Oracle® Cloud

Using Oracle Cloud Infrastructure Storage Software Appliance

Cloud Distribution, Release 16.3.1.4
E68278-11
May 2020

ORACLE

Oracle Cloud Using Oracle Cloud Infrastructure Storage Software Appliance, Cloud Distribution, Release 16.3.1.4

E68278-11

Primary Author: Oracle Corporation

Contents

Preface

Topics

Audience

Using Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution is intended for users of Oracle Cloud Infrastructure Compute Classic who want to provide shared storage capacity for their Oracle Cloud Infrastructure Compute Classic instances, leveraging traditional file services protocols like NFS.

> **Note:**
>
> Users of this document must be familiar with the basics of Oracle Cloud Infrastructure Compute Classic and Oracle Cloud Infrastructure Object Storage Classic. At relevant places in this document, links to the required documentation are included.

Related Resources

For more information, see these Oracle resources:

- Oracle Cloud

  `http://cloud.oracle.com`
- *Using Oracle Cloud Infrastructure Compute Classic*
- *REST API for Oracle Cloud Infrastructure Compute Classic*
- *Using Oracle Cloud Infrastructure Object Storage Classic*

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.

Convention	Meaning
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1
Getting Started

Topics

About Storage Software Appliance – Cloud Distribution

Oracle Cloud Infrastructure Storage Software Appliance is a cloud storage gateway that you can use to access your storage service instance in the cloud over the NFSv4 protocol. The appliance can be deployed in Oracle Cloud.

- The on-premises deployment of the appliance enables applications running in your data center to easily store and access data in your storage service instance in the cloud.
- With Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution, the appliance is provisioned on an Oracle Cloud Infrastructure Compute Classic instance and plays the role of a file server in the cloud. It provides shared, infinitely scalable, low-cost, and reliable file storage capacity for your Oracle Cloud Infrastructure Compute Classic instances running Oracle Linux. Applications running on multiple Oracle Linux instances of Oracle Cloud Infrastructure Compute Classic can access shared file storage by using the NFSv4 protocol. You pay only for the storage space that your applications use, and the capacity expands automatically as your applications write data. To learn more about the features, see Features of the Appliance and watch this short video.

Video

Note that, within Oracle Cloud Infrastructure Compute Classic, you can provide storage capacity for your instances by attaching block storage volumes. But storage volumes can’t be shared because a volume can be attached to only one Oracle Cloud Infrastructure Compute Classic. Besides, there's a limit to the number and the size of the storage volumes that you can attach to each Oracle Cloud Infrastructure Compute Classic instance.

What can I use Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution for?

Oracle Cloud Infrastructure Storage Software Appliance is an effective cloud gateway for many workloads. Use the following guidelines to determine whether the appliance is appropriate for your specific use cases and workloads:

- The appliance supports NFSv4 in asynchronous mode and POSIX Sync mode. The POSIX Sync mode is enabled in the appliance by default.

In the asynchronous mode, there is scope for data loss in the event of a sudden server failure. Avoid using the appliance for workloads and use cases that require synchronous write behavior.

- The appliance is ideal for backup and archive use cases that require the replication of infrequently accessed data to cloud containers. (Not available on Oracle Cloud at Customer)
- Carefully consider use cases that involve frequent changes to existing files. Each time a file is modified and closed, the appliance creates a new version of the file, which is then uploaded to the container in your service instance, replacing the previous version. The appliance will be less efficient and may not perform optimally for this type of workload.
- Don't run applications and executables directly from the appliance mount points, particularly if the appliance cache is not large enough for all the files that the applications will access. Applications typically create temporary files and modify them often, affecting the operational efficiency of the appliance.

How does the appliance work?

You provision Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution on one of your Oracle Cloud Infrastructure Compute Classic instances and configure *filesystems*, each of which maps to a container in your storage service instance in the cloud. You then mount these filesystems, using the NFS v4 protocol, on the *client* Oracle Cloud Infrastructure Compute Classic instances that you want to provide shared file storage for. Currently, the appliance is supported for Oracle Linux instances of Oracle Cloud Infrastructure Compute Classic.

When applications running on the client instances write data to the mount points on the instance, the data is cached on the appliance instance and then uploaded asynchronously to your storage service instance in the cloud. Frequently accessed data is cached on the appliance. You can configure the appliance cache size.

Your applications enjoy read and write performance that’s comparable to the performance when using network-attached storage (NAS). And you get the benefits of low cost, infinite scale, and reliability of your storage service instance in the cloud.

