



Oracle ZFS Storage in OCI Quick Start Guide

Configuration of an Oracle ZFS Storage Instance in Oracle Cloud
Infrastructure (OCI)

September 2021 | Version 1.7
Copyright © 2021, Oracle and/or its affiliates
Public

PURPOSE STATEMENT

This document provides step-by-step instructions for configuring an Oracle ZFS Storage instance in OCI.

DISCLAIMER

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

TABLE OF CONTENTS

Purpose Statement	2
Disclaimer	2
Introduction	4
Limitations	4
Known Issues	5
Image Configuration Summary	5
Overview of Configuration Steps	5
First Steps	6
1. Import OCI Image	6
2. Configure OCI Compute Instance	8
3. Configure Block Storage	12
4. Configure Storage for ZFS Storage Instance	15
5. Share an SMB Filesystem	19
Backups	22
Best Practices	22
Network Best Practices	22
ZFS Storage in OCI Network Routing	22
ZFS Storage in OCI Network Datalinks	22
ZFS Storage in OCI Network Interfaces	22
Block Storage Best Practices	22
System Boot Disk	22
Storage Pools	22
APIs for ZFS Storage in OCI	23
Initial Configuration	23
Initial Setup	23
REST API	24
Additional Configuration	26

INTRODUCTION

Oracle is uniquely positioned to provide products and services that run 24/7 either on-premises or in the cloud and so, has the expertise to optimally run our own products in Oracle's own cloud.

Oracle ZFS Storage in OCI Marketplace provides cloud-based NAS storage and replication services enable on-premises ZFS Storage customers to migrate data and apps from on-premises to OCI. Oracle ZFS Storage instances provide both protocol services and performance for data migration, replication, and sharing.

The Oracle ZFS Storage image in OCI can be configured as a Bare Metal (BM) or Virtual Machine (VM) instance to support the following use cases:

- Migrate data to OCI over NFS, NFSv4, SMB or cross protocols with AD integration using an Oracle ZFS Storage BM or VM instance as a storage gateway
- Share data from ZS BM or VM in OCI over NFS, SMB, or cross protocols back to on-premise
- Replicate data to ZS BM or VM in OCI as a replication target and also reverse the replication back to on-premise
- Migrate and host applications workloads using similar protocols as your on-premise deployments

Sharing data and replicating data can be hosted in the following ways:

- Cloud to Cloud
- On-premise to Cloud
- Cloud to on-premise

After you have reviewed the above supported shapes, review the following summary of recommended shapes and recommended number of NFS and SMB clients.

Network Bandwidth Expectations for NFS/SMB Clients

Shape	Memory	Network Speeds	Maximum Client Bandwidth	Typical Sustained Bandwidth	Number of Clients
VM.Standard2.4	60GB	4.1 Gbps	256 MB/s	192 MB/s	Tens
VM.Standard2.8	120GB	8.2 Gbps	512 MB/s	384 MB/s	Hundred
VM.Standard2.16	240GB	16.4 Gbps	1025 MB/s	768 MB/s	Few Hundred
VM.Standard2.24	320GB	24.6 Gbps	1537 MB/s	1150 MB/s	Hundreds
BM.Standard2.52	768GB	25x2 Gbps	3125 MB/s	2343 MB/s	Thousands

Notes:

- Typical sustained workload mix with 50% read / 50% write.
- Number of clients depends on the desired throughput available to each client. If more throughput is needed per client then fewer clients should be used.
- A bare metal (BM) or virtual machine (VM) instance requires only one volume for operation. You can add more volumes to increase storage capacity for your needs.
- Maximum block volume capacity is 960TB.
- Detailed shape specifications are available at [OCI Shapes](#).

Limitations

The iSCSI protocol is only supported for boot volumes.

Known Issues

- Bare Metal shapes will generate a spurious network problem that can be ignored. The problem will be seen under the ZFS Storage in OCI Maintenance -> Problems tab and description will read: *The driver is suffering from a performance error detected in the driver. A(n) unsupported error has been detected during driver's attach context causing a(n) performance service.* (30773285 - OCI ZFS on BM network interfaces show errors on 2.52 shapes)
- Virtual Machine instances will show network devices speed as 1Gb even though it will use the full bandwidth allowed by the compute shape. (32749253 - VNICs speed is mentioned as 1G at CLI/BUI though VNIC effective bandwidth is more)
- If a new OCI VNIC is added to a running ZFS Storage in OCI VM, a reboot is required before the network device can be used. (32518670 - Adding an additional vnic to the OCI zfssa VM fails)
- The primary network interface used for iSCSI boot should not be modified. Use secondary network interfaces instead. (33001957 - Adding a secondary IP address to bnx0 panics the zfssa BM instance)

Image Configuration Summary

- Two images are available:
 - One is for a virtual machine (VM) instance, which currently is: PV_ZFSSA_8.8.34-1.2.34.4846.1x
 - One is for a bare metal (BM) instance, which currently is: BM_ZFSSA_8.8.34-1.2.34.4846.1x
- A boot volume and data volume are the minimum requirements for configuring a ZFS Storage instance
- A boot volume uses iSCSI protocol, but iSCSI is not otherwise supported

Overview of Configuration Steps

This guide describes the steps to configure Oracle ZFS Storage as a compute instance in Oracle's Cloud Infrastructure (OCI) and contains the following sections:

1. Import OCI Image
2. Configure OCI Compute Instance
3. Configure Block Storage
4. Configure ZFS Storage
5. Share an SMB Filesystem

For more Oracle® ZFS Storage Appliance documentation, go to https://docs.oracle.com/cd/F13758_01/

The final section provides additional management APIs used developed specifically for Oracle ZFS Storage in OCI version.

FIRST STEPS

1. The first step is to get an Oracle Cloud Infrastructure account.
<https://www.oracle.com/cloud/>

This guide assumes a usable compartment, virtual cloud network (VCN) and subnet has already been created and setup for use. An administrator for your OCI tenancy will authorize resources in a specified compartment for you to use.

