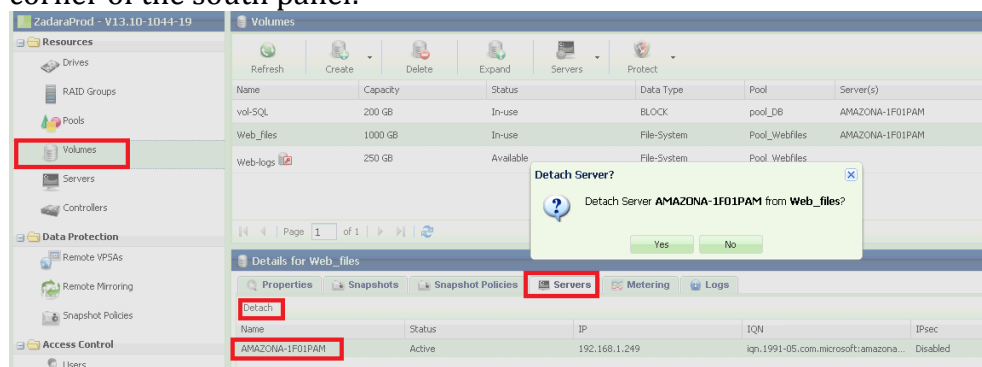


To detach a Volume

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

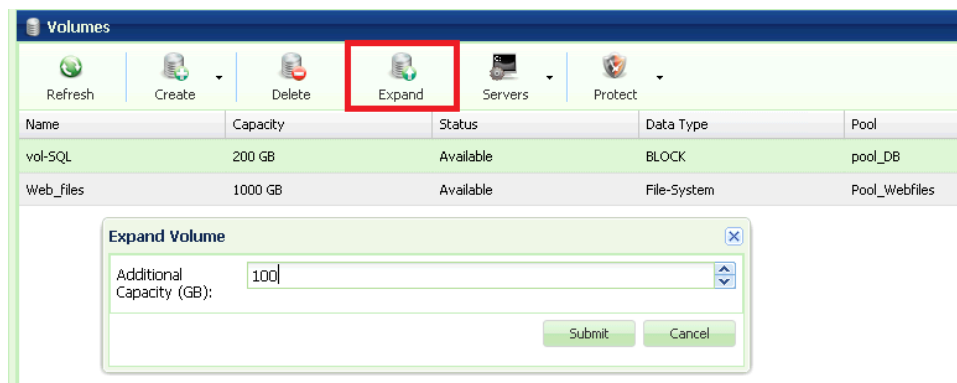
To detach a Volume from a Server, go to the Volumes Page and click the Servers->**Detach from Server(s)** button. You will be requested to select the Servers to detach this Volume.

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:



8.3 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.




8.4 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 which were acknowledged to the Server by the VPSA before the Snapshot was created will be contained in the Snapshot's data set.

8.4.1 Manual creation & deletion of Snapshots

To manually create a Snapshot:

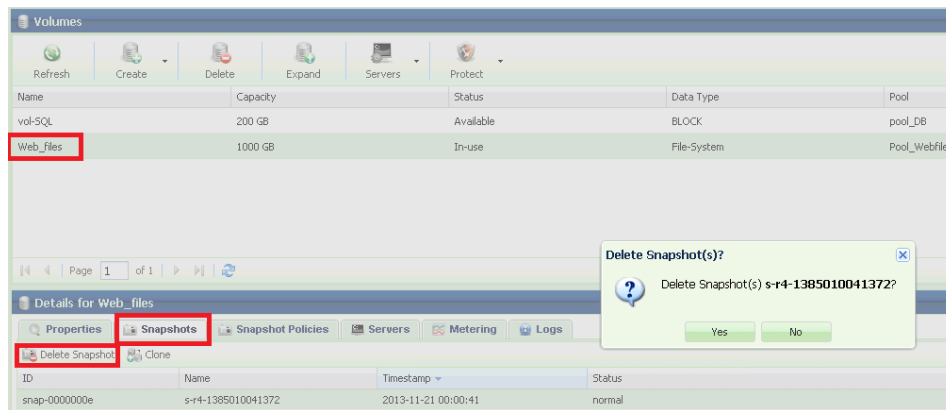
- Go to the Volumes page, press the **Protect** button and select **Create Snapshot**.



- 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 tab to display the list of snapshots associated with this Volume
- Select the snapshot to be deleted in the Snapshot 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 later 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.

8.4.2 Managing Snapshot Policies

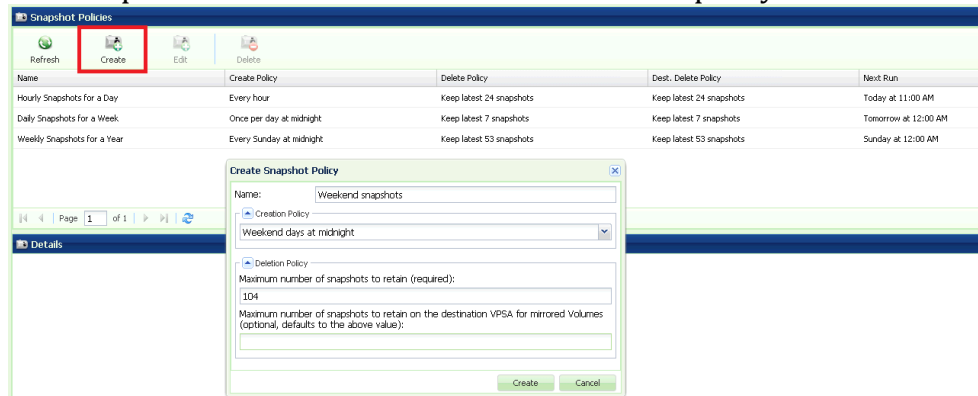
Snapshot policies define the Snapshots lifecycle via enforcement of creation and deletion policies. Snapshot Policies are “global” entities, and you can apply instances of the policies to one or more Volumes. Unapplied policies are idle—they do not consume any resources and never create any snapshots. 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 be deleted only when all relevant delete policies approve its deletion.
- Snapshot creation time is a “rounded” time, regardless of the policy creation time. For example, if you created a Snapshot Policy at 9:02 which has a creation policy to create a snapshot every 10 minutes, the Snapshots will be created at 9:10, 9:20, 9:30 etc. (not at 9:12, 9:22, 9:32 etc.).
- You can decide whether empty snapshots are to be created or not. 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, should a new and empty Snapshot be created.
- Few Snapshots Policies are predefined in the VPSA.

To create a new Snapshot Policy:

- Go to the Snapshot Policies page and press the **Create** button.
- Provide a meaningful name to the 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.

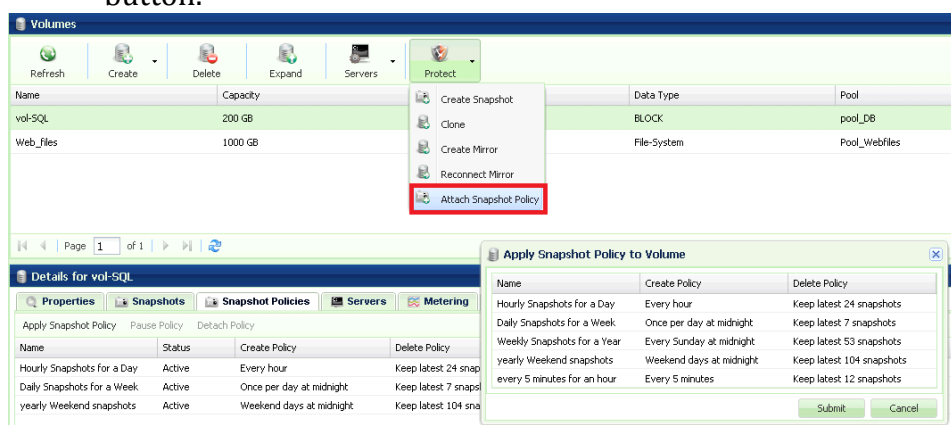


To Edit a Snapshot Policy

- 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, and Allow Empty Snapshots rules.
- Note: 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

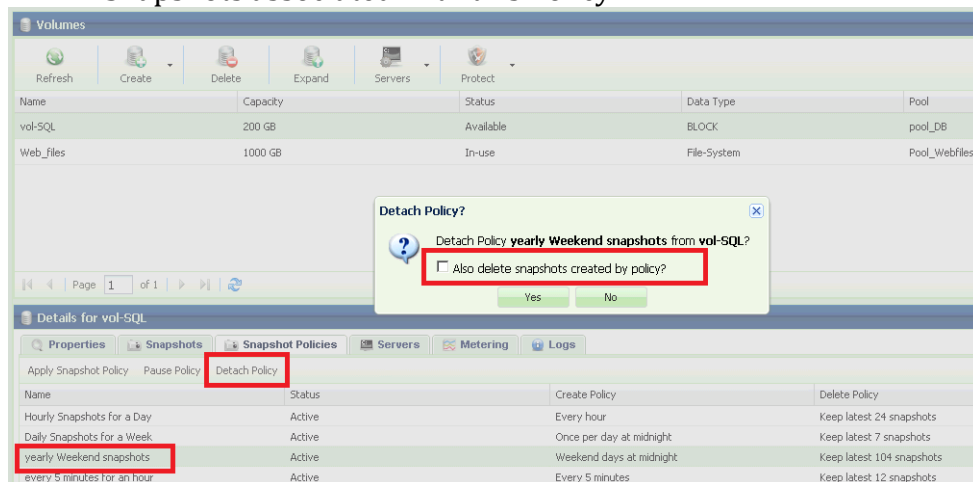
- Go to the Volumes page, select the Volume and select Protect->Attach Snapshot Policy from the menu.
- Select the Snapshot Policy to apply to the Volume, and press the **Submit** button.