The following figure shows two Oracle Cloud Infrastructure Compute Classic instances sharing access over NFS to data stored in your storage service instance in the cloud:

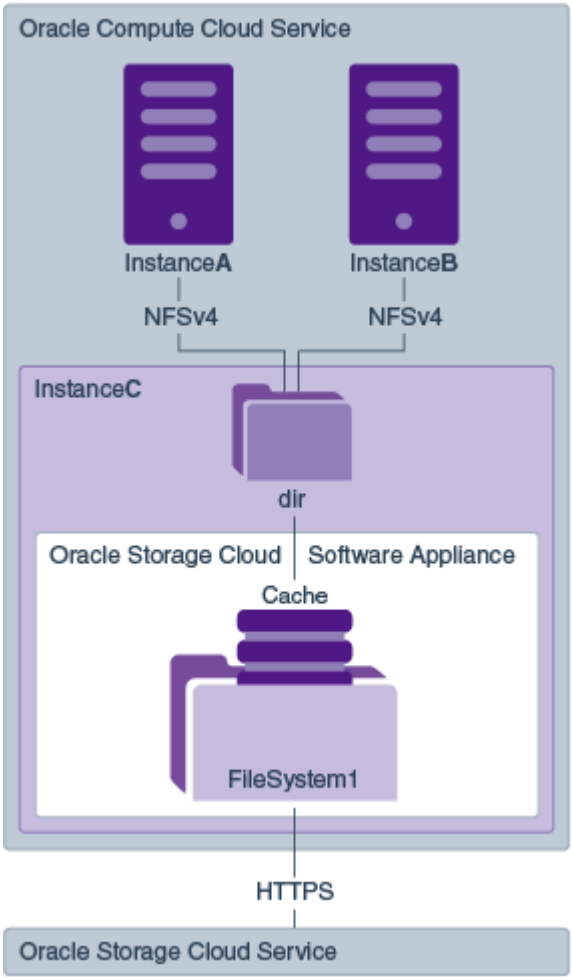

In this example, `instanceA` and `instanceB` share read and write access, through Oracle Cloud Infrastructure Storage Software Appliance provisioned on `instanceC`, to data that's stored in your storage service instance in the cloud.

Features of the Appliance

The following are the features of Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution:

- Shared Storage
- POSIX-Compliant NFS Access to Oracle Cloud Infrastructure Object Storage Classic
- Near-Local Performance, Using a Tunable Cache
- Granular Encryption to Enable Data Security and Storage Efficiency
- End-to-end Data Integrity with Checksum Verification
- Efficient Handing of Large Files
- High Availability
- Support for Data Archival

- Support for File Versions Compaction and End-to-End Delete
- Quick Access to Select Files with Cache Pinning
- Appliance Health Check

Shared Storage

Multiple Oracle Cloud Infrastructure Compute Classic instances can concurrently write data to and retrieve data over NFSv4 from a common set of files.

POSIX-Compliant NFS Access to Oracle Cloud Infrastructure Object Storage Classic

Using Oracle Cloud Infrastructure Storage Software Appliance, your applications can interact with your storage service instance in the cloud through standard protocols, without invoking direct REST API calls to the service. The appliance is compliant with POSIX standards. You can create multiple NFS shares within a single appliance instance. Using a single appliance instance, you can connect to multiple containers on the storage service. The files are copied to the appliance filesystem by using the NFSv4 protocol. The appliance supports NFS in asynchronous and POSIX-Sync modes. The appliance stores the files as objects in your account by using the HTTPS protocol.

Near-Local Performance, Using a Tunable Cache

To your applications, uploading and retrieving data using Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution appears to be almost the same as accessing block storage within Oracle Cloud Infrastructure Compute Classic.

- Data that your applications write is first stored in an upload buffer on the local disks of the Oracle Cloud Infrastructure Compute Classic instance that hosts Oracle Cloud Infrastructure Storage Software Appliance. The files are then uploaded asynchronously to your storage service instance in the cloud. If Oracle Cloud Infrastructure Storage Software Appliance or its host instance stops running for any reason, then cached data isn't lost. The pending upload operations resume automatically when the appliance starts running again.
- Oracle Cloud Infrastructure Storage Software Appliance caches frequently accessed data locally in a read cache, enabling fast data retrieval. You can increase the cache capacity by adding data disks. You can attach up to nine data disks, each up to 2 TB in size.

