

Configuring ZFS Storage for Cloud Snapshot Backups to OCI Object Storage

Configuration Best Practices

May, 2024 | Version 5.1
Copyright © 2024, Oracle and/or its affiliates
Public



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

Disclaimer	2
Introduction	4
Cloud Snapshot Backups	4
Use Cases	4
Advantages	4
Overview	5
Features of Cloud Snapshot Backups and OCI Object Storage	5
Cloud Snapshot Backup Formats (ZFS or Tar)	5
OCI Object Storage Retention Policy Summary	6
OCI Archive Object Storage Support	6
Preparing OCI Object Storage for ZFS Storage Cloud Backups	7
Summary Steps for Preparing OCI Object Storage for ZFS Storage Cloud Backups	7
Oracle Cloud Infrastructure (OCI) Account Setup	7
Creating a Bucket in OCI	9
Oracle ZFS Storage – Configure Cloud Service	11
Creating the Target	12
Target Information	12
Routing to the Target	14
Assigning Authorizations	15
Cloud Snapshot Backup	16
Creating the Initial Snapshot	16
Backup the Initial Snapshot to OCI	17
Backup an Incremental Snapshot to OCI	18
Creating a Tar-format Backup in OCI	20
Cloud Snapshot Restore	20
Recover a Cloud Snapshot to ZFS Storage	20
Recover a Tar-format Cloud Snapshot to a Linux or Solaris System	22
Using OCI Object Retention Policies With ZFS Storage Cloud Backups	24
Monitor Cloud Snapshot Backups and Object Store Bucket	26
Cloud Snapshot Backups Roles and Authorization	29
Customization Tips	31
OCI Command Line Simplification	31

INTRODUCTION

Oracle ZFS Storage, both the on-premises appliance and the virtualized ZFS-HA appliance in OCI, provides the ability to back up ZFS snapshots to Oracle Cloud Infrastructure (OCI) object storage. While this object storage can be provided by either OCI or another on-premises Oracle ZFS Storage Appliance, this paper will focus on the use of OCI object storage. For the on-premises solution, see the paper “Configuring ZFS Storage as an Object Storage Target for Cloud Snapshot Backups”.

Object storage is the primary method of storing large amounts of data in the cloud. With a flat namespace, it scales better than file system storage and supports extended metadata attributes that allow comprehensive search operations. Billions of objects can be stored when needed.

Cloud Snapshot Backups

A snapshot in Oracle ZFS Storage provides an immutable “point-in-time” view of a ZFS project or share. These snapshots can be mounted for read-only access or cloned. Clones may be mounted with read/write access without affecting the original snapshot.

The cloud snapshot backup feature of Oracle ZFS Storage allows you to back up full and incremental snapshots from Oracle ZFS Storage to OCI object storage targets. The selection of snapshots to be backed up to OCI targets is under the administrator’s control.

Use Cases

- Provides low-cost off-site storage for snapshot backups
 - Snapshot backups can be scheduled through workflows
 - Provides recovery by restoring snapshots and rolling back
 - Provides recovery from accidental overwrite or malware like ransomware
 - Supports less-expensive Archive storage for backup objects
 - Provides off-site archive storage for long-term data, such as business compliance requirements. OCI objects can optionally have retention rules set, ensuring that the objects cannot be removed accidentally or by a bad actor
- Provides a straightforward way to migrate data from an on-premises Oracle ZFS Storage Appliance to OCI storage

Advantages

Review the following advantages for storing or migrating data in a hybrid cloud environment with the ZFS Storage Appliance:

- Provides native cloud integration with OCI object storage
- Cloud snapshot backups are fully supported in the ZFS Appliance BUI, CLI, and REST interfaces and integrated with analytics, alerts, logs, and authorization roles
- File system snapshots can be backed up to another on-premises ZFS Appliance or directly to OCI object storage
- Cloud snapshot backups using the ZFS format can be restored back to either an on-premises ZFS Storage Appliance or a ZFS-HA cluster running in OCI
- Cloud snapshot backups using the TAR format can be restored back to ZFS systems or to any system able to process .tar files
- Snapshot backups can be encrypted and compressed and if required, replicated to worldwide on-premises data centers or up to OCI cloud storage

OVERVIEW

This document details the features and benefits of using OCI object storage to back up Oracle ZFS Storage snapshots. It also details the following processes:

- Configuring a ZFS Storage node as a cloud snapshot backup source system
- Configuring an OCI cloud target for archiving cloud snapshot backups
- Sending a ZFS cloud snapshot backup to OCI object storage
- Restoring a ZFS cloud snapshot from OCI object storage

FEATURES OF CLOUD SNAPSHOT BACKUPS AND OCI OBJECT STORAGE

Cloud Snapshot Backup Formats (ZFS or Tar)

ZFS Storage Appliance cloud snapshot backups support both ZFS format and tar format. The ZFS format can be used when restoring cloud snapshot backups back to ZFS Storage, either on-premises or in OCI. The tar format provides more data recovery flexibility because you can restore back to any system that can process tar files, usually Solaris or Linux systems. Both ZFS and tar format can be backed up to the same object storage bucket. ZFS will automatically select the proper format upon restoration back to ZFS Storage.

The ZFS format is block based. Incremental backups will only transfer the blocks that have changed since the last snapshot. Compression applied at the share level will be reflected in the blocks sent to OCI object storage.

The tar format is file based. A change of a single byte in a large file will cause the entire file to be sent in an incremental snapshot. Tar-formatted cloud snapshots do not reflect any compression that may be applied to the ZFS share but will be compressed within the snapshot.

The following table summarizes the features of the ZFS and tar formats.

ZFS or Tar Format Feature Comparison

Feature Description	ZFS Format	Tar Format
Restore on any system?	ZFS Storage Appliance, ZFS-HA cluster, or Oracle Solaris server	Any system regardless of OS
Supports both filesystem and LUN snapshots?	Yes	Filesystem snapshots
Preserves filesystem or LUN properties?	Yes	Filesystem properties
Supports full and incremental backups	Yes	Yes
Supported within the same Oracle Cloud Infrastructure bucket?	Yes	Yes
Supports high-efficiency compression? • If underlying share is compressed, less data is transferred and the backup is faster. • Incremental backups are performed at the block level.	Yes	No
Supports mid-efficiency compression? • Files are read and compressed during the backup operation. • Incremental backups are performed at the file level. Especially not efficient if a large file is modified because the full file will be part of the incremental backup.	No	Yes

OCI Object Storage Retention Policy Summary

OCI object storage retention policies offer administrators the ability to meet regulatory requirements for records management and retention. For an independent assessment of the OCI object storage retention rules feature's ability to meet regulatory requirements for record management and retention, see Cohasset Associate's [SEC 17a-4\(f\)](#), [FINRA 4511\(c\)](#), [CFTC 1.31\(c\)-\(d\)](#) and [MiFID II Compliance Assessment](#).

Different types of retention policies are available to cover a wide range of use cases.