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 associated with this Policy.



Pause\Resume a Snapshot Policy

You can pause an active Volume Snapshot Policy. New Snapshots will not be created. 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 will move to Paused status.
- To resume a Policy: The Pause\Resume button toggles according to the Policy status. Select a Policy in Paused state and press the **Resume Policy** button. The Policy Status will move to active.



Note: When the Enterprise Suite is disabled, all active Snapshot Policies are immediately paused, and when the Enterprise Suite is re-enabled, policies are resumed.

8.5 Cloning a Volume

Cloning a Volume is the process of creating a Read\Write zero-capacity replica of a Volume, with an identical data set as it exists in the Volume, in 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 for the same data set (from the same Snapshot) or with different data sets.

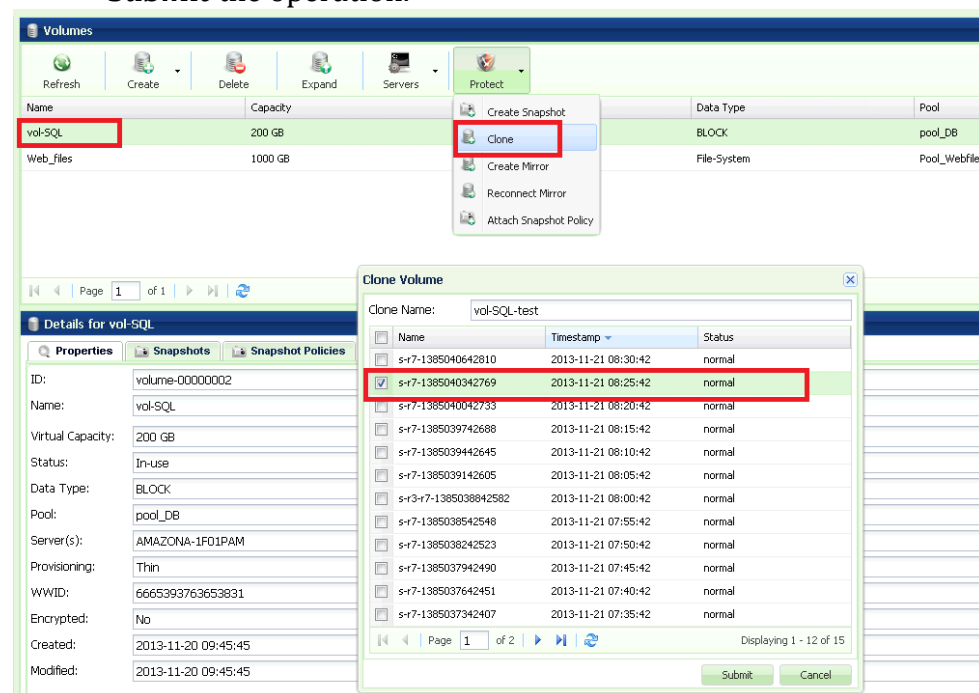
Clones are independent from each other, from the source Volume and from the Snapshots from which they were created. For example, you can delete the original Volume and/or Snapshot and leave the Cloned Volume unaffected. You can also modify Volume attributes of each Clone independently.

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

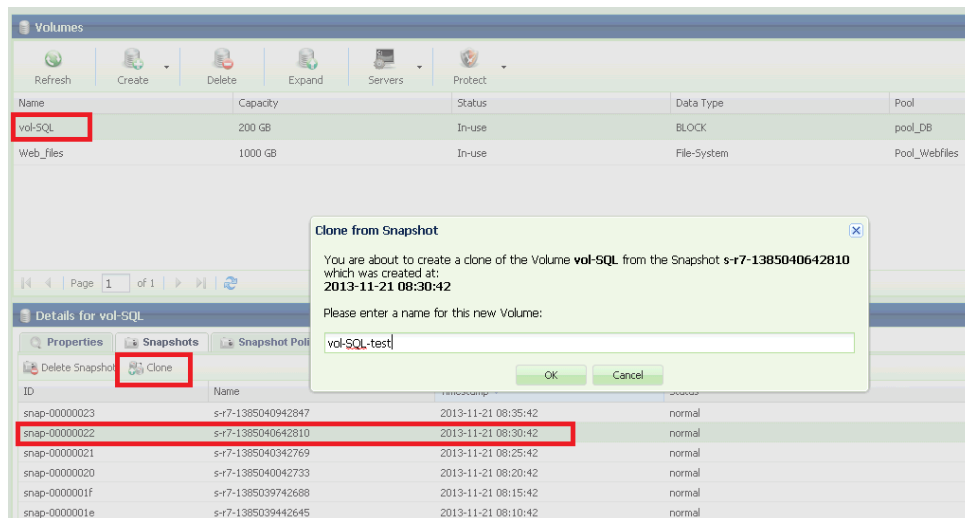
To create a new Clone

Go to The Volumes page, select the Volume to be cloned, and press the Protect->**Clone** button.


- Enter the Cloned Volume name.
- Select the point-in-time Snapshot which you wish to replicate its data set. If you wish clone the current data set of the Volume, don't select any Snapshot and,
- **Submit** the operation.



- 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 the new cloned Volume name.



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



The Export name of a cloned Share will be identical to the Cloned Share display name.

8.6 Managing Encrypted Volumes

Encryption management of Data-at-Rest (data on the Disk Drives) is done by the VPSA on a per-Volume basis, i.e., some Volumes can be encrypted while others cannot.

A VPSA generates a random 128-bit **unique Volume Encryption Key** per encrypted Volume, and uses the Advanced Encryption Standard (AES) to encrypt and decrypt the Volume data.

The Volume Encryption Keys are stored on disk as ciphertext, using AES with a 128-bit Master Encryption Key, which is generated from a user-supplied **Master Encryption Password**.

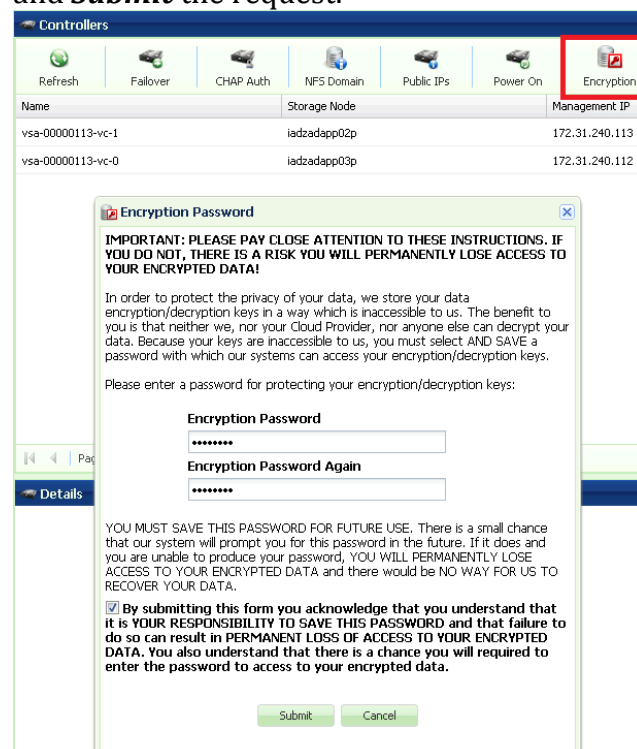
The User owns the Master Encryption Password. It is never stored on any persistent media. Instead, only its SHA1 hash-sum is saved on disk for password validation. Since it is virtually impossible to restore the Master Encryption Password from the SHA1 hash-sum, you are fully responsible to retain and 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.

The above 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 Volumes encryption.