Based on your workload, you can tune the cache limit—both the number of files that can be cached and the maximum cache size.

Granular Encryption to Enable Data Security and Storage Efficiency

Oracle Cloud Infrastructure Storage Software Appliance stores data securely in your account. The appliance provides data security by keeping the data encrypted both at rest in the storage cloud and during transit from cache to storage cloud. To ensure data security, you can configure the appliance to encrypt data on premises before the data is stored in your account, and decrypt files when they are retrieved. You can update the encryption keys at any point in time. By having granular control that enables encryption at the NFS share level, you can ensure that only sensitive data is being encrypted. This minimizes the performance cost associated with encryption. Encryption is supported at the filesystem level. You can configure encryption for each configured filesystem, which ensures that sensitive data is secured in your account. By using more granular controls, you can increase storage efficiency.

End-to-end Data Integrity with Checksum Verification

The built-in data integrity checks ensure that data is validated as it moves through the data path, from Oracle Cloud Infrastructure Storage Software Appliance, to your account, enabling seamless end-to-end data integrity. Checksum verification helps in ensuring the data integrity. Metadata integrity checks are performed to ensure that the metadata is in consistent state.

Efficient Handing of Large Files

Oracle Cloud Infrastructure Storage Software Appliance supports large files that exceed the maximum size allowed by Oracle Cloud Infrastructure Object Storage Classic. Large files are sliced into 1 GB segments and each segment is stored as a separate object. The metadata database maintains the manifest of the segments that comprise a given object, so that multi-segment files can be reconstructed automatically when read through the appliance. The segments are uploaded sequentially.

High Availability

If the Oracle Cloud Infrastructure Compute Classic instance that hosts Oracle Cloud Infrastructure Storage Software Appliance crashes for any reason, the instance is re-created automatically with minimal down time. Appliance configuration data remains intact. Cached data isn't lost. Pending upload operations, if any, at the time of the crash resume automatically. And the cache disks are re-attached automatically to the appliance.

> **Note:**
>
> When the appliance instance is re-created, it gets a new private IP address. So after the instance is re-created, you must mount the filesystems again on your client instances.

Data that you store in your storage service instance in the cloud is replicated automatically on multiple storage nodes within the data center. If one of the nodes fails, copies of the data continue to be available.

Support for Data Archival

(Not available on Oracle Cloud at Customer)

Oracle Cloud Infrastructure Storage Software Appliance supports uploading and restoring objects in containers of the `Archive` storage class in Oracle Cloud Infrastructure Object Storage Classic.

In metered accounts, you can create containers of two *storage classes*, `Standard` (default) and `Archive`. You can use `Archive` containers to store large data sets that you don't need to access frequently, at a fraction of the cost of storing data in `Standard` containers. Note that to download data stored in `Archive` containers, you must first *restore* the objects. The restoration process can take up to four hours depending on the size of the object. A few features, such as bulk upload and deletion are not supported for `Archive` containers. `Archive` containers are ideal for storing data such as email archives, data backups, and digital video masters. For information about

the pricing and other terms for the `Archive` storage class, go to https://cloud.oracle.com/storage?tabID=1406491833493.

Support for File Versions Compaction and End-to-End Delete

Oracle Cloud Infrastructure Storage Software Appliance supports deletion of old file versions from the storage cloud.

Oracle Cloud Infrastructure Storage Software Appliance provides a traditional file system interface for the your storage service instance in the cloud. It allows file operations with byte-level granularity, such as append, re-write, over-write, and truncate. The storage service supports file operations with whole-file granularity. As a result, when a file is modified in a filesystem on Oracle Cloud Infrastructure Storage Software Appliance, it results in a new version of the file being created and uploaded to the storage cloud.

When a file that contains multiple versions exists, the latest or most recent version of the file will always be returned when the file is read. The administrator can configure the number of versions of a file that will be retained in your storage service instance in the cloud. File Version Compaction allows the permanent deletion of unwanted versions. Also, if a file is deleted from the filesystem, then the corresponding object(s) in your service instance will also be deleted, if file version compaction is enabled in the appliance.