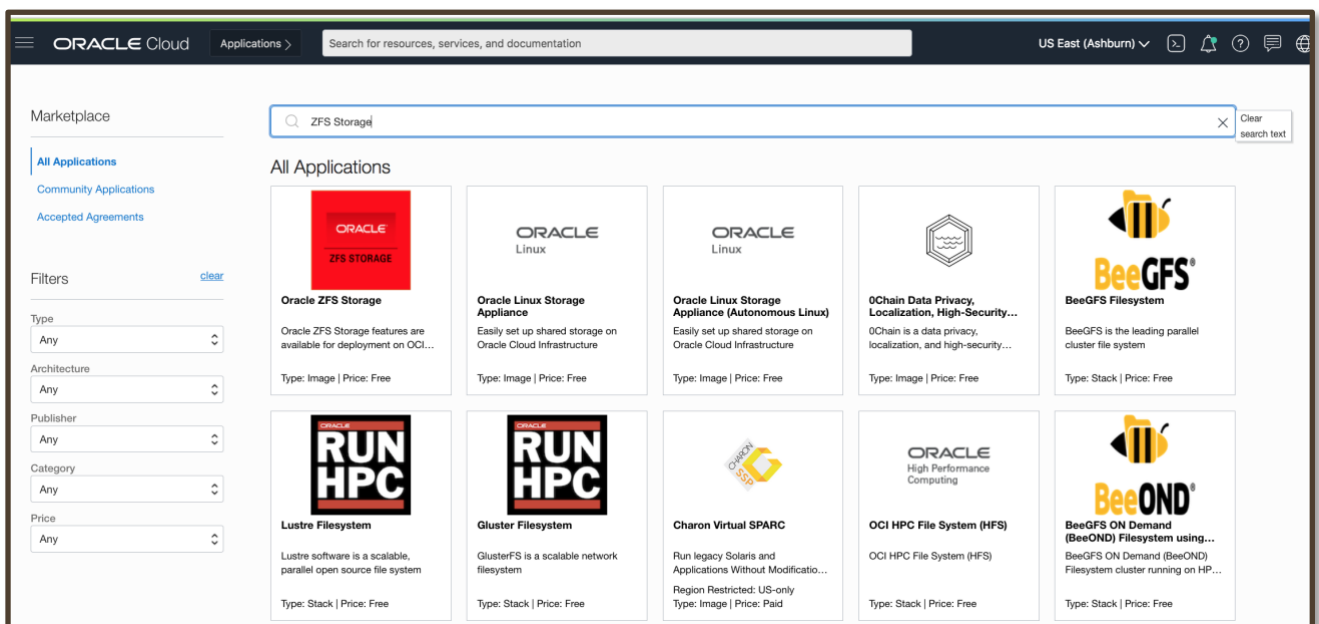
The following information will be needed to configure the OCI compute instance.

1. OCI Compartment ID
2. VCN Compartment and Name
3. Subnet Compartment and Name

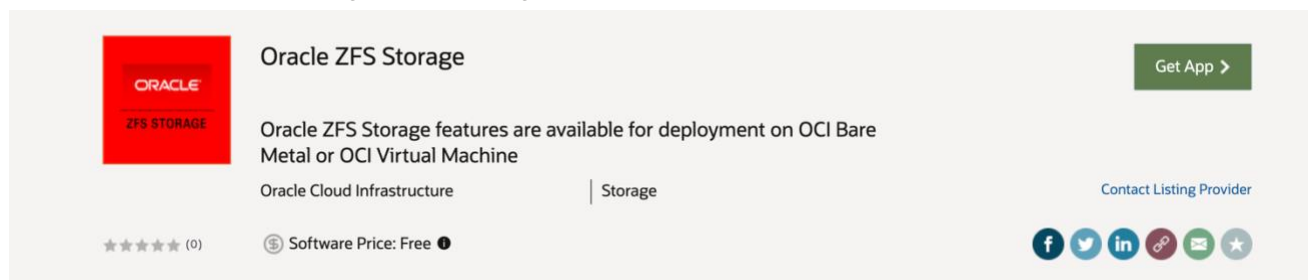
You will also need an ssh client to do the initial configuration and know how to configure the ssh client to use ssh key authentication.

1. IMPORT OCI IMAGE

1. Log into your OCI tenancy and region.
2. In the left hamburger menu, click on Marketplace.
3. In the main screen, click on Marketplace.
4. In the main search screen, enter ZFS Storage.
5. When the Oracle ZFS Storage image appears, click on the link.



6. From the Oracle ZFS Storage screen, click get app.



- The virtual machine (PV image) is selected by default. Select BM image if you are configuring a bare metal instance.

Review the Overview below, review and accept the terms and conditions. Then, click Launch Instance.

Marketplace » Oracle ZFS Storage



Oracle ZFS Storage

Oracle ZFS Storage features are available for deployment on OCI Bare Metal or OCI Virtual Machine instances.

Oracle ZFS Storage features are available for deployment on OCI Bare Metal or OCI Virtual Machine instances.

Categories: Storage

Type
Image

Version
PV_ZFSSA_8.8.34-1... ↕

Compartment

☒ I have reviewed and accept the [Oracle Terms of Use](#).

Launch Instance

Reminder: Patch the instance once installed.

Software Price per OCPU

\$0.00/hr

There are additional fees for the infrastructure usage. ⓘ

2. CONFIGURE OCI COMPUTE INSTANCE

1. Enter the name of the compute instance.
2. Click the Change Shape button. For example, to change to a virtual machine (VM) shape from a bare metal (BM) shape or vice versa.

The screenshot shows the 'Create Compute Instance' form. At the top, it says 'Create an instance to deploy and run applications, or save as a reusable Terraform stack for creating an instance with Resource Manager.' Below this, there's a 'Name' field with 'fishworks' entered. Under 'Create in compartment', 'store' is selected. The 'Placement' section is expanded, showing 'Availability domain' with three options: 'AD 1' (selected with a checkmark), 'AD 2', and 'AD 3'. Below these are links for 'Show advanced options'. The 'Image and shape' section is also expanded, showing an 'Image' field with a red 'ORACLE ZFS STORAGE' icon. To the right of the icon, it says 'Oracle ZFS Storage' and 'Oracle ZFS Storage features are available for deployment on OCI Bare Metal or OCI Virtual Machine'. A 'Return to Marketplace' button is visible on the right.

3. From the Browse All Shapes screen, select the Virtual Machine shape or the Bare Metal shape.
 - For a Bare Metal image, select BM.Standard2.52
 - For a Virtual Machine image, select VM.Standard2.4, VM.Standard2.8, VM.Standard2.16, or VM.Standard2.24

In the example below, the VM shape is selected.