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 Controllers page, and press the **Encryption** button. Read the instructions and warning. Type your Password and **Submit** the request.



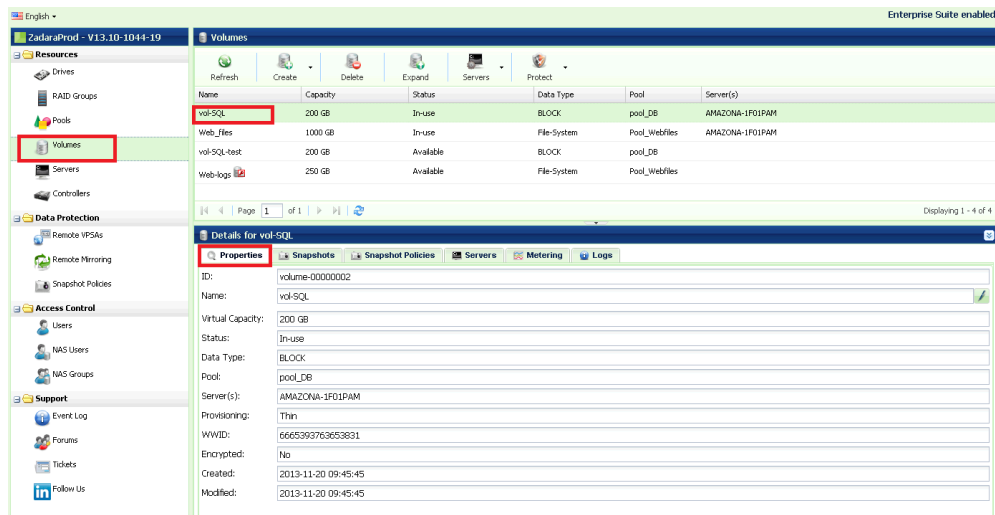
Store your Master Encryption Password in a secured place

To create an Encrypted Volume, follow the steps in section [8.1 - Creating and Deleting a Volume](#).

Encrypted Volumes are displayed with the  icon.

8.7 Viewing Volume Properties

The Volumes Page displays the list of Volumes (Block and NAS) in the VPSA. Select a Volume and see its detailed information in the following south panel tabs:



Properties

Each Volume includes the following properties:

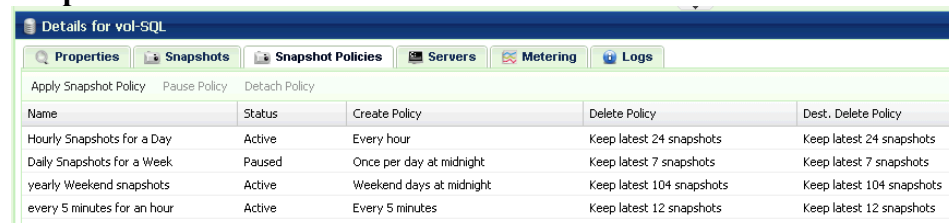
Property	Description
ID	An internally assigned unique ID.
Name	User assigned name. Can be modified anytime.
Virtual Capacity	Capacity of the Volume as seen by the attached Servers.
Available Capacity	Free capacity of the NAS Share.
Mapped Capacity	The used capacity (allocated from the Pool) of the Volume excluding its Snapshots and Clones
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	The NFS Share export path to be used when mounting it.
SMB Export Path	The SMB Share export path to be used when connecting to it from a Windows Server.
Access Type	Access protocols which are used by the Servers which are attached to a NAS Share: NFS, SMB, or Multiple.
atime Update	Yes\No – Indicates whether to update access time of NAS Share files and directories on every access, including read-access.
SMB Only	Yes\No
SMB Guest Access	Yes\No – Allow\Block anonymous user access
Encrypted	Yes\No
WWID	SCSI unique World-wide ID. Use this value on Linux Servers to identify the Volume device when multipathing is configured.
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 check [here](#).

Snapshot Policies



Details for vol-SQL				
Properties	Snapshots	Snapshot Policies	Servers	Metering
Apply Snapshot Policy	Pause Policy	Detach Policy		
Name	Status	Create Policy	Delete Policy	Dest. Delete Policy
Hourly Snapshots for a Day	Active	Every hour	Keep latest 24 snapshots	Keep latest 24 snapshots
Daily Snapshots for a Week	Paused	Once per day at midnight	Keep latest 7 snapshots	Keep latest 7 snapshots
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

List of the Snapshot Policies which 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
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, check [here](#).

Servers

Lists the Servers to which the Volume is attached. For block Volumes, the Lun Number associated with each Server is displayed.

Metering

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

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, 1 Minute, 10 Minutes, or 1 Hour. The **Auto** button lets you see continues updated 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 selected Volume from all attached Servers.
Bandwidth (MB\s)	Total throughput (in MB) of Read and Write SCSI issued to the selected Volume from all attached Servers.
IO Time (ms)	Average response time of all Read and Write SCSI issued to the selected Volume from all attached Servers.

Logs

Displays all event logs associated with this Volume.

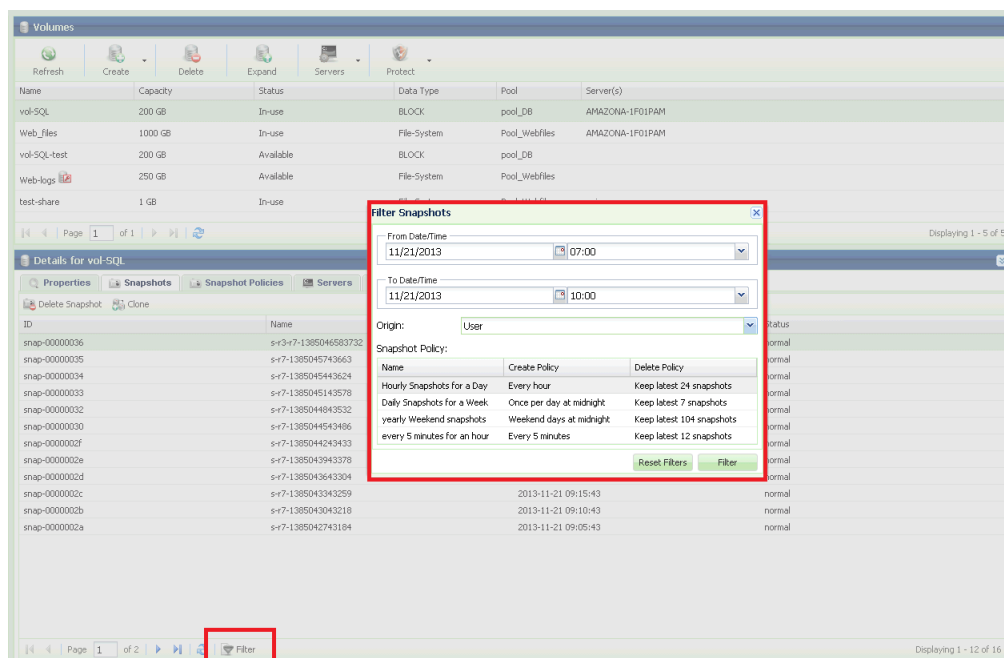
8.8 Filtering Snapshots

With VPSA you can create many snapshots on every Volume. Snapshots can be created manually, or by using Snapshot Policies, or by Remote Mirroring.

Finding a snapshot can be tricky. The Filter Snapshot option will help you find the specific snapshot you need.

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. 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 Snapshot origins.
 - User – Snapshot created manually or via a Snapshot Policy, which was attached to this Volume.
 - Mirror – Snapshots which were created by the Remote Mirroring application (using the Snapshot policy which was defined at the time of the Mirror creation).
- Snapshot Policy – Select a Policy if you'd like to filter only Snapshots that were created by that Policy.