Quick Access to Select Files with Cache Pinning

Oracle Cloud Infrastructure Storage Software Appliance allows you to pin select files to the filesystem cache for quick access. You can pin files to the cache for filesystems connected to any storage class, `Standard` or `Archive`. (Archive support not available on Oracle Cloud at Customer)

When you write a file to your filesystem, it's initially stored in the filesystem cache, and then uploaded to the container on the storage service. After a file has been uploaded to the container, it may get removed from the filesystem cache by the cache manager. The cache is reclaimed using the *Least Recently Used* (LRU) cache management policy to meet the cache threshold that's specified in the filesystem advanced settings. If you want specific files to be always available in the cache for quick access, you can preserve them in the filesystem cache by pinning them to the cache. Once pinned, the files are not removed from the filesystem cache, except if you specifically unpin them.

Appliance Health Check

The appliance health check service is an automated process run on Oracle Cloud Infrastructure Storage Software Appliance. You can monitor the overall system status through the health check and get insights on the appliance performance like local storage usage.

Terminology

The following table defines the key terms used in the context of Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution.

Term	Description
Block storage	**Block storage** is an abstraction that is used by storage volumes and most storage devices such as hard disks, flash drives, and tape. A block is typically a range of contiguous bytes, and a volume is a range of contiguous blocks. Storage protocols such as SCSI, iSCSI, and Fibre Channel provide a method for accessing block storage devices that are attached either locally or remotely. File systems, some databases, and other storage applications use block I/O to access storage volumes to optimize performance and data granularity. Block storage optimizes storage for IOPS and block-based access and provides POSIX-compliant file systems for Oracle Cloud Infrastructure Compute Classic instances. It is limited in terms of scalability and does not support the definition of granular metadata for stored data.
Container	A **container** is a user-created resource in Oracle Cloud Infrastructure Object Storage Classic. It can hold an unlimited number of objects, unless you specify a quota for the container. Note that containers cannot be nested.
FileSystem (or filesystem)	A **FileSystem** in Oracle Cloud Infrastructure Storage Software Appliance connects a directory on the Oracle Cloud Infrastructure Compute Classic instance that hosts the appliance to a container on the storage service. Generally, *file system* (two words) means the mechanism that operating systems use to manage files on disks. This general meaning is distinct from the meaning of *filesystem* (one word) in the context of Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution.
NFS v4	**NFS v4** is version 4 of NFS (network file system), a distributed file system protocol defined in RFC 3530 (https://www.ietf.org/rfc/rfc3530.txt). It enables client computers to mount file systems that exist on remote servers and access those remote file systems over the network as though they were local file systems.
Object storage	**Object storage** provides an optimal blend of performance, scalability, and manageability when storing large amounts of unstructured data. Multiple storage nodes form a single, shared, horizontally scalable pool in which data is stored as objects (blobs of data) in a flat hierarchy of containers. Each object stores data, the associated metadata, and a unique ID. You can assign custom metadata to containers and objects, making it easier to find, analyze, and manage data. Applications use the unique object IDs to access data directly via REST API calls. Object storage is simple to use, performs well, and scales to a virtually unlimited capacity.
Oracle Cloud Infrastructure Compute Classic	**Oracle Cloud Infrastructure Compute Classic** is a secure, reliable, low cost, standards-based infrastructure service that you can use to rapidly provision virtual machines on Oracle Cloud with all the necessary storage and networking resources, manage and scale your virtual machine topology in the cloud easily, and migrate your Oracle and third-party applications to Oracle Cloud. **See Also**: Oracle Cloud Infrastructure Compute Classic Terminology in *Using Oracle Cloud Infrastructure Compute Classic.*
Oracle Cloud Infrastructure Object Storage Classic	**Oracle Cloud Infrastructure Object Storage Classic** provides a low cost, reliable, secure, and scalable object-storage solution for storing unstructured data and accessing it anytime from anywhere. It is ideal for data backup, archival, file sharing, and storing large amounts of unstructured data like logs, sensor-generated data, and VM images.

Workflow for Setting Up the Appliance

The following table outlines the tasks to get you started with Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution.