4. Next, select the Intel shapes box and then the VM.Standard2.4 shape.

Browse All Shapes

A shape is a template that determines the number of CPUs, amount of memory, and other resources allocated to a newly created instance. See [Compute Shapes](#) for more information.

Instance type

Virtual Machine
A virtual machine is an independent computing environment that runs on top of physical bare metal hardware.

Bare Metal Machine
A bare metal compute instance gives you dedicated physical server access for highest performance and strong isolation.

Shape series

AMD
Flexible OCPU count, AMD processors.

Intel
Current generation Intel processors.

Specialty and Previous Generation
Earlier generation AMD and Intel Standard shapes. Always Free, Dense I/O, GPU, and HPC shapes.

	Shape Name	OCPU	Memory (GB)	Network Bandwidth (Gbps)	Max. Total VNICS	
<input type="checkbox"/>	VM.Optimized3.Flex		1	14	4	2
<input type="checkbox"/>	VM.Standard2.1		1	15	1	2
<input type="checkbox"/>	VM.Standard2.2		2	30	2	2
<input checked="" type="checkbox"/>	VM.Standard2.4		4	60	4.1	4
Local Disk: Block Storage Only						
<input type="checkbox"/>	VM.Standard2.8		8	120	8.2	8
<input type="checkbox"/>	VM.Standard2.16		16	240	16.4	16
<input type="checkbox"/>	VM.Standard2.24		24	320	24.6	24

1 Selected

Don't see the shape you want? [View your service limits and request an increase.](#)

Select Shape Cancel

- Click Select Shape.
- Configure networking settings for the instance.
Ask your OCI tenancy administrator what Network and Subnet to use.

Networking [Collapse](#)

[Networking](#) is how your instance connects to the internet and other resources in the Console. To make sure you can [connect to your instance](#), assign a public IP address to the instance.

Network

☒ Select existing virtual cloud network
☐ Create new virtual cloud network
☐ Enter subnet OCID

Virtual cloud network in [Redacted] [\(Change Compartment\)](#)

Subnet

☒ Select existing subnet
☐ Create new public subnet

Subnet in [Redacted] [\(Change Compartment\)](#)

Public IP Address

☐ Assign a public IPv4 address
☒ Do not assign a public IPv4 address
Requires a public subnet

[Show advanced options](#)

- Generate an SSH key and save the key, choose an existing public key file, or paste the contents of a public key. Generate or locate your ssh keys. Existing keys can be found in the your .ssh directory.

```
$ ls .ssh
```

```
config          id_rsa          id_rsa.pub      known_hosts
```

If you already generated keys for the opc user, locate opc/opc.pub.

11. Open a terminal and ssh to the ZFS Storage Appliance *Private IP Address* to set the `opc` user password to enable access to the BUI.

- The instance includes the `opc` user by default. This account provides limited user access. You can transition to a full administrative-capability root account once you have logged in as the **opc user** if you need full administrative access to the instance.
- Use your ssh credentials to log in as the `opc` user to your newly running instance.
- There could be a slight delay before you can ssh into the running instance.

For example:

```
ssh -i <path_to_private_key_file> opc@203.0.113.29
```

```
fishworks:> configuration users
fishworks:configuration users> select opc
fishworks:configuration users opc> set initial_password
Enter new initial_password: *****
Re-enter new initial_password: *****
Initial_password - (set) (uncommitted)
fishworks:configuration users opc> commit
fishworks:configuration users> exit
```

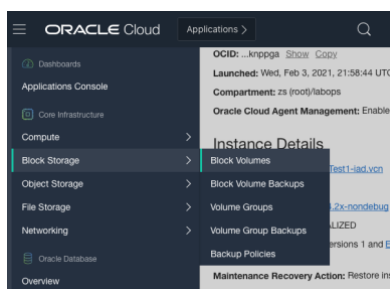
3. CONFIGURE BLOCK STORAGE

In this section, you will do the following steps:

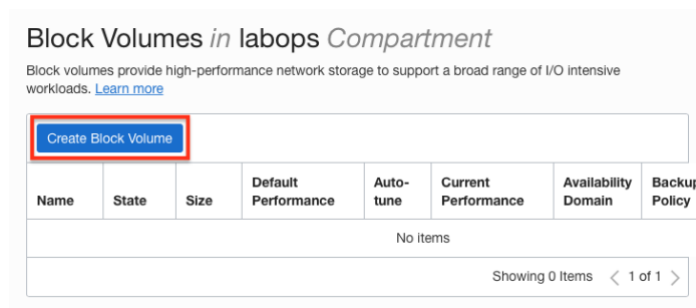
- Create block volume or volumes
- Attach block volumes to ZFS Storage instance
- Add iSCSI target

This section will create a block volume in OCI and attach it to a ZFS Storage Appliance. Once that is done a storage pool is created on the appliance.

1. Go to the *Block Volumes* page for your compartment.



2. Select *Create Block Volume*.



3. Set create block volume properties and select *Create Block Volume*.

Create Block Volume [Help](#)

Name
fishworks-disk01

Create in Compartment
zs (root)/labops

Availability Domain
IZbs:US-ASHBURN-AD-1

Volume Size and Performance
☒ Default ☐ Custom
Volume Size: 1024 GB
Volume Performance: Balanced
IOPS: 25000 IOPS (60 IOPS/GB)
Throughput: 480 MB/s (480 KB/s/GB)

Backup Policies
 Select Backup Policy in labops [\(Change Compartment\)](#)
 No Backup Policy Selected

Encryption
☒ **Encrypt using Oracle-managed keys**
 Leaves all encryption-related matters to Oracle.
☐ **Encrypt using customer-managed keys**
 Requires you to have access to a valid Key Management key.
[Show Tagging Options](#)

☒ View detail page after this block volume is created

[Create Block Volume](#) [Cancel](#)

4. Attach storage to compute instance. Go back to Compute Instances and select the ZFS storage appliance compute instance.