9 Managing Access Control

9.1 Adding & Deleting Users

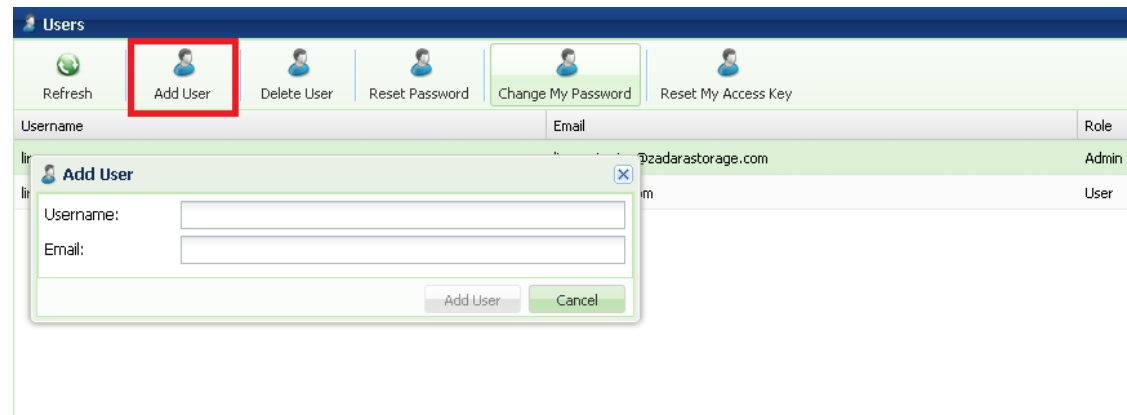
The VPSA's User Management system supports multiple users. There are two distinct roles:

- **Admin** - The Admin user can add and delete users and reset user's password through the VPSA GUI. There is a single Admin user per VPSA, and that is the user who created the VPSA via the VPSA Management Console.
- **User** - A user who was added by the Admin user. This user has full rights to manage the VPSA either through the GUI or REST APIs. Each user has its own Password and Access Key.

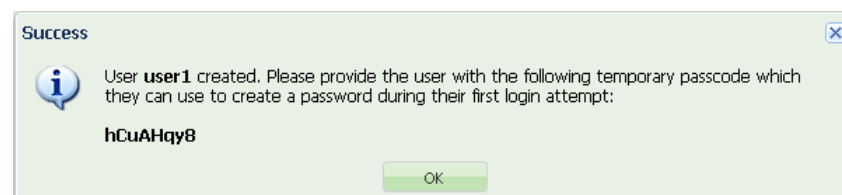
Adding a new User

Log in to the VPSA with Admin user credentials. Go to Users page and click Add User.

Enter the Username and Email address, and press Add User button to confirm 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 temp passcode the first time logging to the VPSA.



Deleting a User

Log in to the VPSA with Admin user credentials. Go to Users page, select the user from the users list and click the **Delete User** button.

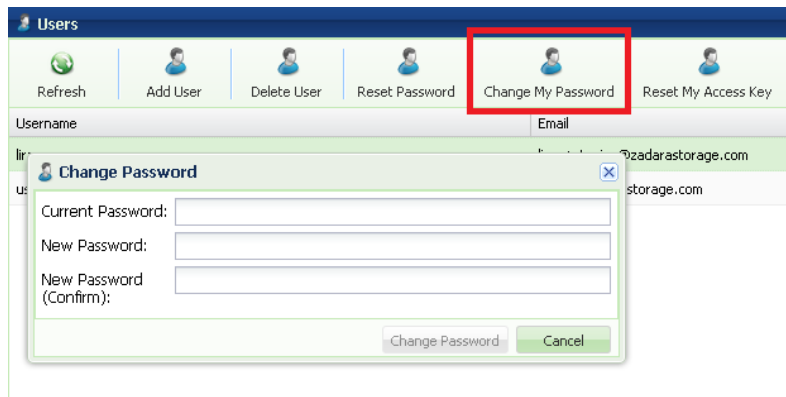
The User will be deleted. No other entities (created or managed by that user) will be affected by this operation.

9.2 Managing User Passwords

The VPSA stores in its database a cryptographic hash value (using a one-way SHA-1 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 in the database.

Changing your password

Log in to the VPSA, go to the Users page and click **Change my 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 User and to all regular Users. If you log as the Admin User and see the full list of the available VPSA Users, the **Change my Password** button will be enabled only when your Admin User is selected.

Resetting User Password

This operation is available only to the Admin User. You can reset any User's password. A new temporary passcode will be created and sent to the Admin User 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 Users page, select a user from the users list and click **Reset Password** button

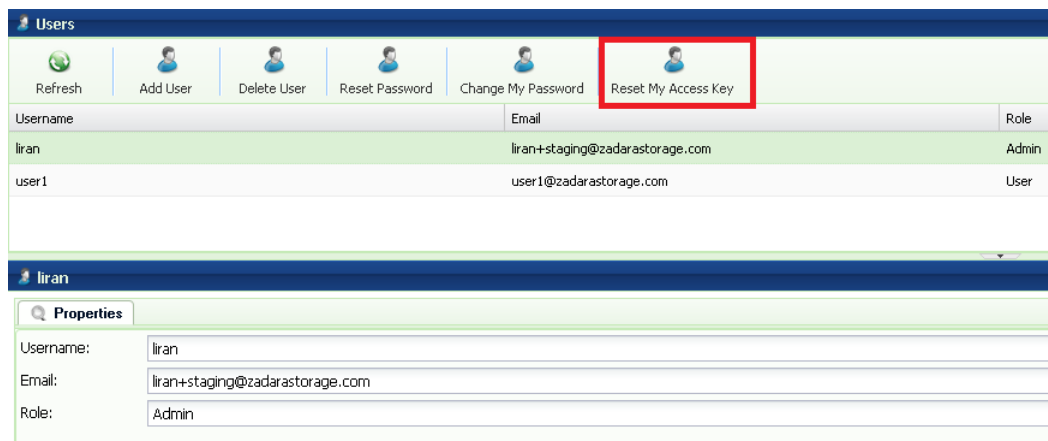
Resetting API Key

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.

Go to the Users page, and press the **Reset My Access Key** button.



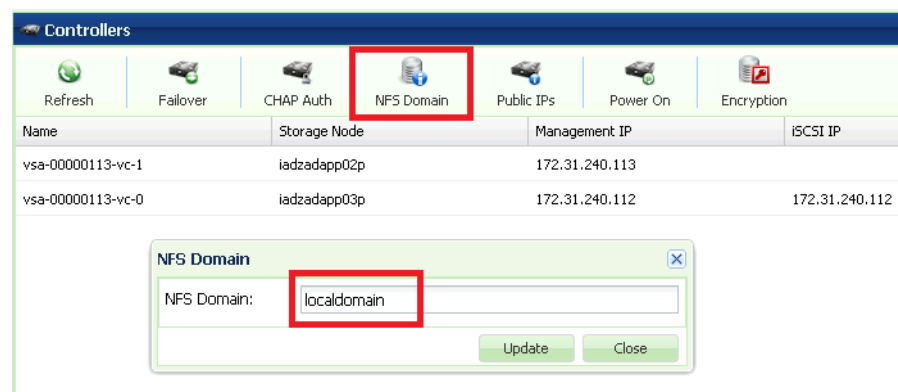
9.3 Managing NAS Users Access Control

9.3.1 Creating NFS Users

By default, "root" user and group, at the NFS client, are mapped to "root" user and group in the VPSA NFS server. All other client-side users are mapped to user "nobody" and group "nogroup."

In order to set a basic NFS authentication so that users and groups at the client will be mapped to the corresponding users and groups at the VPSA NFS server, perform the following steps:

- Go to VPSA GUI > Controllers and press the **NFS Domain** button. The NFS Domain dialog will appear:



- Enter NFS domain name identical to the domain name set in the Client and press the **Update** button. Typically, the default domain name on the Linux client side is "localdomain" and therefore this is also the default value in the VPSA.

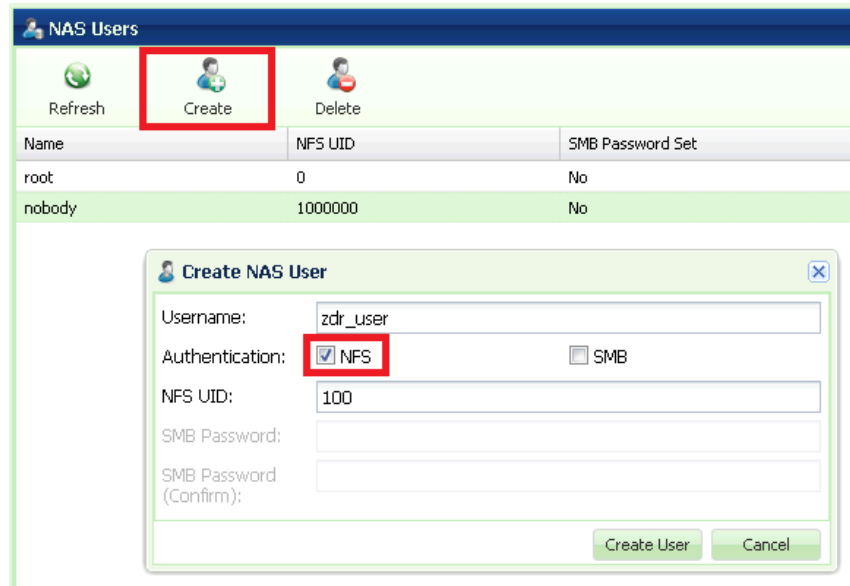
NOTE1: In a Linux client, the domain name is usually set in the `/etc/ldap.conf` file. It is mandatory to have this value set.