Task	More Information
1. Obtain the appliance image, provisioning tool, and configuration template.	Obtaining the Storage Software Appliance Image
2. Prepare the Oracle Cloud Infrastructure Compute Classic environment.	Preparing the OCI Compute Classic Environment
3. Define the parameters of the Oracle Cloud Infrastructure Compute Classic instance that will host Oracle Cloud Infrastructure Storage Software Appliance.	Preparing the Appliance Configuration File
4. Using an Oracle-provided machine image and a provisioning script, launch an Oracle Cloud Infrastructure Compute Classic instance with Oracle Cloud Infrastructure Storage Software Appliance pre-installed on the instance. Oracle Cloud Infrastructure Storage Software Appliance is a cloud storage gateway that you can use to access your storage service instance in the cloud over the NFSv4 protocol.	Provisioning the Appliance
5. Create filesystems on Oracle Cloud Infrastructure Storage Software Appliance. A **FileSystem** in Oracle Cloud Infrastructure Storage Software Appliance connects a directory on the Oracle Cloud Infrastructure Compute Classic instance that hosts the appliance to a container on the storage service. During the appliance provisioning process, you can opt to create the first filesystem by using the provisioning tool. After provisioning the appliance, you can add filesystems at any time through the management console of the appliance.	Managing FileSystems
7. Mount the Oracle Cloud Infrastructure Storage Software Appliance filesystems on the Oracle Cloud Infrastructure Compute Classic instances (that is, *client* instances) for which you want to provide shared storage over NFS v4.	Mounting Appliance FileSystems on Client Instances

2 Obtaining the Storage Software Appliance Image

To provision Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution, you must first upload and register the appliance image in Oracle Cloud Infrastructure Compute Classic. Then you must specify the appliance parameters in a configuration file and use a provisioning tool to create an Oracle Cloud Infrastructure Compute Classic instance with the appliance deployed on the instance.

1. Go to the Oracle Cloud Downloads page at http://www.oracle.com/technetwork/topics/cloud/downloads/index.html.
2. Look for **Oracle Cloud Infrastructure Storage Software Appliance**, and click that link.
3. Look for **Download Oracle Cloud Infrastructure Storage Software Appliance** and click that link.
4. On the resulting page, accept the license agreement.
5. Download the `tar.gz` file that contains the appliance image, configuration template, and provisioning script for **Cloud Distribution**.
6. Extract the contents of the `tar.gz` file to a local directory of your choice.
7. Copy the configuration template (`fscs_conf_template.yml`) and provisioning script (`fscs.sh`) to the host on which you plan to run the appliance provisioning tool. Provide *Execute* permission to `fscs.sh` with `chmod` command. See About the Appliance Provisioning Tool.
8. Upload the appliance image file (`OSCSA-CD-`*`version`*`.tar.gz`) to the Oracle Cloud Infrastructure Object Storage Classic account that's associated with your Oracle Cloud Infrastructure Compute Classic site.

   See **Uploading Machine Image Files to Oracle Cloud Infrastructure Object Storage Classic** in *Using Oracle Cloud Infrastructure Compute Classic*.
9. Register the appliance image in Oracle Cloud Infrastructure Compute Classic.

   See **Registering a Machine Image in Oracle Cloud Infrastructure Compute Classic** in *Using Oracle Cloud Infrastructure Compute Classic*.

Important:

Make a note of the name that you specify for the machine image while registering it, as shown in the following example:

Next Step: Preparing the OCI Compute Classic Environment

3

Preparing the OCI Compute Classic Environment

Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution provides Oracle Cloud Infrastructure Compute Classic instances shared access to file-based storage in the cloud over NFSv4. Before setting up the appliance, complete the following preparatory tasks in Oracle Cloud Infrastructure Compute Classic.

1. Generate the necessary SSH key pairs.

   To access the Oracle Cloud Infrastructure Compute Classic instance (on which Oracle Cloud Infrastructure Storage Software Appliance is deployed) securely using SSH, you must generate SSH key pairs and upload the public keys to Oracle Cloud Infrastructure Compute Classic.

   See Generating an SSH Key Pair in *Using Oracle Cloud Infrastructure Compute Classic* and the following example:

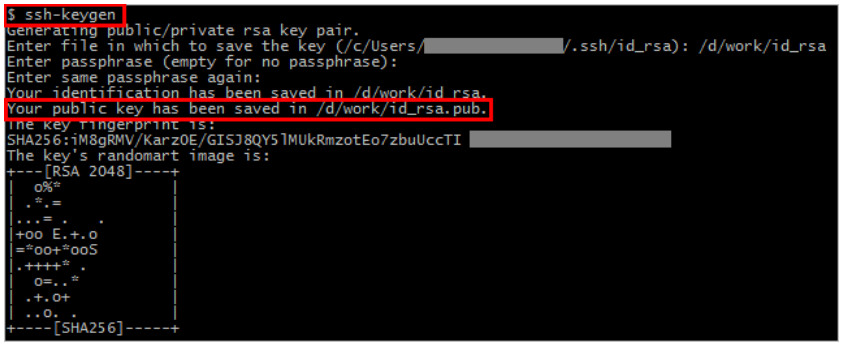

2. Add the public keys to Oracle Cloud Infrastructure Compute Classic.

   See Adding an SSH Public Key in *Using Oracle Cloud Infrastructure Compute Classic* and the following example:

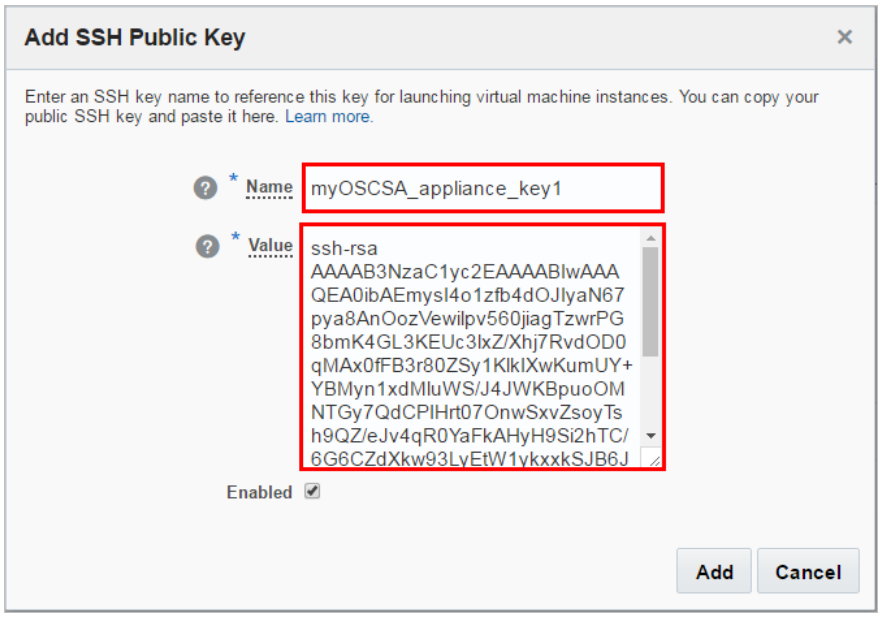

3. Identify (or create) a *security list* in Oracle Cloud Infrastructure Compute Classic for the appliance instance.

   A security list is a firewall for one or more Oracle Cloud Infrastructure Compute Classic instances. Once set up, you can open specific ports to instances in the security list by using security rules. You must identify (or create) a security list for the Oracle Cloud Infrastructure Compute Classic instance that will host Oracle Cloud Infrastructure Storage Software Appliance.

   > **Tip:**
   >
   > Your Oracle Cloud Infrastructure Compute Classic site has a predefined security list named `/Compute-`*`identity_domain`*`/default/default`. There may also be other security lists created by you or other users in the account. You can add the appliance instance to any of these security lists. But for better security and to ensure that you don't inadvertently interfere with the security settings of other users, create and use a new security list.

   See Creating a Security List in *Using Oracle Cloud Infrastructure Compute Classic* and the following example:

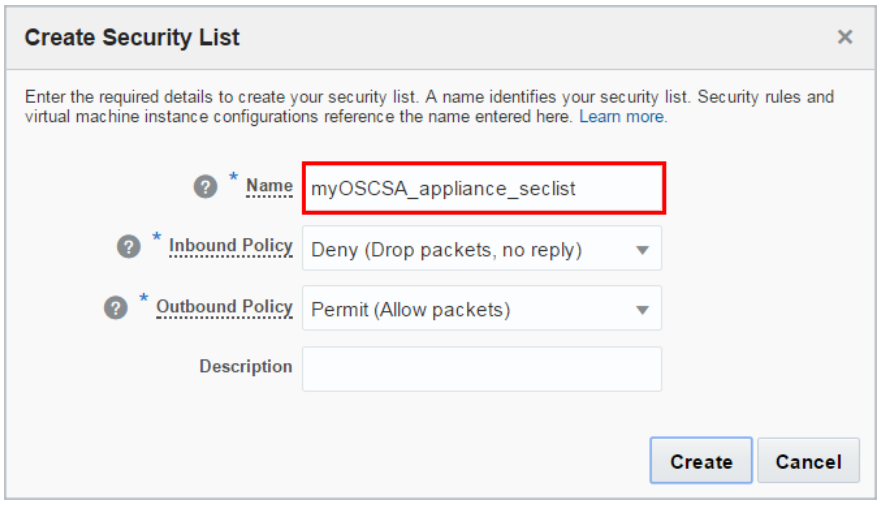

4. To enable NFS access to the appliance instance, create a *security application* in Oracle Cloud Infrastructure Compute Classic with `protocol`=TCP and `port`=2049.