ORACLE Cloud

Applications >

Dashboards

Applications Console

Core Infrastructure

Compute

Block Storage

Object Storage

File Storage

Networking

Oracle Database

Overview

Autonomous Data Warehouse

Autonomous JSON Database

Autonomous Transaction Processing

Instances

Dedicated Virtual Machine Hosts

Instance Configurations

Instance Pools

Cluster Networks

Autoscaling Configurations

Custom Images

Boot Volumes

Boot Volume Backups

OS Management

fishworks-disk01

fishworks

Create Instance

Name	State	Public IP	Shape	OCPU Count	Memory (GB)	Availability do
fishworks	Running	-	VM.Standard2.2	2	30	AD-1

Showing 1 Item < 1 of 1 >

5. Scroll down to Resources and select *Attached Block Volumes* and then select *Attach Block Volume*.

The screenshot shows the 'Attached Block Volumes' page. On the left, under 'Resources', 'Attached Block Volumes' is selected. The main content area has a red box around the 'Attach Block Volume' button. Below the button is a table with columns: Name, State, Volume Type, Device path, Type, Access, Size, AD, and Created. The table is empty, with a message 'There are no block volumes attached to this instance.' and a pagination bar showing 'Showing 0 Items < 1 of 1 >'.

6. Set Attach Block Volume Properties.

- For VM and Bare Metal compute instances select *iSCSI*.
- Select the block volume from your compartment by volume name or volume OCID.
- Select Read/Write
- Click Attach

The screenshot shows the 'Attach Block Volume' configuration page. It includes sections for 'Volume attachment type' (with radio buttons for 'Let Oracle Cloud Infrastructure choose the best attachment type', 'iSCSI' (selected), and 'Paravirtualized'), a note 'This instance only supports iSCSI attachments. Learn more.', and a checkbox for 'Require CHAP Credentials'. The 'Volume' section has radio buttons for 'Select volume' and 'Enter volume OCID' (selected), a text input for 'Volume OCID' containing 'ocid1.volume.oc1', and performance metrics: 'Target Performance: Balanced', 'VPU: 10', 'IOPS: 25000 IOPS (60 IOPS/GB)', and 'Throughput: 480.00 MB/s (480 KB/s/GB)'. There is an optional 'Device path' field. The 'Access' section has radio buttons for 'Read/Write' (selected), 'Read/Write - Shareable', and 'Read Only - Shareable', each with a brief description.

Repeat step 1 through 6 to add more block volumes for data disks.

Note: If you want to expand existing block volumes, see the following MOS note (*How to Expand a Zpool on a ZFS Storage on OCI Marketplace Deployment*): [Oracle Support Document 2800583.1](#)

4. CONFIGURE STORAGE FOR ZFS STORAGE INSTANCE

1. Add iSCSI target to ZFS Storage in OCI.

After confirming that the iSCSI is done, select iSCSI Commands and Information.

Attached Block Volumes

[Block volumes](#) provide high-performance network storage to support a broad range of I/O intensive workloads.

Attach Block Volume

Name	State	Volume Type	Device path	Type	Access	Size	AD	Created
fishworks-disk01	Attached	Block Volume	-	iscsi	Read/Write	1 TB		

View Block Volume Details

ISCSI Commands & Information

Copy Attachment OCID

Copy Resource OCID

Detach

2. Get iSCSI IP Address and port and IQN.

iSCSI Commands & Information

Use OS tools to edit your /etc/fstab volume to have the _netdev and nofail options from the OS. Failure to run commands will cause instance boot failure.

Commands for connecting

```
sudo iscsiadm -m node -o new -T iqn.2015-12.com.oracleiaas:e1288ba0
sudo iscsiadm -m node -o update -T iqn.2015-12.com.oracleiaas:e1288ba0
sudo iscsiadm -m node -T iqn.2015-12.com.oracleiaas:e1288ba0
```

Commands for disconnecting

```
sudo iscsiadm -m node -T iqn.2015-12.com.oracleiaas:e1288ba0
sudo iscsiadm -m node -o delete -T iqn.2015-12.com.oracleiaas:e1288ba0
```

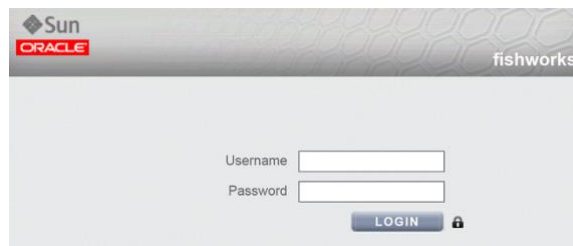
IP address and port: [redacted] [Copy](#)

Volume: iqn.2015-12.com.oracleiaas:e1288ba0 [redacted]

IQN: [redacted] [Copy](#)

[Close](#)

3. Open a web browser and navigate to the ZFS Storage URL. For example: <https://203.0.113.29:215>
This will be the IP address of the *Primary Private Address* on port 215.
Accept the web browser security warning. Then the login screen appears.



Sun ORACLE fishworks

Username

Password

LOGIN

Type in the *Username* opc and the *Password* created on the ssh terminal and click *LOGIN*.

This guide will use the BUI for the remaining configuration.

- Go the Oracle ZFS Storage Appliance BUI. Select *Maintenance* and then select *Workflows* and then select the arrow on *OCI iSCSI Target Attach* workflow to execute the workflow.



Oracle Public Cloud User@fishworks LOGOUT HELP

Configuration Maintenance Shares Status Analytics

HARDWARE SYSTEM PROBLEMS LOGS WORKFLOWS

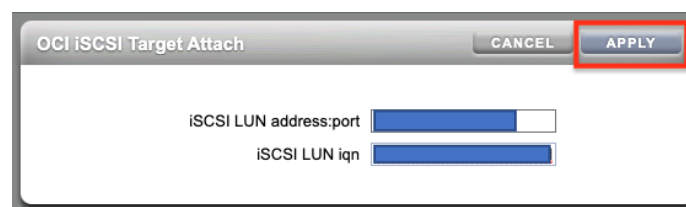
About Workflows

Workflows allow you to upload scripts to the appliance that can be executed in response to explicit user action, and be given explicit, validated parameters. Normally, a workflow executes as the user who executes the workflow, but workflows may be optionally executed as the user who owns the workflow, allowing for a mechanism for arbitrarily fine authorizations.