NOTE2: Make sure `idmapd` service is running (In Ubuntu `idmapd` service. In RedHat `rpcidmapd`) and that `/sys/module/nfs/parameters/nfs4_disable_idmapping` is set to "N" (to make it persistent set:

`GRUB_CMDLINE_LINUX_DEFAULT="nfs.nfs4_disable_idmapping=N"` in `/etc/default/grub` and then run `update-grub`).

- Go to VPSA GUI > NAS Users, and press the **Create** button.

- a. Set the user name.
- b. Check NFS authentication.
- c. Select an NFS UID (in the range 1-999,999).
- d. If you wish to grant this user access to SMB shares as well, check the SMB check-box and provide the User SMB Password (to be used later when mounting a share on a Windows Client).



9.3.2 Creating SMB Users

- Go to VPSA GUI >NAS Users, and click the **Create** button.
- Set the user name.
- Check SMB authentication.
- Provide the User SMB Password. You will be asked to provide this user name 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, check the NFS check-box and select an NFS UID on the range of 1-999,999.

The screenshot shows the 'NAS Users' management interface. At the top, there are three buttons: 'Refresh', 'Create', and 'Delete'. Below them is a table with three columns: 'Name', 'NFS UID', and 'SMB Password Set'. The table contains two rows: 'root' with NFS UID '0' and 'SMB Password Set' 'No', and 'nobody' with NFS UID '1000000' and 'SMB Password Set' 'No'. A 'Create NAS User' dialog box is open in the foreground. It has fields for 'Username' (filled with 'zdr_smb'), 'Authentication' (with radio buttons for 'NFS' and 'SMB', where 'SMB' is selected and highlighted with a red box), 'NFS UID' (empty), 'SMB Password' (masked with dots and highlighted with a red box), and 'SMB Password (Confirm)' (masked with dots and highlighted with a red box). At the bottom of the dialog are 'Create User' and 'Cancel' buttons.

Name	NFS UID	SMB Password Set
root	0	No
nobody	1000000	No

9.3.3 Editing SMB Users Password

It is possible to edit the password of an SMB User at any time. Go to NAS Users page and select **Edit SMB Password**:

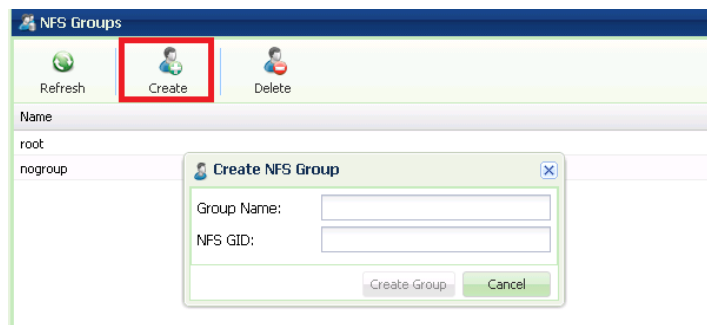
The screenshot shows the 'NAS Users' management interface. At the top, there are three buttons: 'Create', 'Delete', and 'Edit SMB Password'. The 'Edit SMB Password' button is highlighted with a red box. Below the buttons is a 'Change SMB Password' dialog box. It has two input fields: 'New SMB Password' and 'New SMB Password (Confirm)'. At the bottom of the dialog are three buttons: 'Change Password' (highlighted with a red box), 'Remove Password' (highlighted with a red box), and 'Cancel'.

- 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 an NFS ID, you can select **Remove Password** to erase the user SMB password.

9.4 Creating NFS Groups

You can create and view NAS Groups via the NAS Groups page. Please note that NAS Groups are applicable only for NFS access control.

To create an NFS Group, go to the NFS Groups page and select **Create**:



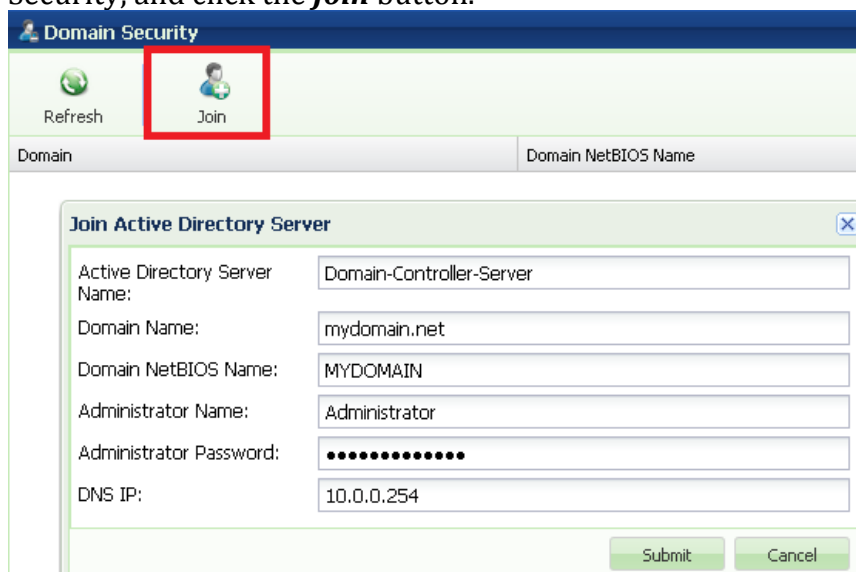
Enter a valid NFS Group Name and a valid NFS Group ID (on the range of 1-999,999) which matches the Group Name and ID on your Linux Server.

9.5 Enabling Active Directory Authentication

By joining the VPSA to the Active Directory (AD), users can use the same set of username and password that are stored in the AD to login the SMB shares.

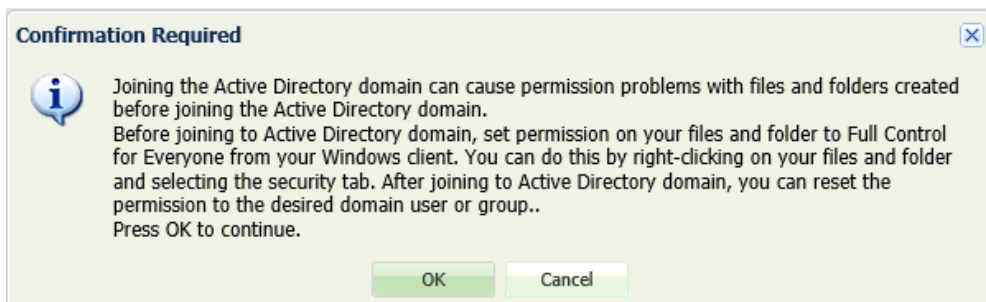
9.5.1 Joining the VPSA to Active Directory

To join the VPSA to the Windows Active Directory, Go to VPSA GUI > Domain Security, and click the **Join** button.



- In the opened dialog, fill out the following information:
 - Active Directory Server Name
 - Domain Name
 - Domain NetBIOS Name
 - Administrator Name (of the domain)
 - Administrator Password
 - DNS IP (The IP must be the same as the DNS server of your Active Directory)

- Click the Submit button and confirm the following warning message requesting to ensure proper permissions of files and folders created on the VPSA shares prior to joining the AD:



NOTE: The joining of the VPSA to the active Directory may fail if the time on both 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.

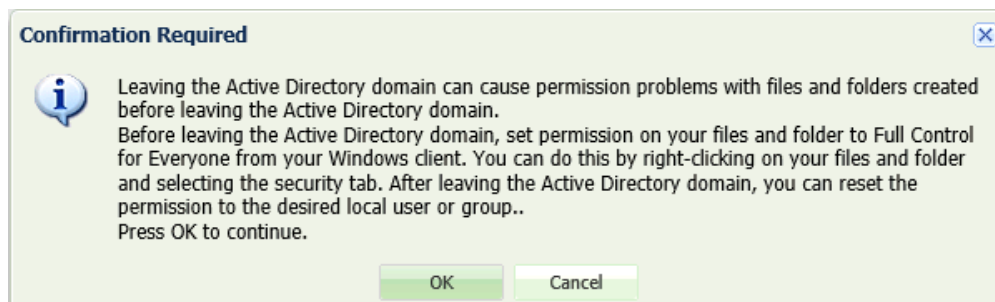
9.5.2 Leaving an Active Directory

To leave the Active Directory, Go to VPSA GUI > Domain Security, and click the **Leave** button (Join & Leave are a toggled button).



Enter the Domain Administrator's Name and Password and press submit.

confirm the following warning message requesting to ensure proper permissions of files and folders created using AD, before leaving it:



10 Managing Remote Mirroring

VPSA Asynchronous Remote Mirroring provides the ability to replicate your VPSA's data asynchronously to a different VPSA, either locally, within the same Zadara Cloud, or remotely to a VPSA located in a **remote region** or even a different cloud provider.