POLICY	USE CASE	COMPLIANCE TYPE
Time-bound – specify a time duration	- Your industry might require data retention for a defined time duration - Data retention regulations might also require locked retention settings - Locked retention settings mean, the only change allowed is to increase the retention duration	Regulatory
Time-bound – specify a time duration that might change	- Your internal business requirements might require data retention - While data retention is required, that time period could change	Data Governance
Indefinite – object modification is prevented until retention rule is removed	- Requirement is to preserve certain business data in response to potential or on-going litigation - A legal hold does not have a defined retention period and remains in effect until removed	Legal Hold
Retention lock	Your company data retention regulations might also require that you lock the retention settings	Data Governance

OCI Archive Object Storage Support

OCI object storage supports different tiers of object storage, including the inexpensive archive object storage. The ZFS Storage cloud backup supports writing to both the standard and archive object tier.

The use of Object Lifecycle Policies or auto-tiering to Infrequent Access object storage in OCI is not supported by the cloud backup feature. If less expensive object storage is desired, use an archive bucket for the backup's data target.

When using archive storage, separate buckets must be used for metadata and the backup objects. This is detailed in the section [Using OCI Object Retention Policies With ZFS Storage Cloud Backups](#).

PREPARING OCI OBJECT STORAGE FOR ZFS STORAGE CLOUD BACKUPS

Summary Steps for Preparing OCI Object Storage for ZFS Storage Cloud Backups

- Confirm your access to OCI Object Storage in your corporate OCI tenancy or create a free OCI cloud storage trial at <https://www.oracle.com/cloud/free/>
- If using a free trial account:
 - Sign into your account after it is activated
 - Note the limitations of the object storage trial at the top of the screen
- Collect information needed to define a Cloud Backup target in ZFS Storage
 - Review user profile and tenancy information and locate OCIDs
 - Generate or upload API key
- Create a bucket to be used as a backup target – see [Oracle Cloud Infrastructure \(OCI\) Object Storage Setup](#)
 - Select Storage then Object Storage from the “hamburger” menu
 - Select the (root) storage compartment that is created automatically in your tenancy
 - Create a bucket with auto-generated name or give it a new name if you prefer
 - Select tier
 - Select encryption key manager
 - Determine the URL for the bucket’s location

Oracle Cloud Infrastructure (OCI) Account Setup

OCI provides a low-cost solution for data archiving that integrates well with ZFS Storage. OCI object storage provides standard, infrequent access, and archive tiers. When stored in any of these storage tiers, data is encrypted automatically. This cannot be disabled. Multiple cloud regions are available, and data can be migrated between geographic regions.

Sign up for a free 30-day cloud trial to start an on-premises data migration to OCI object storage:


<https://www.oracle.com/cloud/free/>

Collecting the OCIDs

To configure Oracle ZFS Storage for cloud backups, it is necessary note the Oracle Cloud Identifier (OCID) of a few different OCI resources.

User OCID and API Keys


The OCI Object Storage User profile screen includes user certificates that are created when the account is created. The OCID of the user is required when the cloud bucket target is created on Oracle ZFS Storage. Navigate to the Identity & Security screen in the OCI console, select Domains, then choose the domain the user is a part of – often the OracleIdentityCloudService domain but possibly the Default domain. Select Users, then choose the OCI user that will own the OCI buckets for the backups. You can choose to view the user OCID or copy it to the local clipboard.




(Note that tenancies that have not been updated to use Identity Domains may have a different method of accessing the user profiles.)

An API public key must be generated or copied and associated with the user. This key is required to modify and view objects in Oracle ZFS Storage and to grant the user access through CLI if desired.

Scroll down in the OCI Object Storage User profile screen until the Resources menu is shown on the left side of the screen. Choose the “API keys” item, then click on the “Add API key” button.




The next screen will generate an API key pair automatically, but advanced users may wish to upload or paste a public key from an already-generated pair. If key is generated here, the user **MUST** download the private key at this point – there is no chance to recover it after this point. Without this public key, there is no way to use this key pair for authentication.



Tenancy OCID


Each OCI tenancy is assigned one unique and uneditable object storage namespace that spans all compartments within a region. Navigate to the top page in the OCI console. To the right will be an identifier labeled “Your tenancy”. Click

on the tenancy name displayed. The OCI Tenancy detail screen is then shown. You can choose to view the tenancy OCID or copy it to the local clipboard.




Creating a Bucket in OCI

OCI provides a management interface into all cloud components, including object storage. After logging in to OCI, click the “hamburger” menu in the upper left corner and choose storage.



On the Storage screen, select “Buckets” under “Object Storage and Archive Storage”.



The OCI object storage interface provides a management view of the object storage tenancy and current compartment. Buckets are created to store data objects and can be reviewed accordingly. More options are available to provide bucket details.

ORACLE Cloud

Applications >

Search for resources, services and documentation

US East (Ashburn) ▾

🔍 🔔 ? 💬 🌐 👤

Object Storage

Object Storage

Data Transfer - Import

Data Transfer - Export

List Scope

COMPARTMENT

▾

Tag Filters

add | clear

no tag filters applied

Buckets in Compartment

Create Bucket

Name ▲	Storage Tier	Visibility	Created
bucket-cs2	Standard	Private	Wed, Aug 14, 2019, 14:25:30 UTC ⋮
bucket-new-cs	Standard	Private	Mon, May 6, 2019, 21:03:03 UTC ⋮
bucket-ooow2019	Standard	Private	Tue, Jul 9, 2019, 22:19:37 UTC ⋮
demobucket	Standard	Private	Thu, Feb 28, 2019, 03:28:55 UTC ⋮
dp-archive	Archive	Private	Wed, May 13, 2020, 14:05:37 UTC ⋮

Click the “Create Bucket” button to open a dialog. The bucket defaults to creating standard objects but can be changed to immediately write archive objects. Archive object storage is much cheaper than standard object storage, but archive objects are not immediately available for read access. It may take up to an hour for an archive object to become available.

The cloud backup feature in Oracle ZFS Storage supports writing to and reading from archive buckets with no special configuration required, but if the backup is to be directly written to a bucket set as archive, a second bucket must be created in the Standard tier to hold the metadata for the archive object. When an archive object is read, it is promoted back to a standard object and will incur the Standard tier costs while it remains at that tier. The minimum retention period for an archive object is 90 days. If you delete or overwrite objects in the Archive tier before the retention requirements are met, you are charged the prorated cost of storing the data for the full 90 days.

Create Bucket [Help](#)

Bucket Name
oci_target

Default Storage Tier
☒ Standard
☐ Archive

The default storage tier for a bucket can only be specified during creation. Once set, you cannot change the storage tier in which a bucket resides. [Learn more about storage tiers](#)

☐ Enable Auto-Tiering
Automatically move infrequently accessed objects from the Standard tier to less expensive storage. [Learn more](#)

☐ Enable Object Versioning
Create an object version when a new object is uploaded, an existing object is overwritten, or when an object is deleted. [Learn more](#)

☐ Emit Object Events
Create automation based on object state changes using the [Events Service](#).