Workflows Total: 10

NAME	DESCRIPTION	VERSION
Cleanup DIMM Problems	Cleanup DIMM problems after part replacement	1.0
Clear locks	Clear locks held on behalf of an NFS client	1.0.0
Configure for Oracle Enterprise Manager Monitoring	Sets up environment to be monitored by Oracle Enterprise Manager	1.2
Configure for Oracle Solaris Cluster NFS	Sets up environment for Oracle Solaris Cluster NFS	1.0.0
Mitigation controls for CVE-2017-5015, CVE-2017-5754	Mitigation controls for CVE-2017-5015, CVE-2017-5754	1.0.0
OCI iSCSI Target Attach	Attaches an OCI iSCSI target to be used in storage pools	1.0
OCI iSCSI Target Detach	Detaches an OCI iSCSI target from the system	1.0
OCI iSCSI Target List	List all configured OCI iSCSI targets	1.0
Unconfigure Oracle Enterprise Manager Monitoring	Removes the artifacts from the appliance used by Oracle Enterprise Manager	1.0
Unconfigure Oracle Solaris Cluster NFS	Removes the artifacts from the appliance used by Oracle Solaris Cluster NFS	1.0.0

- Enter Address and IQN from the OCI iSCSI Commands & Information into the workflow and then click apply.



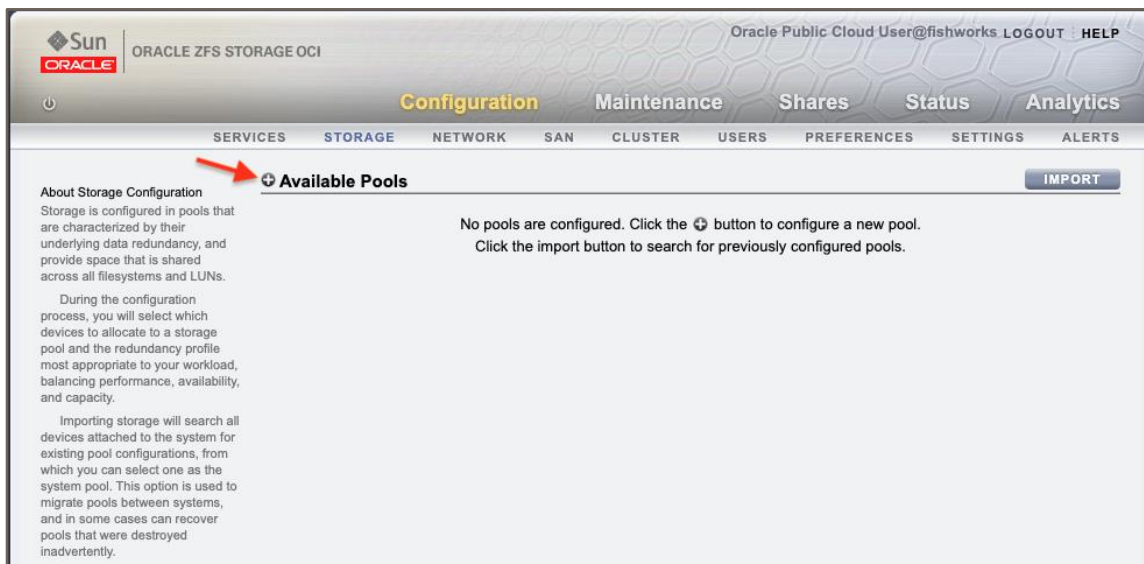
OCI iSCSI Target Attach

CANCEL APPLY

iSCSI LUN address:port

iSCSI LUN iqn

- Go to *Configuration Storage* and create a storage pool by clicking the  icon next to *Available Pools*.



7. Name the new storage pool and click APPLY.

CANCEL APPLY

Each storage pool is identified by a unique name. Please enter the name that will be used to identify this storage pool.

Pool name

8. Use the drop-down list next to *Data Devices* to select a disk to allocate to the pool, then click *COMMIT*.

Oracle Public Cloud User@fishworks LOGOUT HELP

Confirm that all devices are present and minimally functional, and allocate them to a storage pool.

ABORT COMMIT

Verify and allocate devices Step 1 of 2

Verify that storage is correctly attached and functioning. If devices are missing or malfunctioning, they will not be available for use and cannot be added without reconfiguring the pool. It is recommended that you fix any problems before configuring storage on the appliance. Mixing device types and speeds is strongly discouraged.

Model 0

Data Devices ✓ 1 (1T)

Log Devices -

Cache Devices -

9. Select COMMIT to accept the striped data profile.

ORACLE ZFS STORAGE OCI

Oracle Public Cloud User@fishworks LOGOUT HELP

Confirm that all devices are present and minimally functional, and allocate them to a storage pool.

ABORTCOMMIT

Choose Storage Profile

◀ Step 2 of 2 ▶

Configure available storage into a pool by defining its underlying redundancy profile. Carefully read the profile descriptions to understand how each balances the inherent trade-offs between availability, performance, and capacity, and select the profile that best fits your workload. If available, NSPF indicates no single point of failure, which affords certain profiles the ability for a pool to survive through loss of a single disk shelf.

Storage Breakdown

Data

1008G

Reserve

16G

Data profile: Striped

Data is distributed evenly across all disks without redundancy, maximizing performance and capacity, but providing no protection from disk failure whatsoever. Striping is recommended only for workloads in which data loss is an acceptable trade off for marginal gains in throughput and storage space.

10. The *Configure Storage* page is redisplayed with a description of the new pool.

ORACLE ZFS STORAGE OCI

Oracle Public Cloud User@fishworks LOGOUT HELP

ConfigurationMaintenanceSharesStatusAnalytics

SERVICESSTORAGENETWORKSANCLUSTERUSERSPREFERENCESSETTINGSalERTS

About Storage Configuration

Storage is configured in pools that are characterized by their underlying data redundancy, and provide space that is shared across all filesystems and LUNs.

During the configuration process, you will select which devices to allocate to a storage pool and the redundancy profile most appropriate to your workload, balancing performance, availability, and capacity.

Importing storage will search all devices attached to the system for existing pool configurations, from which you can select one as the system pool. This option is used to migrate pools between systems, and in some cases can recover pools that were destroyed inadvertently.

Available Pools

IMPORT

NAME	DATA PROFILE	LOG PROFILE	STATUS	ERRORS	ENCRYPTED
demo	Striped	-	Online	0	

demo

ADDREMOVEUNCONFIG

REVERTAPPLY

Pool Name

demo

Data Profile

Striped

Log Profile

-

Pool Status

Online

Data Errors

No known persistent errors

Scrub Schedule

30 days

Scrub Status

Never scrubbed

SCRUB

Allocation

Data

1000G

Data + Reserve

1 disks

Spare

0 disks

Log

0 disks

Cache

0 disks

Meta

0 disks

Device Status

0 errors

No device faults have been detected in the storage pool.