Asynchronous Mirroring doesn't impact IO response time from the Server perspective as the Server IO returns immediately after being written to the local VPSA storage. 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 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 where to resume the sync.

Remote VPSA communication is strongly authenticated and secured using cryptographic protocols that are designed to provide communication security over the Internet

For efficient bandwidth utilization, Mirrored data is compressed before being shipped to the remote VPSA in different region.

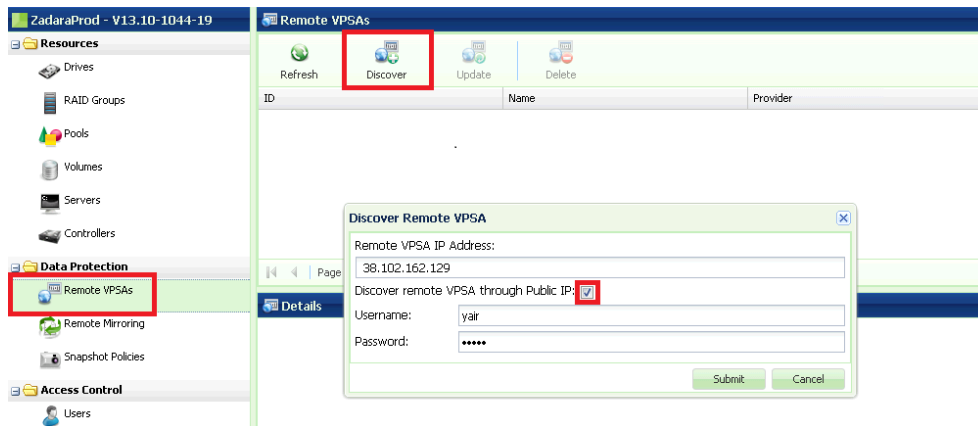
You can establish a many-to-many remote mirroring relationship for different Volumes between different VPSAs, i.e. a VPSA can mirror Volumes to many remote VPSAs and be at the same time the Destination VPSA for other Volumes in any other VPSA.

10.1 Connect to a remote VPSA

The first step for building a DR plan is to establish a trusted relationship between your VPSAs.

If the VPSAs are located in different Zadara Storage Clouds, you will need to first assign each VPSA a Public IP. See [here](#) for more details.

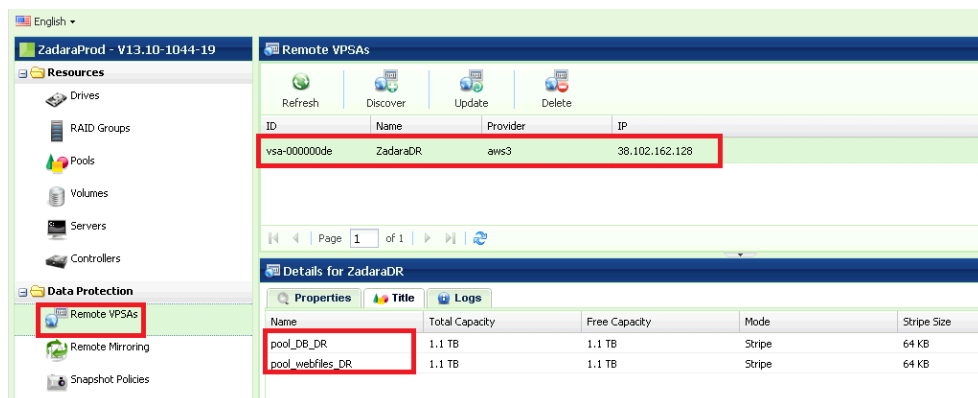
Go to the Remote VPSAs page and click the **Discover** button.



Enter the following details:

- **Remote VPSA IP Address:**
 - If the remote VPSA is located in a different Zadara Storage Cloud on 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.
 - Check the “Discover remote VPSA through Public IP” checkbox.
 - If the other VPSA is located in the same Zadara Storage Cloud:
 - Enter the remote VPSA Management IP address.
 - Don’t 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.

10.2 Viewing remote VPSA Properties

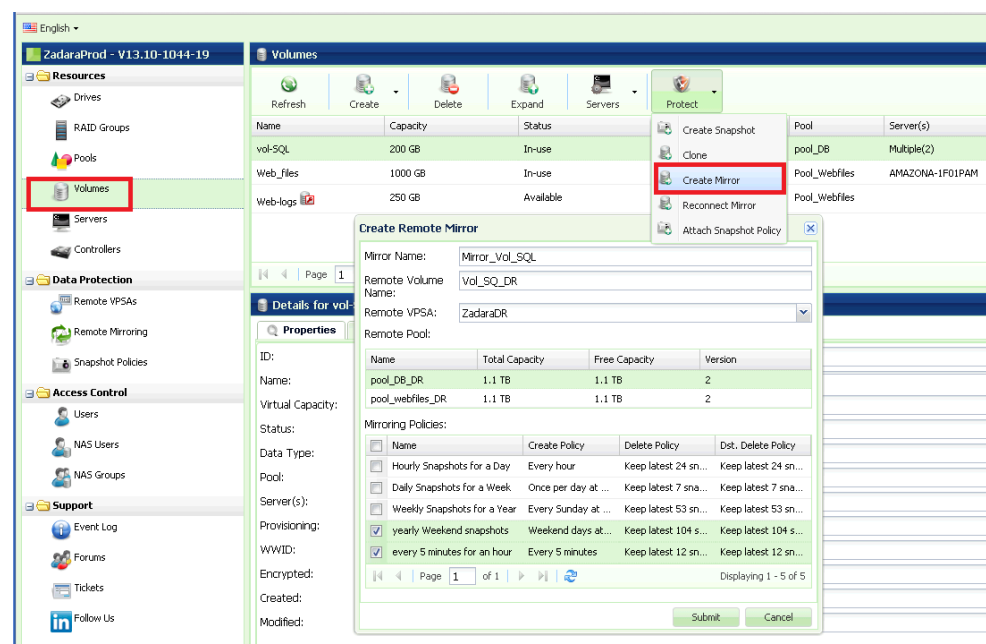


You can view all the remote VPSAs for which this VPSA has established a trusted relationship. For each VPSA the following details are provided:

- **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.
- **Pool list** – Each VPSA publishes the list of Pools that can be used to provision the remote Volume.
 - Note: This list is not updated automatically. Click the **Update** button to update the remote Pools info from the remote VPSA.

10.3 Creating a Remote Mirror

To create a Remote Mirror, go to the Volume page, and click the Protect->**Create Mirror** button. Alternatively, you can create the Remote Mirror from the Remote Mirroring page, by clicking **Create**. You will see a similar dialog:



Provide the following attributes to in the Creation dialog

- **Mirror Name**.
- **Remote Volume Name** – You can leave it empty, in which case the remote Volume name will be identical to the mirrored Volume name.
- **Remote VPSA** – Select any of the discovered remote VPSAs. Once a remote VPSA is selected, the Remote Pool list is populated accordingly.
- **Remote Pool** – Choose the Pool on the DR VPSA where the remote Volume and its snapshots will be provisioned.

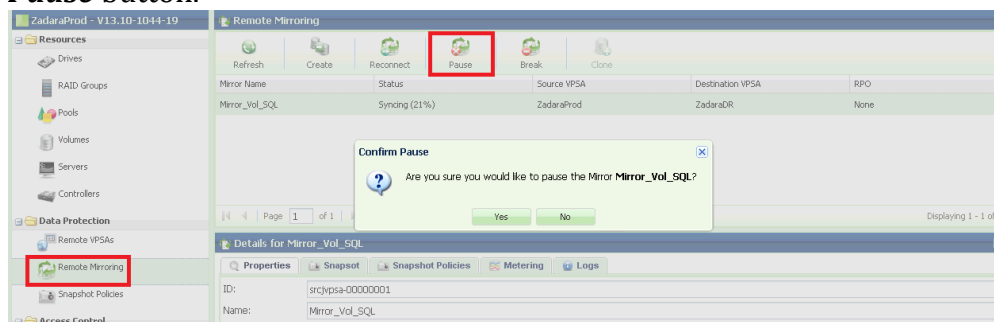
- **Mirroring Policies** – Select one or more existing Snapshots Policies. The Snapshot policies define the number of Snapshots to retain on Source and Destination Volumes and the minimum time between Snapshots (and hence the Return Point Objective). See [here](#) for more details on Snapshot Policies.