☐ Uncommitted Multipart Uploads Cleanup
Create a lifecycle rule to automatically delete uncommitted multipart uploads older than 7 days. [Learn more](#)

Encryption
☒ Encrypt using Oracle managed keys
Leaves all encryption-related matters to Oracle.
☐ Encrypt using customer-managed keys
Requires a valid key from a vault that you have access to. [Learn more](#)


Neither Auto-tiering to the Infrequent Access tier nor Object Lifecycle Management rules are supported for ZFS Storage cloud backups. Instead, consider writing directly to Archive storage.

It's easiest to allow OCI to manage the encryption keys to the objects, but the bucket can be set to use customer-managed keys. See the [OCI documentation on encryption](#) for more detail on providing your own keys.

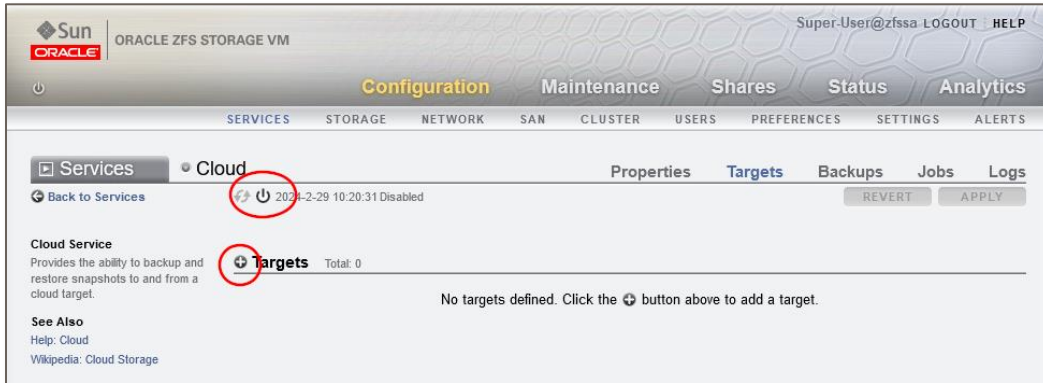
Click the “Create” button at the bottom of the screen to create the bucket.

ORACLE ZFS STORAGE – CONFIGURE CLOUD SERVICE

In the Oracle ZFS Storage Browser User Interface (BUI), Navigate to Configuration->Services and click on Cloud, as shown below.



In the Cloud properties screen, click the “on/off” icon to enable the Cloud service. The indicator next to the Cloud service name will turn green.



Creating the Target

Next, click the “Plus” icon next to the word “Targets” to enter the settings for a new OCI cloud target.

The screenshot shows the 'Add Cloud Target' dialog box. It has a 'CANCEL' button and an 'ADD' button. The form contains the following fields and options: 'Name' (text input), 'Location' (text input), 'Use data bucket' (checkbox, with a note: 'Check this box if your data and metadata are in different buckets. Data can be saved either in a standard or archive bucket.'), 'Bucket' (text input), 'Data Bucket' (text input), 'Instance Principal Authentication' (checkbox), 'User' (text input), 'Tenancy' (text input), 'Private Key' (text area), 'Use web proxy' (checkbox, with a note: 'If your system communicates to the web through a proxy, check this box and enter the configuration information below.'), 'Host : port' (text input), 'Username' (text input), 'Password' (text input), 'Write limit bandwidth' (checkbox, with a value of '0' and a unit of 'M/s'), and 'Read limit bandwidth' (checkbox, with a value of '0' and a unit of 'M/s').

Target Information

- **Name the Target**

Give the Target a relevant name. This must be unique within your Oracle ZFS Storage cluster.

- **Determine Location URL for OCI Bucket**

Oracle Object Storage is a regional service. It can be accessed through a dedicated regional API endpoint. The OCI Object Storage API endpoints use a consistent URL format of `https://objectstorage.<region-identifier>.oraclecloud.com`. For example, endpoint in US West (us-phoenix-1) is `https://objectstorage.us-phoenix-1.oraclecloud.com`.

The Region Identifier for all OCI regions can be found at <https://docs.cloud.oracle.com/en-us/iaas/Content/General/Concepts/regions.htm>.

- **Use data bucket**

Select this check box if your data and metadata will be in different buckets. Data will be saved to the Data Bucket, and metadata will be saved to the Bucket. If you use an archive bucket, the metadata must be saved to a standard bucket.

- **Bucket**

An established bucket name. Data can be saved in either a standard or archive bucket. The combination of Location, Bucket, and Data Bucket must be unique for your system. If you selected Use data bucket, this field is for the metadata bucket; if you did not select Use data bucket, this field is for both the metadata and data.

- **Data Bucket**

If you selected Use data bucket, this field is for the bucket that will hold the backup objects; if you did not select “Use data bucket”, this field is not available. The combination of Location, Bucket, and Data Bucket must be unique for your system.

- **Instance Principal Authentication**

Oracle ZFS-HA instances in OCI running OS8.8.63 or later may choose to use an OCI Instance Principal for authentication. In this case, the User and Tenancy fields will be greyed out. An instance principal is an Oracle Identity Access Management feature, enabling instances to become authorized principals that can perform actions on service resources, such as Object Storage. See the [OCI documentation](#) for more details.

- **User**

The user name OCID from the Oracle Cloud Infrastructure account.

- **Tenancy**

The tenancy name OCID from the Oracle Cloud Infrastructure account.

- **Private Key**

The private key for the account, which must match the public key uploaded to the Oracle Cloud Infrastructure account. Enter the key in the Privacy Enhanced Mail (PEM) format.

- **Use web proxy (optional)**

Select this check box to use a proxy for system communications with the web and complete the fields for host port name and number. The user name and password fields are optional.

- **Write limit bandwidth (optional)**

Select this check box to limit the traffic write bandwidth when uploading a cloud backup to the cloud target. Enter a value and select a unit of measurement. For example, 5 M/s limits writes to the cloud target to 5 megabytes per second.

- **Read limit bandwidth (optional)**

Select this check box to limit the traffic read bandwidth when restoring a cloud backup from the cloud target. Enter a value and select a unit of measurement. For example, 4 M/s limits reads from the cloud target to 4 megabytes per second.

Click the Add button at the bottom of the window when the settings have been entered. If the settings are correct and the bucket can be accessed, the indicator next to the target name will turn green. Otherwise, an error message will be displayed indicating the issue found with connecting to the bucket.

Note that once created, only the target name, web proxy, and bandwidth limitations can be edited. If any of the other fields are incorrect leading to ZFS Storage being unable to connect to the target, the target must be deleted and a new target created.

Routing to the Target

Depending on the infrastructure ZFS Storage is deployed in, the system may have issues reaching the target bucket consistently. Not uncommon are multiple alerts noting the target is alternately available and unavailable. This is usually caused by inconsistent routing through the interfaces available to the ZFS storage.

This can be avoided by defining a static route on ZFS Storage to force the object backups to go through the appropriate interface with access to the OCI object storage endpoint.