5. SHARE AN SMB FILESYSTEM

Complete the following steps to set up a simple filesystem share over Server Message Block (SMB) with Windows user access.

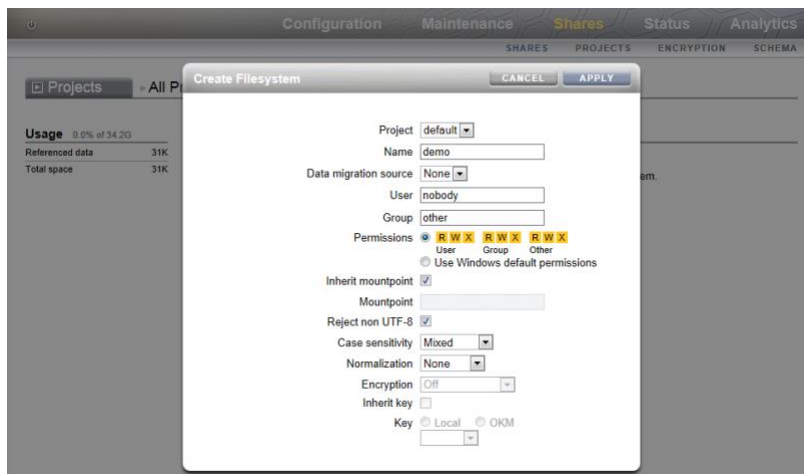
1. Navigate to the Shares screen.

Click the add item icon  next to Filesystems to create a new filesystem.



2. Name the filesystem and change the permissions for Group and Other to allow anyone to read, write, and execute on the filesystem.


In this example, the filesystem is named demo. The filesystem is part of the default project. Click APPLY to save the changes.



3. In the Shares screen, mouse over the entry for the new filesystem and click the edit icon to edit the filesystem attributes.

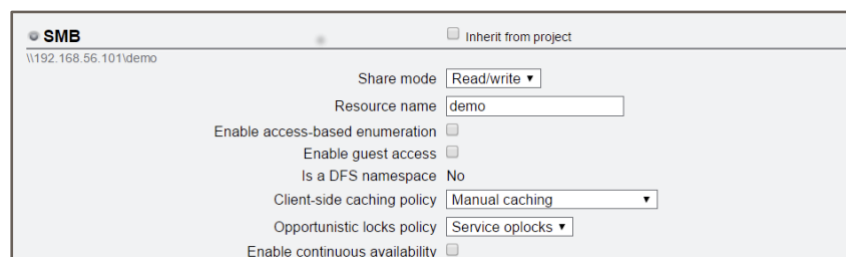


4. Click Protocols.

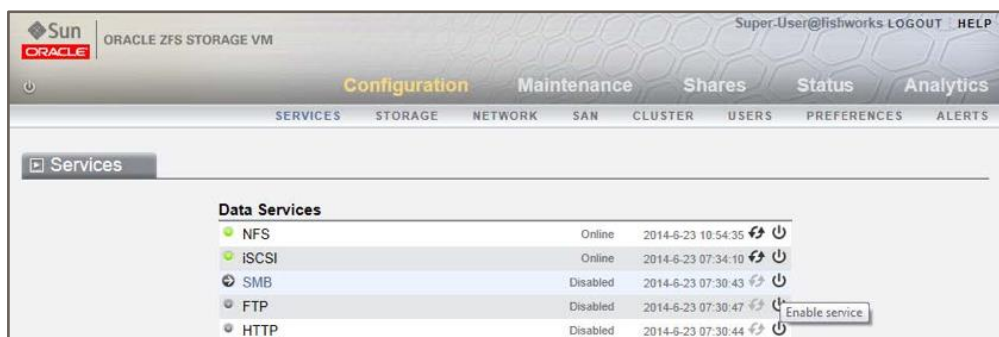


5. In the SMB section, clear the checkbox for Inherit from project, select Read/Write in the Share mode drop-down list, and set the Resource Name.


In this example, the Resource Name is demo. Click APPLY to save the changes.



6. Click Configuration to access the Configuration Services screen.
7. Enable the SMB service by clicking the power icon.



The state will change from Disabled to Online.

8. Configure a user with access to the filesystem share.
 - a. Click USERS in the navigation bar, and click the add item icon  next to Users to create a new user.
 - b. Select Local Only, set the Username and Password, and click ADD. Log out of the BUI by clicking LOGOUT near the top of the screen.

Add User [CANCEL] [ADD]

Properties

Type: ☐ Directory ☒ Local Only

Username:

Full Name:

Password:

Confirm:

Require session annotation: ☐

Kiosk user: ☐

Kiosk screen:

Roles | Exceptions

1 Total

NAME	DESCRIPTION
<input checked="" type="checkbox"/> basic	Basic administration

- From a Windows client, connect to the IP address of your ZFS Storage instance, and log in with the credentials you set in step 8 to access the shared filesystem.

BLOCK VOLUME BACKUPS

OCI Block Volume Service allows you to create snapshots of both boot volume and block volumes.

- Boot Volume snapshots
 - <https://docs.oracle.com/en-us/iaas/Content/Block/Tasks/backingupbootvolume.htm>
- Block Volume Backups
 - <https://docs.oracle.com/en-us/iaas/Content/Block/Concepts/blockvolumebackups.htm>

BEST PRACTICES

Network Best Practices

- The primary OCI VNIC should be used for iSCSI traffic.
- A secondary OCI VNIC should be created for NAS traffic.
- A secondary IP on the secondary OCI VNIC should be used for NAS traffic. Since the secondary IP address can be moved, it allows for easier migration of NAS traffic to different ZFS Storage in OCI instances.
- For information on configuring a secondary OCI VNIC, see [Managing VNICs](#).

ZFS Storage in OCI Network Routing

- It is recommended to set the multihoming model to strict.
- Create a default IPv4 route on the primary network interface with the destination set to 169.254.0.0/16 for iSCSI traffic to increase network throughput.