10.4 Pause & Continue Remote Mirror

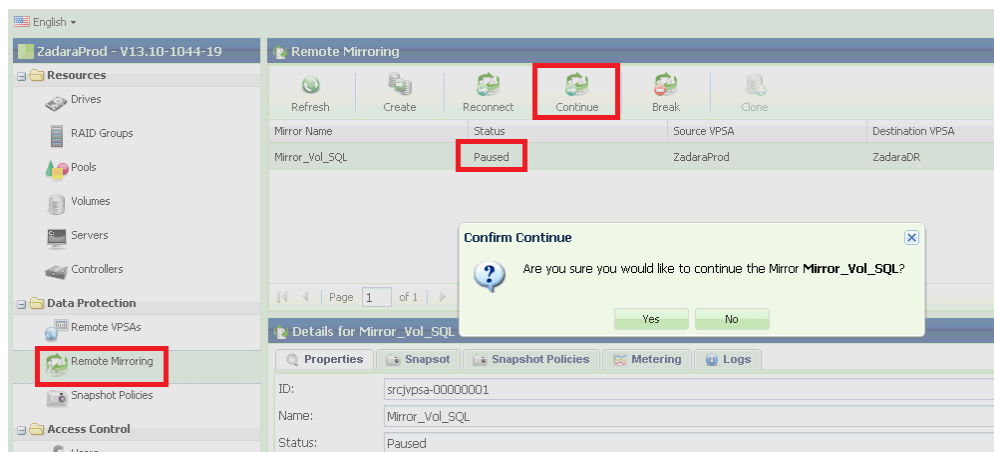
It is possible to pause a Remote Mirror. A paused Mirror will stop syncing data immediately, and stop creating new Snapshots. The status of the Mirror will move to Paused.

Note: A Mirror will be paused automatically if the Enterprise Suite is disabled.

To pause a Mirror, select the Mirror in the Remote Mirroring page and press the **Pause** button.



To Continue the Mirror operation, select the Mirror in the Remote Mirroring page and press the **Continue** button.



10.5 Managing Mirror Lifecycle

The Mirror controls the Remote Volume. 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 “Dest Volume.”

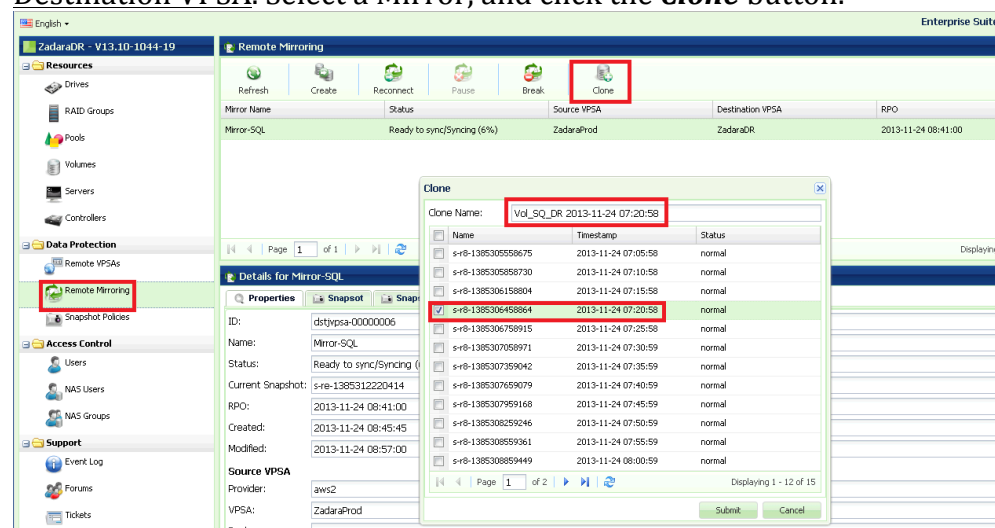
You can view Mirror destination Volumes on the “Dest Volumes” tab on the Pool Page, but they do not appear in the Volumes page.

10.5.1 Clone Destination Volume for Dev & Test of Remote Mirror

For Dev & Test of the Remote Mirror, you can Clone the destination Volume using the data set of any Snapshot that was completely synced. You cannot create a Clone of the Snapshot that is currently being synced.

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 Remote Mirroring page on the Destination VPSA. Select a Mirror, and click the **Clone** button.



- Select the point-in-time Snapshot which 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 & the timestamp of the selected Snapshot. You can modify this name at any time.
- You can find the newly created Volume in the Volume Page:

Volumes				
Refresh	Create	Delete	Expand	Protect
Name	Capacity	Status	Data Type	Pool
Vol_SQ_DR 2013-11-24 07:20:58	200 GB	Available	BLOCK	pool_DB_DR

10.5.2 Breaking a Mirror

Breaking a Mirror is the process of deleting the Mirroring relationship between the Source Volume and the Destination Volume, while leaving sufficient information for future Mirror reconnect. The Destination Volume then becomes a “regular” Volume and the source and the destination Volumes are now independent. A mirror can be broken from the source or from the destination VPSAs.

You can perform a future Mirror reconnect in both directions.

To break a Mirror, go to the Remote Mirroring page, select a Mirror, and click the **Break** button. After confirming the operation, the Mirror Object in the Remote Mirroring page will disappear from the source and the destination VPSAs.

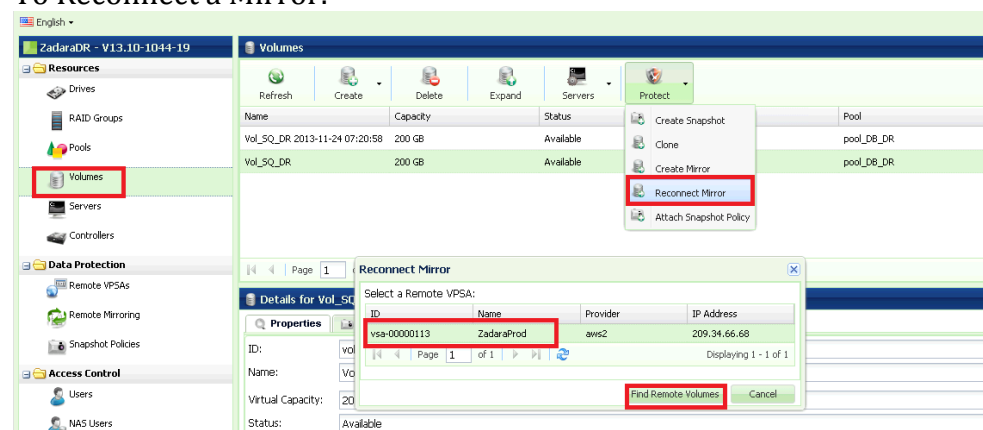
10.5.3 Reconnecting a Mirror

As previously described, the VPSA “leaves” sufficient metadata on each Volume after it breaks a Mirror, for future possible Reconnect of the Mirror relationship. This metadata allows the VPSA to identify a remote Volume on a Remote VPSA, which used to have a Mirroring relationship anytime in the past, with a given Volume, and find the most recent Snapshot that is in-sync on both Volumes. This then enables it to reconnect the Mirror relationship and resume the sync process from the most updated data set.

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.

Note: 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:



- Go to the Volumes Page, select the Volume you wish to act as the Source Volume of the Mirror, and click the **Protect->Reconnect Mirror** button. Alternately, you can go to the Remote Mirroring Page and click **Reconnect**, and select the Source Volume.
- Select the Remote VPSA that contains a Volume that used to be a mirror pair of the selected Source Volume in the past.

- Press the **Find Remote Volumes** button. The VPSA will query the Remote VPSA and display suggested Remote Volumes which can be Destination Volumes of the Mirror, with the following info:
 - **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 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.

The screenshot shows the 'Reconnect Mirror' dialog box. It has a title bar with a close button. Below the title bar is a section labeled 'Suggested Remote Volumes:'. Inside this section is a table with four columns: 'Remote Volume', 'New Data', 'Last Sync', and 'Snaps to Del.'. There is one row of data: 'vol-SQL', 'No', '2013-11-24 09:20:00', and '1'. The entire table is highlighted with a red border. Below the table is a pagination bar showing 'Page 1 of 1' and 'Displaying 1 - 1 of 1'. At the bottom of the dialog are 'Continue' and 'Cancel' buttons.

Remote Volume	New Data	Last Sync	Snaps to Del.
vol-SQL	No	2013-11-24 09:20:00	1

- Press **Continue**.
- Enter a name for the new Mirror.
- Select Snapshot Policies for the new Mirror.
- Press **Submit** to Reconnect the Mirror.