To add a static route, start by identifying the IP address(es) of the endpoints by examining the file at https://docs.oracle.com/en-us/iaas/tools/public_ip_ranges.json to find the region and the OBJECT_STORAGE endpoint subnets that are appropriate for your OCI tenancy. The example below shows the relevant endpoints for the US-Phoenix region,

```
{
  "last_updated_timestamp": "2024-04-17T11:11:06.887305",
  "regions": [
    {
      "region": "us-phoenix-1",
      "cidrs": [
        . . .
        {
          "cidr": "134.70.8.0/21",
          "tags": [
            "OSN",
            "OBJECT_STORAGE"
          ]
        },
        {
          "cidr": "134.70.16.0/22",
          "tags": [
            "OSN",
            "OBJECT_STORAGE"
          ]
        }
      ]
    }
  ],
}
```

Next, determine the interface to be used to connect to the cloud target. In the ZFS Storage browser interface, navigate to the Configuration->Network screen. In the ZFS-HA instance shown below, the `jh-zfs-a` interface is used for local client data access is in a local subnet with no external access. The `jh-zfs-adm-a` interface is used for administration and its `xxx.xxx.216.0/20` subnet is accessible by systems at the customer's on-premises location. The choice made here was to route the object storage backups through the `jh-zfs-adm-a` interface.



Next, go to the “Routing” screen and determine the default gateway used by the chosen interface. In this example, the gateway address of xxx.xxx.208.1 is used for the desired interface.

Network						Configuration	Addresses	Routing
						Multihoming model ☐ Loose ☐ Adaptive ☒ Strict		
+ Routing Table Entries Total: 8								
ALL : STATIC : DYNAMIC : DHCP : SYSTEM : DIRECT								
DESTINATION	GATEWAY	FAMILY	TYPE	STATUS	INTERFACE			
default		IPv4	static	active	jh-zfs-a IPv4 static, 222.255/20, via vtinet1			
default	2.208.1	IPv4	static	active	jh-zfs-adm-a IPv4 static, 216.181/20, via vtinet2			

Finally, create a new route by clicking the “plus” icon to bring up the dialog box.

Insert Static Route

CANCEL

ADD

Family

☒ IPv4

☐ IPv6

Kind

☐ Default

☒ Network

Destination

134.70.8.0/21

Gateway

208.1

Interface

jh-zfs-adm-a (vtinet2)

Enter the endpoint subnet and gateway address as previously determined and choose the desired interface from the pulldown menu. Click the Add button to create the new route.

Note that when there are multiple object storage endpoint subnets available in a region, such as those shown above for the us-phoenix-1 region, a new route should be added for each subnet shown.

Assigning Authorizations

For the appropriate user role(s) and if all services are not globally authorized, add authorizations for the Cloud filter for the Services scope. Like other services, there are three available authorizations: administer (enable and disable the service), configure (change service settings), and restart (restart the service).

Also, for the Cloud Targets scope, select all cloud targets or a specific cloud target, and select authorizations:

- backup - Back up snapshot data to cloud
- delete - Delete cloud backups
- restore - Restore cloud backups to local shares

For the Projects and Shares scope, globally set filters or specify filters, and add authorizations:

- backup - Read share data to back up to cloud target if cloud backup is authorized for the cloud target
- restore - Write cloud backup data if cloud restore is authorized for the cloud target

Use the appropriate user's role with the corresponding required authorizations for configuring and managing the cloud data service and cloud backups. More details are available in the [foo](#) section later in this document.

CLOUD SNAPSHOT BACKUP

The cloud backup feature in Oracle ZFS Storage is based on ZFS snapshots, which are an immutable point-in-time copy of shares and LUNs hosted on ZFS Storage. See the [ZFS Storage Appliance Administration Guide](#) as well as the technical brief [Oracle ZFS Storage Appliance: Best Practices for Implementing Snapshot Retention](#) for more on ZFS snapshots and their use in an overall data protection plan.

Creating the Initial Snapshot


The process below will create a ZFS snapshot and copy it to the OCI target created above.

- **Choose or create a share** to be backed up to OCI. If the share is empty, write some data to it.


Projects provide a way to group shares that have common properties or which may have a use case or client application in common. Shares inherit many of their properties from the project they belong to.

This example will back up a single share, *backup01* which is in the project *BackupPrj*, to the OCI target.


- **From the Shares tab, edit the share's properties** by bringing the mouse pointer over the line for *backup01* and clicking on the pencil icon. At the General properties screen that will open, click on the Snapshots tab.



- **Click the plus (+) icon to create a new snapshot.**




- Give the snapshot a name in the popup that appears and click Apply.



A dialog box titled "Create Snapshot" with "CANCEL" and "APPLY" buttons. It contains a "Name" field with the text "backup01-fullsnap1" and a "Retention policy" dropdown menu set to "Off".

- The created snapshot will now be listed on the screen.




The screenshot shows the "Snapshots" tab in the ZFS Storage Administration GUI for the path "pool-a/local/BackupPrj/backup01". The "Properties" section shows "Inherit from project" checked, "zfs/snapshot visibility" set to "Hidden", and "Enable retention policy for Scheduled Snapshots" unchecked. Below, a table lists the snapshot:

NAME	CREATION	UNIQUE	TOTAL	CLONES
backup01-fullsnap1	2024-3-1 17:27:54	0	186M	-

ZFS snapshots are very space efficient. This is reflected in the Unique and Total sizes shown for the snapshot. While it refers to 168MB of data, when first taken this snapshot uses no additional space. As data is deleted or changed in the mounted share, this snapshot will show a growing unique size as it holds on to the blocks that were changed.


Backup the Initial Snapshot to OCI

- Click on the “tape reel” icon to back up this share to the OCI bucket that was defined as the target.



This screenshot is identical to the previous one, but with a red circle highlighting the "tape reel" icon in the action column of the snapshot table, indicating the next step in the process.

- Choose the target and the format for this backup. If multiple targets have been defined on the ZFS Storage, they'll be listed in the pulldown menu. The defaults for the other fields are fine for this initial backup. See the [ZFS Storage Administration Guide](#) for details on these settings. Click the Apply button to begin the backup.




The 'Create Backup' dialog box contains the following fields and controls:

- Target:** A dropdown menu with 'jh-oci-target' selected.
- Format:** Radio buttons for 'ZFS' (selected) and 'Tar'.
- Incremental:** An unchecked checkbox.
- Parent:** A dropdown menu.
- Require parent exists:** A checked checkbox.
- Current Backups:** A section header with 'Total: 0'.
- No Backups:** A text label at the bottom.

A Backup Details dialog box is displayed and dynamically shows the backup transfer rate, amount of data transferred, and status `in-progress`. The status changes to `completed` after completion. If the `write limit` property was set for the cloud target, the write traffic bandwidth is limited to the set value when uploading the cloud backup to the cloud target. Click OK to close the dialog box.

To view further progress details, including the completion percentage, from the Configuration menu, select Services, then Cloud, then click the Jobs tab.

The Jobs list contains active and recently completed jobs in the cloud data service in chronological order. The Updates column contains the date and time that the job was created, and the date and time that the backup was started/updated. The Status column indicates the completion percentage, the amount of data transferred, and the backup transfer rate.



The screenshot shows the 'Configuration' window with the 'Cloud' service selected. The 'Jobs' tab is active, displaying a table of backup jobs.