ZFS Storage in OCI Network Datalinks

- Link Speed, Link Duplex and Flow Control should all be set to Auto.
- Link speed for VM instances will be reported as 1GB but will actually use the full amount of bandwidth allocated to the instance. (See known issues)
- All network datalinks should have the MTU set to 9000 for best performance.

ZFS Storage in OCI Network Interfaces

- The primary network interface used for iSCSI traffic should not be modified because it can cause a system panic. (See known issues)
- Consider using separate subnets for storage administrators and NAS clients for enhanced security.
- NAS client interfaces should uncheck 'Allow Administration' for enhanced security.

Block Storage Best Practices

System Boot Disk

- System disk contains read only OS image, logs, core dumps and configurations.
- Configuration data can be backed up using 'Maintenance System Configs'
- Does not include OS image, logs, core dumps, replication or share data.
- Logs and core dumps can be saved using 'Maintenance System Bundles'
- Entire system disk can be backed up using OCI boot volume backups.

Storage Pools

- Pool disks contain all configuration data under 'Shares'
- All disks in each pool should be same size especially if they are under 800GB.
- All data disks in each pool should have the same performance settings.
- Suggest creating a volume group containing all data disks for each storage pool.
- Block volume backups must use volume groups to keep pool data consistent.
- For best system resource usage recommend only one pool per VM.
- All data disks provided by OCI have multiple copies so striped pools provide data protection. ZFS will detect bit rot but data will have to be restored from backup if bit rot is detected.
- Consider backing up data disks or using a parity or mirrored storage profile to protect against bit rot or a block volume outage.

Backup of ZFS Configuration

We recommend that after your ZFS Storage in OCI instance is configured, that you create a backup of the configuration with the following steps:

- From the Appliance BUI, go to Maintenance→System.
- Under the Configurations section, click Backup.
- This will create a backup of the Appliance configuration, that can be downloaded and stored separately for recover purposes.

For information about the configuration backup content, what is included and what is not included, see [Backing Up the Configuration](#).

SECURITY REFERENCES

For information about setting permissions on shares and recommended security practices, see the following references:

- [Access Control Lists for Filesystems](#)
- [Oracle® ZFS Storage Appliance Security Guide, Release OS8.8.x](#)

APIS FOR ZFS STORAGE IN OCI

Initial Configuration

The on-premise ZFS Storage Appliance uses a manual based initial configuration over the serial console that sets up the initial network settings and the root password. Console configuration is still supported if the ssh keys are not supplied to the instance during launch. A serial console connection can be made by using the *Console Connection* on the OCI compute instance page.

If user SSH keys are defined for the instance then they will be automatically applied to the opc and root users during the initial configuration. The password will be set to a long random value to disable password based access. Initial access will only be available via ssh until the password is set.

The initial install will try and use DHCP settings to set the initial configuration properties usually populated from the console. If DHCP does not provide all the values then a *Console Connection* still needs to be made to do the initial install.

Initial Setup

Initial setup for on-premise ZFS Storage Appliance is a guided manual setup that configures storage, DNS, naming services, NTP, network and phone home.

OCI compute instances have a special metadata called "user_data" field that can be used for configuring the compute instance on initial boot. If this field is applied then the ZFS Storage in OCI will run the configuration script. See [Working with CLI Scripting](#) in the [Oracle® ZFS Storage Appliance Administration Guide](#)

ZFS Storage in OCI will also look for "config_data" metadata and if it exists it will be used instead of "user_data". The reason for this is OCI does not allow "user_data" to be modified, and "config_data" can be modified and deleting "config_data" after the system has been configured can help stop leaking of configuration information through OCI compute metadata.

The install script can define or override initial configuration properties that are usually populated by DHCP to guarantee that manual configuration via the *Console Connection* is not needed.

DNS is usually configured during initial configuration but the script can add additional DNS server. NTP is auto configured for the OCI environment and enabled by default so it does not need to be configured. If more than one NIC is used it is recommended to change the routing multihoming to strict.

Example User Data Script to configure a system. This script sets up initial configuration and also set up static routing for iSCSI. Note: Since OCI only allows one user script the configuration properties are defined as comments so that when the CLI script is ran they will be ignored.

```
#@hostname=fishworks
#@domain=example.com
#@ip_addr=203.0.113.29
```

```
#@ip_mask=255.255.255.0
#@router=203.0.113.29
#@dns_servers=203.0.113.125

script

print('setting up DNS...');
run('top configuration services dns');
run('create');
set('address', '203.0.113.126');
run('commit');

run('top configuration net routing');
prop('multihoming', 'strict');
run('commit');
.
```

REST API

The ZFS Storage Appliance (ZS) has a REST API for managing all aspects of an on-premise ZS but some additional functionality is required for integrate with cloud orchestration services. This REST API expands the ZS REST API to manage OCI resources that are not available from the on-premise ZS REST API.

The new REST API endpoint will be available from the `/api/oci/v2` resource. This resource path will be hidden and its only supported use will be by Oracle IT within OCI. BUI and CLI interfaces for these resources will not be made available since their purpose is for use with cloud orchestration software.

Instance Information

OCI compute instances are able to query OCI to get information about itself. The instance endpoint simply returns the instance information returned from the OCI service: <http://169.254.169.254/opc/v1/instance/>

The content is defined by the OCI instance information service defined at <https://docs.cloud.oracle.com/iaas/Content/Compute/Tasks/gettingmetadata.htm>

The purpose of supplying instance information is to enable cloud monitoring and orchestration use cases.

Example Request

```
GET /api/oci/v2/instance
```

Example Response

```
{
  "instance": {
    "compartmentId":
"ocidl1.compartmen.oc1..aaaaaaaadrcsphnhoji76iyemczsoh2ggmokaaka7fxzm772hrjpxqfne
slq",
    "displayName": "fishworks",
    "timeCreated": 1573235633999,
    "agentConfig": {
      "managementDisabled": true,
      "monitoringDisabled": true
    },
    "canonicalRegionName": "us-phoenix-1",
```



```

    "shape": "BMStandard2.52",
    "availabilityDomain": "iZbs:PHX-AD-1",
    "id":
"ocid1.instance.oc1.phx.anyhqljsuna655qccpnt2hmojteyoljyh43fvfalzuljgurzogxrkj5tc
ghq",
    "state": "Running",
    "definedTags": { },
    "region": "phx",
    "faultDomain": "FAULT-DOMAIN-1",
    "image":
"ocid1.image.oc1.phx.aaaaaaaaxqr2he5lcipvthyxru2l1l1stb3jzgfg6gyqx17xb45qt6evn2lwa
",
  }
}

```

iSCSI Targets

A hardware based ZS does not act as an iSCSI initiator and has no API to manage the available iSCSI targets. The `iscsi-targets` resource allows clients to manage the `iscsi-targets` attached to the ZS.