The screenshot shows the 'Reconnect Mirror' dialog box. It has a title bar with a close button. Below the title bar is a section labeled 'Mirror Name:' with a text input field containing 'Mirror-SQL-DR-to-Src'. Below that is a section labeled 'Mirroring Policies:'. It contains a table with three columns: 'Name', 'Create Policy', and 'Delete Policy'. There are three rows of policies: 'Hourly Snapshots for a Day' (checked), 'Daily Snapshots for a Week', and 'Weekly Snapshots for a Year'. Below the table is a pagination bar showing 'Page 1 of 1' and 'Displaying 1 - 3 of 3'. At the bottom of the dialog are 'Submit' and 'Cancel' buttons.

Name	Create Policy	Delete Policy
<input checked="" type="checkbox"/> Hourly Snapshots for a Day	Every hour	Keep latest 24 snapshots
<input type="checkbox"/> Daily Snapshots for a Week	Once per day at midnight	Keep latest 7 snapshots
<input type="checkbox"/> Weekly Snapshots for a Year	Every Sunday at midnight	Keep latest 53 snapshots

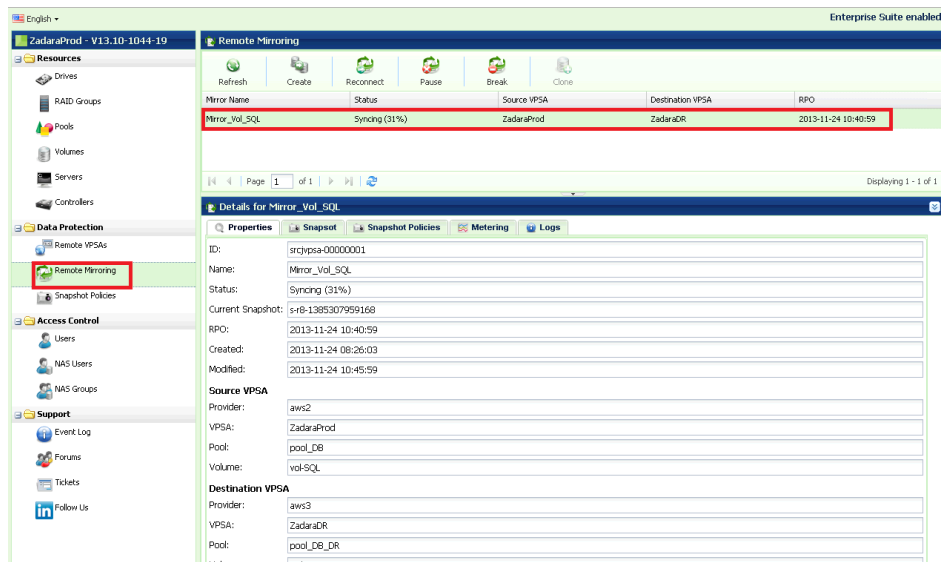
Note: Reconnect Mirror is blocked in the following cases:

- The Destination Volume is attached to a Server
- The Destination Volume has active Snapshot Policies

10.6 Viewing Remote Mirror Properties

Remote Mirroring Page displays the list of Remote Mirrors that the VPSA participates in either as the Source or the Destination. Mirrors are not symmetric and hence the info displayed is slightly different between Mirrors, which the VPSA is the source of, and those that it is the destination of.

Select a Mirror and check the detailed information in the following south panel tabs:



Properties

Each Remote Mirror includes the following properties:

Property	Description
ID	An internally assigned unique ID.
Name	User assigned name. Can be modified anytime.
Status	<ul style="list-style-type: none"> 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.
RPO	Return Point Objective – This is the timestamp of the most recent fully synchronized Snapshot.
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

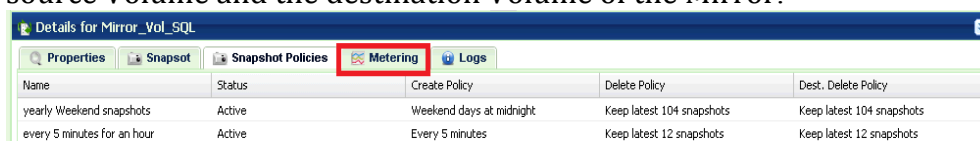
	the source VPSA.
Source VPSA\Volume	Source Mirror Volume name.
Destination VPSA\Provider	The name of the Cloud Provider where the destination VPSA resides.
Destination VPSA\ VPSA	Destination VPSA name.
Destination VPSA\Provider	Pool name where the destination Volume is provisioned.
Destination VPSA\ Pool	Destination Mirror Volume name.

Snapshots

Lists the point-in-time Snapshots of the Mirror on this VPSA. Please note that Mirror 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 check [here](#).

Snapshot Policies

List of active Snapshot Policies used by this Mirror to manage Snapshots on the source Volume and the destination Volume of the Mirror.



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

Source VPSA manages the Mirror Snapshots Policies; therefore modifications of the Mirror Snapshots Policies are allowed only on the Source VPSA.

The source VPSA will update the Remote VPSA regarding any change in the Dest Delete Policy.

You may make modifications while the Policy is active on a Mirror, and changes will be effective immediately. For example, if you change the policy to retain fewer snapshots, some older snapshots will be deleted immediately.

The following info is provided per Snapshot Policy on the Source VPSA:

- **Create Policy** – Minimum time between Snapshots.
- **Delete Policy** – How many snapshots to retain on Source Volume.
- **Dest Delete Policy** – How many snapshots to retain on the Destination Volume.

The following info is provided per Snapshot Policy on the Destination VPSA:

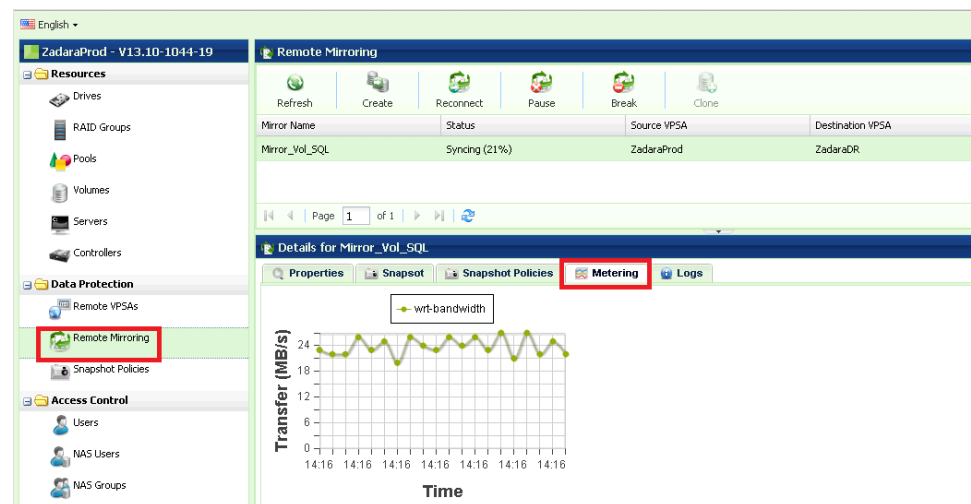
- **Create Policy** – N\A.

- **Delete Policy** – How many snapshots to retain 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 Chart provides live info of the Mirror transfer **throughput** associated with the selected Remote Mirror. You can view the Mirror metering info on 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 one of the following: 1 Second, 1 Minute, 10 Minutes, or 1 Hour. The **Auto** button lets you see continues updated live metering info (refreshed every 3 sec).



Logs

Displays all event logs associated with this Remote Mirror

11 Managing TechSupport Tickets

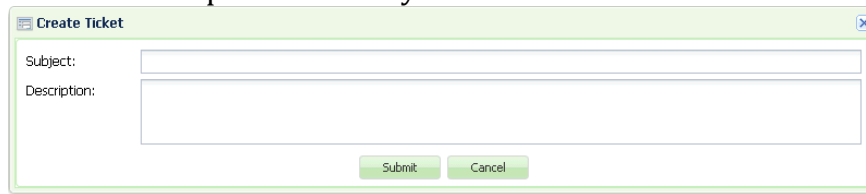
You can manage your Zadara TechSupport tickets directly from your VPSA. Any support request is redirected to the Zadara Support portal at <http://support.zadarastorage.com/home>.



To Open a Support Ticket

- Go to the Tickets page and click **Create**.
- Enter the Ticket Subject and description and press **Submit**.

- A ticket will be created and a set of logs (ZSnap) will be uploaded to the Zadara portal for analysis of the issue.



The image shows a 'Create Ticket' dialog box. It has a title bar with the text 'Create Ticket' and a close button (X). The main area contains two labels: 'Subject:' and 'Description:'. The 'Subject:' label is followed by a single-line text input field. The 'Description:' label is followed by a multi-line text area. At the bottom right of the dialog, there are two buttons: 'Submit' and 'Cancel'.

To Manage Support Tickets

- You can view the list of open support tickets, displaying the 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.