TARGET	UPDATES	OPERATION	STATUS
jh-oci-target	2024-3-4 13:09:46 Created 2024-3-4 13:09:50 Updated	backup	completed: 186M @ 46.4MB/s

To view job details, double-click on the job in the Jobs list or click its Edit icon.

Backup an Incremental Snapshot to OCI

ZFS snapshots on Oracle ZFS Storage are very space-efficient, tracking only the blocks that have been changed or deleted in the file system since the last snapshot.

- **Take another snapshot** of the share. 298MB of data has been added to the example backup01 share. This will be an incremental snapshot and the name reflects that.



The 'Create Snapshot' dialog box contains the following fields and controls:

- Name:** A text input field with 'backup01-incr01' entered.
- Retention policy:** A dropdown menu with 'Off' selected.

The snapshot list now shows the two snapshots:

Snapshots Schedules Total: 2				
SHOW ALL MANUAL SCHEDULED				
NAME	CREATION	UNIQUE	TOTAL	CLONES
backup01-incr01	2024-3-4 13:39:14	0	484M	-
backup01-fullsnap1	2024-3-1 17:27:54	120K	186M	-

The original full snapshot used to have zero unique blocks, and it now has 120K that represents changed or deleted space since that snap was taken. This 120K is space owned solely by this snapshot and is space that is not accounted for when looking at the mounted share.

The new incremental refers to a total of 484M of disk space, but the snapshot currently takes up zero space as shown in the Unique column. Of that 484M, 298M is new data included in the incremental snapshot.

Create a cloud backup for the new snapshot.

When sending a ZFS snapshot to OCI object storage, two options control what is sent to the OCI target bucket: a toggle to indicate whether the backup is an incremental, and when an incremental is indicated, another toggle indicates whether the parent snapshot must exist on the same cloud target.

Depending on whether the Require parent exists property is selected as described next, the parent snapshot must be available on both the local system and backed up to the same cloud target, or the parent snapshot does not also have to be available on the same cloud target. The parent snapshot must also be in the same format as the incremental snapshot: zfs or tar.

For an incremental snapshot and to require that it must have a parent snapshot on the same cloud target, select the check box for Require parent exists. For a parent snapshot and to require that you cannot later delete the cloud backup with the parent snapshot if it has cloud backed-up incremental snapshots (children), select the check box for Require parent exists.

For both parent and incremental snapshots, this property can affect cloud backup restore and delete operations, as described in the following table.

Snapshot Type	Property Selected?	Restore Effect	Delete Effect
Parent	Yes	No effect.	Cannot delete cloud backup with parent snapshot if cloud backed-up incremental snapshots exist.
Parent	No	No effect.	Can delete cloud backup with parent snapshot if cloud backed-up incremental snapshots exist, which would save space on the cloud target.
Incremental	Yes	Parent snapshot must exist on the same cloud target, as well as on the local system.	No effect.
Incremental	No	Parent snapshot must exist on the local system. This allows you to create a cloud backup of an incremental snapshot without its parent snapshot on the same cloud target.	No effect.

Here the type is set to Incremental so that only the new data in the share will be sent to the backup target. Note that when an Incremental backup is selected, the pulldown to choose the parent snapshot is activated.

The 'Create Backup' dialog box shows the following configuration: Target is 'jh-oci-target', Format is 'ZFS' (selected), Incremental is checked, Parent is 'backup01-fullsnap1', and 'Require parent exists' is checked. Below the configuration, a section titled 'Current Backups' shows 'Total: 0' and 'No Backups'.

After the incremental snapshot has loaded, we can see a total of 298M has been uploaded, the size of just the new data.

The 'Backup Details' dialog box shows the following statistics: Format is 'zfs', Rate is '149MB/s', Transferred is '298M', and Status is 'completed'.

Creating a Tar-format Backup in OCI

This example details creating a backup in tar format that can be restored to either a ZFS Storage system or to a computer running Linux or Solaris.

- Create a new snapshot. Because the backups inherit the snapshot name when viewed on a ZFS Storage system, here a new incremental snapshot is created with the name backup01-tar-inc. It is functionally similar to the previous incremental snapshot and at this point has the same number of unique blocks, but the new name will help identify our tar backup among the ZFS-formatted backups.
- **Click on the “tape reel” icon for the latest snapshot** to back up the snapshot to the OCI bucket that was defined as the target. Click the radio button to change the format to Tar. Note that the Incremental checkbox and the parent pulldown become inactive. Click Apply to start the transfer.

The 'Create Backup' dialog box shows the following configuration: Target is 'jh-oci-target', Format is 'Tar' (selected), Incremental is unchecked, Parent is a dropdown menu, and 'Require parent exists' is checked. Below the configuration, a section titled 'Current Backups' shows 'Total: 1' and a table of current backups.


TARGET	COMPLETED	FORMAT	STATUS
jh-oci-target	2024-3-8 21:34:49	zfs	completed

CLOUD SNAPSHOT RESTORE


Recover a Cloud Snapshot to ZFS Storage

This example details recovering an incremental snapshot from the OCI bucket back to ZFS Storage.

- **Identify the snapshot(s) to be restored** by navigating to Configuration→Services→Cloud→Backups, select backup and click “restore” icon. Here the latest incremental snapshot for the *backup01* share is being selected.




- **Select a pool, project and share name** to restore the snapshot into. This snapshot will be restored to the same storage pool but a different project, *RestorePrj* and a different share name, *restore01*, will be used.

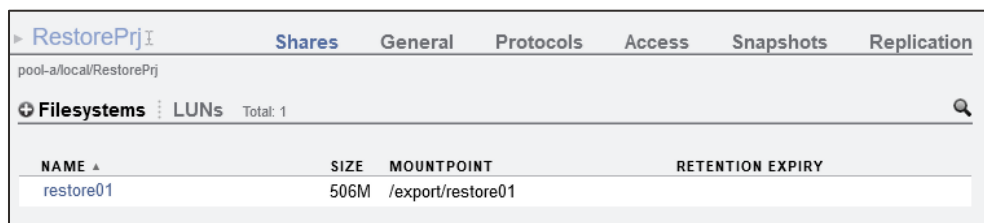


The snapshot can also be restored to an existing share by checking the “Use existing share” box and entering an existing local share name. See the ZFS documentation on [Restoring a Cloud Backup](#) for more details on using existing shares.


Click Apply to begin the restore process. The full share size of 483M is shown to have been transferred.



- Verify that the new share has been recovered.



Note that its list of snapshots shows that both the incremental and parent snapshots were restored in the process. From here, it would be possible to also recover the share to any previous incremental snapshots (if any) as well as to the original parent snapshot.



Recover a Tar-format Cloud Snapshot to a Linux or Solaris System

This example details recovering an incremental snapshot from the OCI bucket back to ZFS Storage. It is assumed here that the system that the object is being recovered to [has the OCI CLI utility installed](#) and that a proper config file is in place so that some parameters do not have to be specified on the command line itself.