Important! Dual attachment of LUNs to multiple ZS instances is not supported. All iSCSI LUNs should be attached to only one system at a time. Failure to do so will result in data corruption!

The purpose of this resource is to enable initial setup and migration of iSCSI LUNs between ZS instances.

iSCSI Target Resource Properties

PROPERTY	TYPE	DESCRIPTION
addr	string	iSCSI target address:port
iqn	string	iSCSI target iqn
pool	immutable string	The pool name if the LUN associated with the iSCSI target is part of a zfs pool.

List iSCSI Targets

The `get` command will return the currently attached iSCSI LUNs available for use as storage pools on the ZS. The available number of iSCSI targets should match the number of storage devices available for pools. If a device is configured as part of a storage pool, the pool name will be provided.

Example Request

```
GET /api/oci/v2/iscsi-targets
```

Example Response

```

{
  "targets": [{
    "iqn": "iqn.2015-12.com.oracleiaas:48558e46-6cf0-4b27-b8c2-f8bd69812305",
    "addr": "169.254.2.5:3260",
    "pool": "p1"
  }, {

```

```

    "iqn": "iqn.2015-12.com.oracleiaas:bff0f183-736e-49c3-9478-b1352407262f",
    "addr": "169.254.2.6:3260",
    "pool": "p1"
  }
}

```

Create iSCSI Targets

This command will attach all the specified OCI iSCSI targets to the ZS. The most efficient method to bring all storage online is to attach all LUNs using a single command. When the command returns the LUNs associated with the iscsi targets should be available on the ZS system. To import or clear any pools associated with the new devices, use the pool resource documented below.

Generate a list of all volumes attached to a compute instance that can be used as the POST data for the command, run the following:

```

INSTANCE_ID=ocidl.instance.oc1.phx.anyhqljruav6g7ic63hz2jkm5ywxxyz3y23xh
oci compute volume-attachment list --instance-id $INSTANCE_ID \
  --query 'data [*].{iqn:"iqn",addr:"ipv4"}'

```

Example Request

```

POST /api/oci/v2/iscsi-targets
[
  {
    "iqn": "iqn.2015-12.com.oracleiaas:48558e46-6cf0-4b27-b8c2-f8bd69812305",
    "addr": "169.254.2.5:3260"
  }, {
    "iqn": "iqn.2015-12.com.oracleiaas:bff0f183-736e-49c3-9478-b1352407262f",
    "addr": "169.254.2.6:3260"
  }, {
    "iqn": "iqn.2015-12.com.oracleiaas:6a9b6b07-aaf3-44ce-82ba-b2ae59b11d4e",
    "addr": "169.254.2.7:3260"
  }
]

```

The response will be a JSON object containing a "target" property whose value is the list of iscsi-targets that were created.

Remove a single iSCSI Target

The delete command on a specified IQN resource will only remove the specified iSCSI LUN. This is useful for modifying an existing system.

```
DELETE /api/oci/v2/iscsi-targets/<iqn>
```

Remove All iSCSI Targets

The delete command will remove all iSCSI LUNs from the system. Any LUN that is part of a pool will not be deleted.

Example Request

```
DELETE /api/oci/v2/iscsi-targets
```

Additional Configuration

Apply a configuration from a mirrored system disk. This will have the same effect as saving a system configuration backup and applying it to the system.

Property	Type	Description
addr	string	iSCSI addr of target containing cloned boot disk
iqn	string	iSCSI iqn of target containing cloned boot disk
locked	boolean	Flag determining if system configuration is locked. When locked a warning will be given on login that system configuration changes will not be propagated to the destination system. Storage pools will also not be imported if the system reboots to avoid dual import when using shared disks.
version	string	system version to import (optional, default is to use latest version on system disk)
guid	string	System pool guid (optional, default will be to use the first importable system disk guid found)
status	immutable string	One of STARTING, RUNNING, FAILED, COMPLETED showing full job status.
config_status	immutable string	Set to COMPLETED or FAILED once the system config is ready for client I/O. Some background tasks may still be running such as collecting lock data for phone-home support.
started	immutable string	job start time in iso8601 formatted UTC date
finished	immutable string	job end time in iso8601 formatted UTC date
message	immutable string	Human readable status and error messages for debugging.

Lock Configuration

Lock configuration on original source compute instance.

```
PUT /api/oci/v2/config
{
  "locked":true
}
```

Start Configuration Import Job

Example command to start a configuration import job. The addr and iqn properties are the values of a iscsi settings of the system disk attached to the ZS.

The configuration from the old ZS attached system disk will be applied to the new ZS system.

```
PUT /api/oci/v2/config
{
  "addr": "169.254.2.18:3260",
  "iqn": "iqn.2015-12.com.oracleiaas:36bcc330-c1a0-4494-ad66-76c88cb1d044"
}
```

Get Configuration Job Properties

Get details of a running config job.

```
GET /api/oci/v2/config
```

```
{
  "config": {
    "addr": "169.254.2.18:3260",
    "iqn": "iqn.2015-12.com.oracleiaas:36bcc330-cla0-4494-ad66-76c88cb1d044",
    "status": "COMPLETED",
    "started": "2020-09-25T16:50:27",
    "config_status": "COMPLETED",
    "finished": "2020-09-25T16:51:50",
    "version": "ak-nas-2013.06.05.8.23.3_13.7-2",
    "message": "2020-09-25T16:51:50 Restored system configuration\n2020-09-25T16:51:50 Waiting for background logs to be copied\n2020-09-25T16:51:50 completed",
    "guid": "17894726420108290233"
  }
}
```

CONNECT WITH US

Call +1.800.ORACLE1 or visit oracle.com.
Outside North America, find your local office at oracle.com/contact.



blogs.oracle.com



facebook.com/oracle



twitter.com/oracle

Copyright © 2021, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

SOLUTION BRIEF

September 2021

Author: [zfssa-storage-feedback_ww_grp@oracle.com](#)