- Identify the snapshot(s) to be restored.** The backups are not listed in OCI object storage the same way that they are on ZFS Storage. The command to list a bucket's objects is:


```
"oci os object list --bucket-name <target_bucket> --namespace <your_namespace>"
```

Here the bucket name is the name given to the target bucket in OCI and the namespace is the identifier for your tenancy's object namespace. (This is often the name of your tenancy but may be different. Ask your tenancy administrator what your object namespace might be.)

The OCI command will return a JSON object with all the contents of the bucket, which may contain objects from multiple sources. You can use OCI CLI parameters to present the data in a table format as well. For example:

```
local% # oci os object list -bn jh-backup -ns my_ns --output table --query \
'data[*].{"Name":"name","Size":"size","Time created":"time-created"}'
```

Name	Size	Time created
zfs/backups/tar/cda3106f31746130/8c81de53e21df174/000000001	499393398	2024-03-11T15:15:34.817000+00:00
zfs/backups/zfs/39f49a2a6e439496/a17d59d91fba2b83/000000001	99926174952	2024-03-05T17:21:44.976000+00:00
zfs/backups/zfs/cda3106f31746130/6154202a6d20a2b2-adbb8088c2b62f1a/000000001	312095088	2024-03-08T21:34:49.396000+00:00
zfs/backups/zfs/cda3106f31746130/adbb8088c2b62f1a/000000001	194770932	2024-03-04T18:09:50.433000+00:00
zfs/backups/zfs/cda3106f31746130/cc7c5e0cd4d24ce7-6154202a6d20a2b2/000000001	4672	2024-03-18T19:52:58.332000+00:00
zfs/dependencies/zfs/cda3106f31746130/6154202a6d20a2b2-cc7c5e0cd4d24ce7	0	2024-03-18T19:52:58.352000+00:00
zfs/dependencies/zfs/cda3106f31746130/adbb8088c2b62f1a-6154202a6d20a2b2	0	2024-03-08T21:34:49.439000+00:00
zfs/log/4b5ba1ce-8508-4725-dc7a-e4a5184c7f34/000000001	122	2024-03-05T16:42:04.959000+00:00
zfs/log/4b5ba1ce-8508-4725-dc7a-e4a5184c7f34/000000002	124	2024-03-05T17:21:46.841000+00:00
zfs/log/c218dd93-141d-4c47-b932-8175fe1b89df/000000001	122	2024-03-01T22:52:35.861000+00:00
zfs/log/c218dd93-141d-4c47-b932-8175fe1b89df/000000002	124	2024-03-01T22:52:37.229000+00:00
zfs/log/c218dd93-141d-4c47-b932-8175fe1b89df/000000022	139	2024-03-18T19:52:58.062000+00:00

zfs/log/c218dd93-141d-4c47-b932-8175fe1b89df/00000023	141	2024-03-18T19:52:58.399000+00:00
zfs/log/c218dd93-141d-4c47-b932-8175fe1b89df/00000024	140	2024-03-18T19:56:59.613000+00:00
zfs/manifests/tar/cda3106f31746130/8c81de53e21df174	794	2024-03-11T15:15:34.854000+00:00
zfs/manifests/zfs/39f49a2a6e439496/a17d59d91fba2b83	797	2024-03-05T17:21:46.720000+00:00
zfs/manifests/zfs/cda3106f31746130/6154202a6d20a2b2-adbb8088c2b62f1a	848	2024-03-08T21:34:49.476000+00:00
zfs/manifests/zfs/cda3106f31746130/adbb8088c2b62f1a	761	2024-03-04T18:09:50.498000+00:00
zfs/manifests/zfs/cda3106f31746130/cc7c5e0cd4d24ce7-6154202a6d20a2b2	834	2024-03-18T19:52:58.375000+00:00
zfs/source/4b5balce-8508-4725-dc7a-e4a5184c7f34	163	2024-03-05T16:39:14.212000+00:00
zfs/source/c218dd93-141d-4c47-b932-8175fe1b89df	166	2024-03-01T18:20:06.925000+00:00
zfs/source/cc866fa5-a519-4f92-ab7f-8484713a5711	165	2024-03-01T18:20:07.238000+00:00
zfs/target	85	2024-03-01T18:20:06.879000+00:00
+-----+-----+-----+		

When a cloud snapshot backup is created, a unique ID string is generated that can be used to restore the snapshot. The above output is described as follows:

Component	Property Selected?
Name	zfs/backups/type – Identifies the snapshot backup ID and type of snapshot (zfs/tar). These are the backup objects that can be restored. zfs/dependencies – identifies any incremental backup denoted with a /, which means that this backup cannot be removed unless the incremental backup snapshot is removed zfs/log – Identifies log data of the snapshot backup zfs/manifests – Identifies the metadata for each of the backups taken zfs/target – identifies the cloud snapshot target zfs/source – identifies the source name of the backup
Size	Cloud snapshot backup size in bytes
Time Created	Identifies the date/time that the backup was created

If there are multiple backup objects returned, the time-created field in the JSON item can be used to narrow down the correct backup.

In some cases, an OCI bucket may hold backups from multiple ZFS Storage systems. The substring directly after the type is an identifier for the source system’s storage pool.

The manifest objects contain a JSON object listing the metadata for the object. If needed, the manifest can be downloaded and examined.

```
% oci os object get --bucket-name jh-backup -ns my-ns --name
"zfs/manifests/tar/cda3106f31746130/8c81de53e21df174" --file ./backup.man
Downloading object [#####] 100%

% cat .\backup.man
{"source": {"id": "c218dd93-141d-4c47-b932-8175fe1b89df", "name": "jh-zfs-a", "product":
"SUNW,nori", "version": "", "data-name": "pool-a/local/BackupPrj/backup01", "data-type":
"zfs:filesystem", "data-version": "50", "snap-name": "backup01-tar-inc", "format-name": "tar",
"format-version": "v1", "mountpoint": "/export/backup01", "creation": "2024-03-11T14:38:57",
"space-data": 507707392, "space-unique": 0, "guid": "8c81de53e21df174", "data-owner": "c218dd93-
141d-4c47-b932-8175fe1b89df"}, "id": "cda3106f31746130/8c81de53e21df174", "started": "2024-03-
11T15:15:30", "completed": "2024-03-11T15:15:34", "length": 499393398, "size": 499393398,
"segments": 1, "data": "jh-backup", "version": 1, "backups": [{"segment": "000000001", "etag":
"oq0/ZajyF/rb9Qszc8Q4Cg==15", "length": 499393398}]}
```

The manifest above shows that the storage pool with the identifier cda3106f31746130 is from the ZFS system “jh-zfs-a”. Other fields that may be useful in the manifest object include:

name : the system name of the ZFS Storage system the backup came from

data-name : the ZFS share that was backed up
snap-name : the name of the snapshot the cloud backup was taken from

Note that if separate buckets are used for the backups and metadata, the buckets used in the CLI commands will be different for examining the manifests and backup objects.

- **Recover the backup object to a tar file** using the identified full name string.

```
% oci os object get -ns my_ns -bn jh-backup --name \  
" zfs/backups/tar/cda3106f31746130/8c81de53e21df174/000000001" \  
--file ./zfs-backup.tar
```

The recovered file can then be expanded with the `tar` command.

When large snapshots are backed up, they may be broken up into multiple parts. In such cases, the objects must be downloaded individually and joined together again on the target system before they can be expanded, as shown in the following example:

```
% oci os object get -ns my_ns -bn jh-backup --name \  
" zfs/backups/tar/cda3106f31746130/8c81de53e21df174/000000001" \  
--file ./zfs-01.tar
```

```
% oci os object get -ns my_ns -bn jh-backup --name \  
" zfs/backups/tar/cda3106f31746130/8c81de53e21df174/000000002" \  
--file ./zfs-02.tar
```

```
% cat zfs-01.tar zfs-02.tar | tar xvf -
```

USING OCI OBJECT RETENTION POLICIES WITH ZFS STORAGE CLOUD BACKUPS

It is possible to set object retention policies in OCI to hold on to objects for a pre-determined period. See [Data Retention Rules for Object Storage](#) for more on object retention in OCI. Of special note is the ability to lock retention rules so that the rules may not be modified or deleted, which will affect the ability to remove objects in the buckets the rules apply to.


It is highly recommended that when first setting up retention policies that short retention periods be used and that the rules remain unlocked. It is possible to accidentally wind up having to

Object retention requires the use of two separate buckets in OCI – one for metadata and one for the backup data.

- From OCI Cloud management console, go to the Object Storage pane and create two buckets:
 - One bucket for data. For example, `ret-bucket` (data).
 - One bucket for metadata. For example, `ret-meta-bucket` (metadata).

ret-bucket	Standard	Private	Wed, Jun 16, 2021, 16:17:39 GMT	:
ret-meta-bucket	Standard	Private	Wed, Jun 16, 2021, 16:56:06 GMT	:

- From ZFS Storage Appliance BUI, go to Configuration→Cloud→Targets to create a target and identify the data bucket and metadata bucket. The example below shows the configuration according to the names used above:



Edit Cloud Target [CANCEL] [APPLY]

Name: ret-target

Location: https://objectstorage.us-ashburn-1.oraclecloud.com

Bucket: ret-meta-bucket

Data Bucket: ret-bucket

User: ocid1.user.oc1..[REDACTED]

Tenancy: ocid1.tenancy.oc1..[REDACTED]

Use web proxy: ☒ If your system communicates to the web through a proxy, check this box and enter the configuration information below.

Host : port: www-proxy.us.oracle.co : 80


Username: root

Password: (set)

Write limit bandwidth: ☐ 0 M/s

Read limit bandwidth: ☐ 0 M/s

- From OCI cloud management console, go to Object Storage pane and set the retention policy. For example, the ret-bucket has a time-bound retention policy of one year.



ret-bucket

[Edit Visibility] [Move Resource] [Re-encrypt] [Add Tags] [Delete]

Bucket Information [Tags]

Visibility: Private

Namespace: zs

Default Storage Tier: Standard

Auto-Tiering: @ Disabled [Edit]

Approximate Count: 1 objects [Edit]

ETag: 2f41cc38-b2d0-4315-8326-7afaaa645c8

OCID: ...4eaplwda [Show] [Copy]

Encryption Key: Oracle managed key [Assign]

Created: Wed, Jun 16, 2021, 16:17:39 UTC

Compartment: [REDACTED]

Approximate Size: 1.85 GiB [Edit]

Emit Object Events: @ Disabled [Edit]

Object Versioning: @ Disabled [Edit]


Retention Rules

[Create Rule]

Name	Retention Duration	Lock State	Scheduled Lock Time	Last Modified
retention-rule-1yr	1 Year	Unlocked	-	Wed, Jun 16, 2021, 21:04:21 UTC

Showing 1 item

- Test with cloud snapshot backup. For example, send a cloud snapshot backup to the ret-bucket target in OCI (cloud).



Create Backup [CANCEL] [APPLY]

Target: ret-target

Format: ☒ ZFS ☐ Tar

Incremental: ☐

Parent: cloudbackup-snap

Require parent exists: ☒

Current Backups Total: 2

TARGET	COMPLETED	FORMAT	STATUS
OCcloud	2021-3-29 14:17:50	zfs	completed
OCcloud-demo	2021-5-25 13:48:10	zfs	completed

MONITOR CLOUD SNAPSHOT BACKUPS AND OBJECT STORE BUCKET

ZFS Storage analytics and logging provide some deep insights into the performance and history of the cloud backups.


- **Monitor cloud backup throughput with ZFS Analytics.**

You can monitor cloud snapshot transfers with ZFS Storage Analytics.

Cloud snapshot backup throughput performance to OCI cloud object storage from on-premises storage will depend upon connectivity to OCI. OCI's FastConnect feature provides a dedicated, private connection between your data center and OCI. FastConnect provides higher-bandwidth options, and a more reliable and consistent networking experience compared to internet-based connections.

When backing up from ZFS-HA instances in OCI, bandwidth may be affected by factors such as the shape used for the controllers. E4.Flex shapes, for instance, allocate bandwidth at 1 Gbps per OCPU.

Go to the Analytics screen and select the “Plus” icon next to Add statistics. Select Cloud bytes or Cloud requests.




- **Monitor cloud backup job details and alert information.**


Information about the current cloud snapshot backup job is detailed at the top of the BUI when then the job is in progress. Backup job details can also be displayed from Configuration→Services→Cloud→Backups or Configuration→ Services→Cloud→Jobs, which includes most recent backup jobs.

Backup job details can also be reviewed in the alert logs in Maintenance→Logs.


- Review details of last cloud snapshot backup job details in Configuration→ Services→Cloud→Jobs. Click the edit icon on the line of the job to be examined.




- Review details of recent cloud snapshot backups in Configuration→ Services→Cloud→Backups. Click the edit icon on the line of the backup to be examined.



- Review previous cloud snapshot backup log entries in Maintenance→Logs.

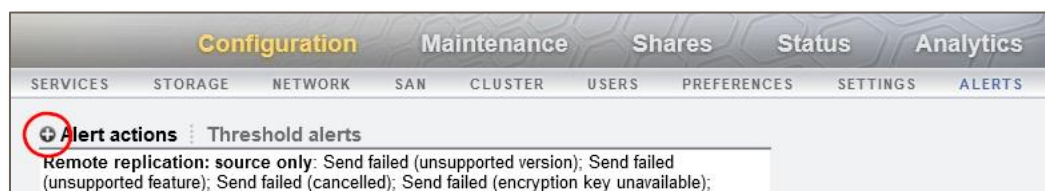


- Select a cloud snapshot backup alert and click on the “show alert details” icon for more detail on that alert.



- **Send external alert messages when a problem is encountered**

Oracle ZFS Storage has the capability of sending emails or SNMP traps when something out of the ordinary is seen with a cloud backup job. To enable an alert, go to Configuration→Alerts and click on the Plus icon next to Alert Actions.



Choose Cloud Backup from the Category pulldown menu, click on Subset, then choose the events that are considered important enough to raise an external alarm. Choose to send email or an SNMP trap from the Alert Actions pulldown menu and populate the fields appropriate for your tenancy. Multiple actions can be taken for each alert. Actions can be added by clicking the Plus icon next to Alert Actions towards the bottom of the window. When done, click Add to enable the action.

Add alert action [CANCEL] [ADD]

Events

Category: Cloud Snapshot

Events: ☐ All ☒ Subset

- ☐ Target online
- ☒ Target offline
- ☒ Target failed
- ☐ Backup to target started
- ☐ Backup to target finished
- ☒ Backup to target failed
- ☐ Restore from target started
- ☐ Restore from target finished
- ☒ Restore from target failed
- ☐ Delete backup started
- ☐ Delete backup finished
- ☒ Delete backup failed

Alert actions

Send email: Send to: [TEST]

Subject:

See [Configuring Alerts](#) in the ZFS Storage documentation for more details on alerts in ZFS Storage.

CLOUD SNAPSHOT BACKUPS ROLES AND AUTHORIZATION

Oracle ZFS Storage provides the ability to limit access to certain features through the use of roles and the assignment of authorizations to these roles. Users can then have one or more of these roles assigned to them.

A pre-configured category of authorizations (Cloud Snapshot) allows a storage administrator to determine who can add, remove, and restore cloud snapshot backups. These authorizations can be applied to existing roles, or a new role can be created to specifically allow certain actions.

See “[Understanding Users and Roles](#)” in the ZFS Storage documentation for details on using this feature.

- **To add a new role to manage Cloud Backups**, go to Configuration→Users and click on the Plus icon next to Roles.



In the window that appears, start by choosing a type – Local or Directory. Choosing Local will apply only to this ZFS Storage Appliance or ZFS-HA cluster; Directory roles will be applied to LDAP or Active directory groups. See “[Adding a Role](#)” in the ZFS Storage documentation for more details. Enter a name and a description for the role.

Select “Cloud targets” in the Scope pulldown menu, then use the checkboxes to determine the actions a user with this role can perform. Optionally, on systems that have more than one cloud target defined, you may choose a specific target for the role to apply to. In the example below, a Cloud backup operator role is defined allowing any user with this role assigned to create and restore from cloud snapshots, but may not remove any backups from OCI object storage.

The screenshot shows the 'Add Role' dialog box. The 'Properties' section indicates the role is local to the appliance. The 'Type' is set to 'Local'. The 'Name' is 'CloudSnapsOperator' and the 'Description' is 'ckup and restore backups'. The 'Authorizations' section shows the 'Scope' set to 'Cloud targets'. Below this, a list of authorizations is shown with checkboxes: 'backup' (checked), 'delete' (unchecked), and 'restore' (checked). The 'Cloud target name' is set to '*'. The 'ADD' button is visible in the top right corner of the dialog box.

Properties	
This role is local to this appliance.	
Type	Local
Name	CloudSnapsOperator
Description	ckup and restore backups

Authorizations	
Scope: Cloud targets	
Cloud target name: *	
Selected: 2 / Total: 3	
☒	backup: Backup snapshot data to cloud
☐	delete: Delete cloud snapshots
☒	restore: Restore cloud snapshots to local shares

You can then create or modify an existing user to assign the CloudSnapsOperator role to.

Add User

CANCELADD

Properties

This is a locally-defined appliance administrator.

TypeLocal

UsernameBackupOperator

User IDauto

Full NameBackup Operator

Password

Confirm

Require session annotation

Kiosk user

Kiosk screenhttps://osc-zs5-2-n2-02:215/#status/dashboard

SMB enabled user

RolesExceptions

Total: 7

1-6

NAME	DESCRIPTION
☐ CloudSnapsAdmin	Can backup, delete, and restore backups
☒ CloudSnapsOperator	Can backup and restore backups
☐ OEMrole	read / write account for OEM
☐ OEMtest	Restricted read / write account for OEM
☐ basic	Basic administration
☐ ikorole	iko demo

If a Cloud Snapshots Operator attempts to delete a backup, the system will halt the request given the user does not have authorization to perform the action.

OK

You are not authorized to perform this action. If you wish to proceed, contact an administrator to obtain the proper credentials.

CUSTOMIZATION TIPS

OCI Command Line Simplification

The OCI command line tool can be used to manage target buckets and monitor OCI object stores, you might consider setting up profiles to simplify the OCI command line execution.

Profiles can be used to simplify command line input when working with a specific target. Instead defining each parameter separately on the command line to access a target, these parameter values can be grouped into a profile so that they will be specified simply by using the `--profile` argument on the command line.

The profile's settings are defined under its title (for example, `[zfssa]`). A special name of `[DEFAULT]` may be used and values in the default section will be used unless overridden by the values in the named profile.

Two files are used to define various parts of a profile. These files usually reside in the `.oci` directory in the local user's home directory, notated as `~/.oci`.

The `.oci/config` file will contain the values for the `user`, `fingerprint`, `key_file`, and `region` fields, while the `.oci/oci_cli_rc` file will contain the values for the `compartment_id`, `namespace`, and `endpoint` fields.

- **~/.oci/config**

Create a profile named `zfssa` in the config file by adding lines like the following:

```
[zfssa]
user=ocid1.user.oc1..opc-user
fingerprint= eb:5c:e1:c1:8a:57:26:de:ad:be:ef:30:28:b0:af:92
key_file=~/.oci/oci_api_key.pem
tenancy=ocid1.tenancy.oc1..nobody
```

If any of the four fields shown above are not present in the config file, an error will be shown.

The values for `key_file` and `fingerprint` fields are from the keys generated as described in the [CLI Installation and Certificate Creation](#) section of this document above.

The `user` field must be in the Oracle Cloud ID (OCID) format. When connecting to OCI object storage on ZFS Storage, this field takes the format of `ocid1.user.oc1..<zfs_username>`.

The `tenancy` value is required but has no equivalent in ZFS Storage. The value shown above is an acceptable dummy value.

- **~/.oci/oci_cli_rc**

Create a profile named `zfssa` in the `oci_cli_rc` file by adding lines like the following:

```
[zfssa]
endpoint=http://ZFSSA-name-or-IP/oci
compartment_id=cloud target
namespace=cloud-target
```

Any values for the three fields shown above not included in the `oci_cli_rc` file must be provided on the command line.

The `endpoint` field requires a URL that uses the IP address or a DNS-defined name for the target ZFS storage followed by `/oci`.

The `compartment_id` and `namespace` fields both have a value of the ZFS share created to be the backup target in the [Share Creation On the Target ZFS Storage Appliance](#) section of this document.

Using a Profile on the Command Line

To use a defined profile, call it by name using the `-profile <profile_name>` argument.

The following two command lines creating a bucket in the cloud-target share are equivalent when the `zfssa` profile is defined as above:

- `local% oci os bucket create --endpoint "http://zfssa.example.com/oci" -ns cloud-target -c cloud-target -name data-bucket`
- `local% oci os bucket create --profile zfssa --name data-bucket`

CONNECT WITH US

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com).

Outside North America, find your local office at [oracle.com/contact](https://www.oracle.com/contact).

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2024, 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