   See Creating a Security Application in *Using Oracle Cloud Infrastructure Compute Classic* and the following example:

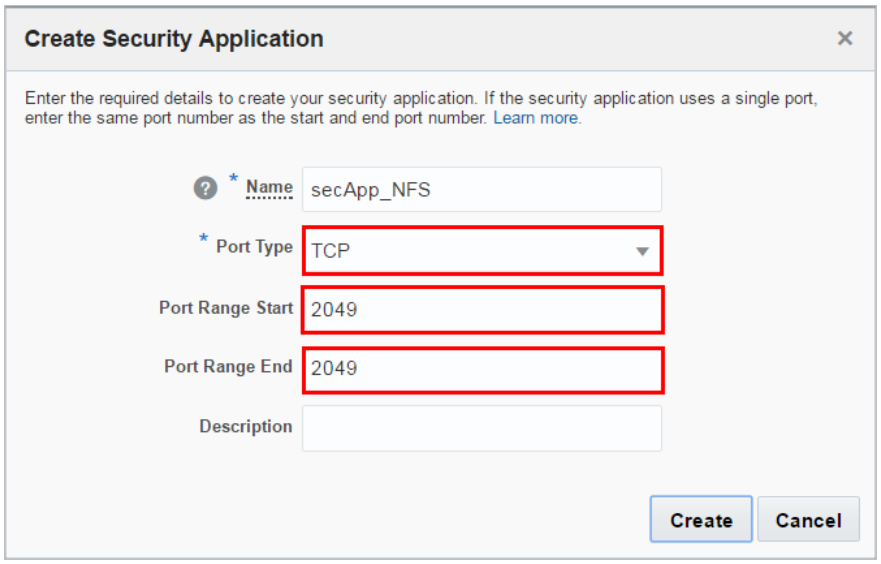

5. Create the following security rules in Oracle Cloud Infrastructure Compute Classic to open the required ports on the appliance instance, as described in Creating a Security Rule in *Using Oracle Cloud Infrastructure Compute Classic*.

   - To permit HTTPS traffic to the management console of Oracle Cloud Infrastructure Storage Software Appliance, create a security rule with the following settings:
     - **Security application**: `/oracle/public/https`
     - **Source**: To permit access from any host external to your Oracle Cloud Infrastructure Compute Classic site, specify the predefined security IP list, `/oracle/public/public-internet`, as the source. To permit access for only selected hosts, create a security IP list containing those hosts and

specify that security IP list as the source. See Creating a Security IP List in *Using Oracle Cloud Infrastructure Compute Classic*.

– **Destination**: The security list that you identified (or created) in step 3.

Example:

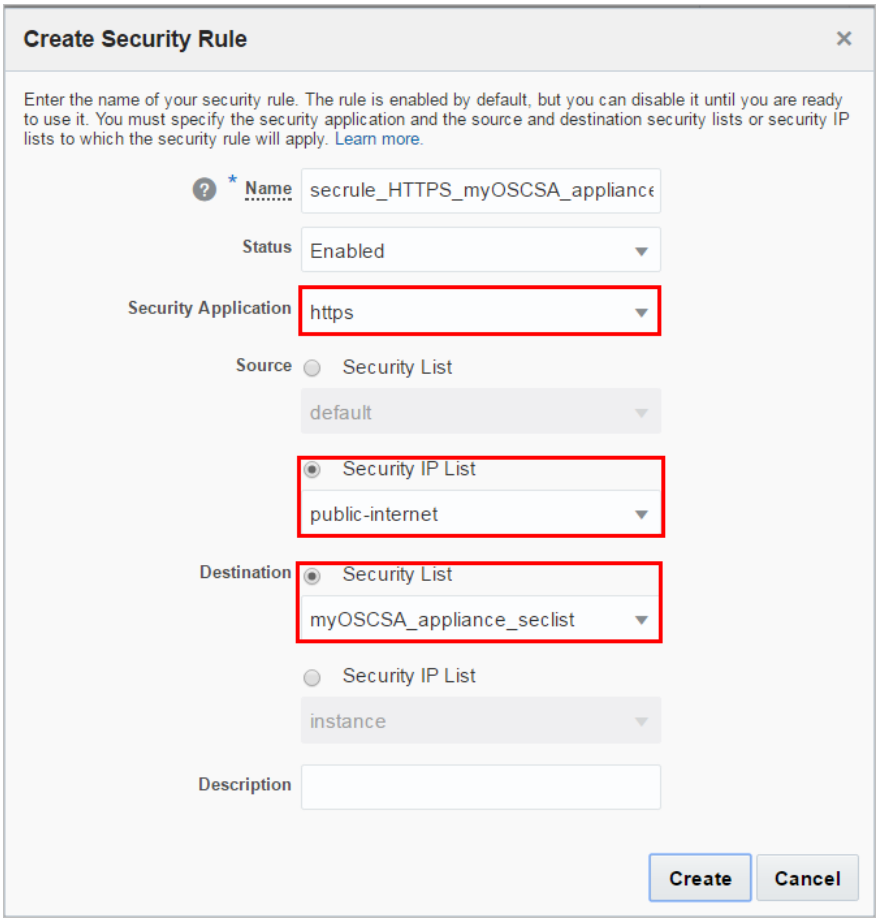

- To permit NFS access from your client Oracle Cloud Infrastructure Compute Classic instances to the appliance instance, create a security rule with the following settings:
  – **Security application**: The security application that you created in step 4 for port 2049.
  – **Source**: The security list containing the client instances for which you want to provide shared storage.
  – **Destination**: The security list that you identified (or created) in step 3.

  Example:

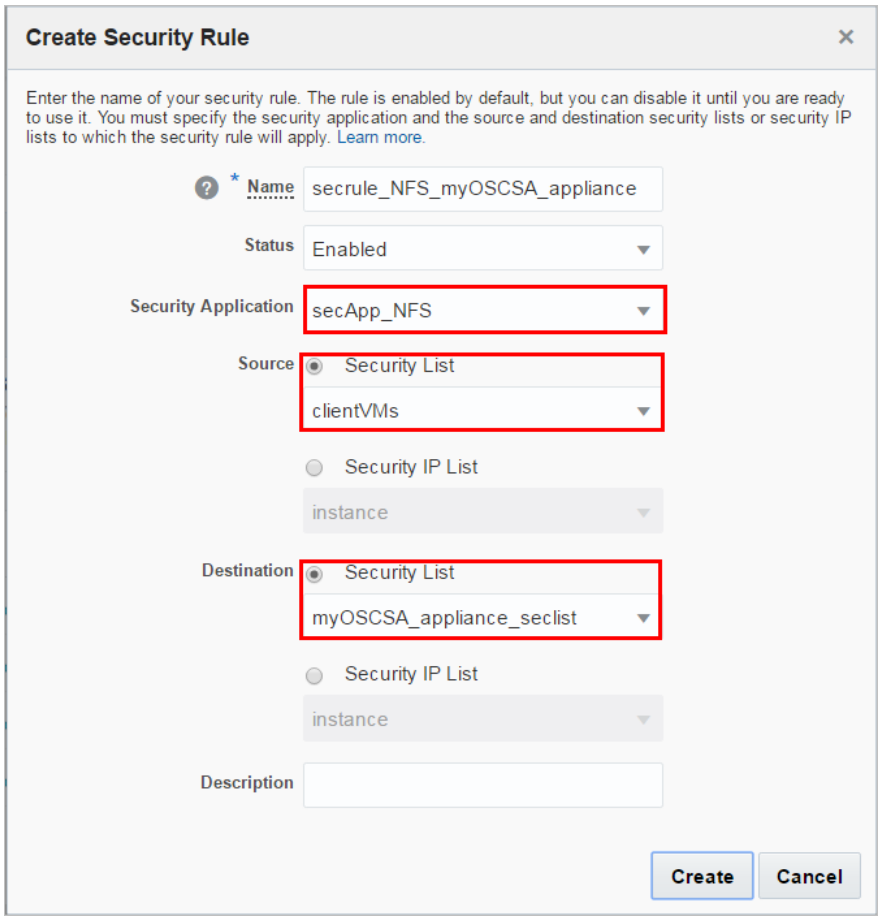

- To permit SSH connections to the appliance instance, create a security rule with the following settings:
  - **Security application**: The predefined security application, `/oracle/public/ssh`
  - **Source**: To permit access from any host external to your Oracle Cloud Infrastructure Compute Classic site, specify the predefined security IP list, `/oracle/public/public-internet`, as the source. To permit access for only selected hosts, create a security IP list containing those hosts and specify that security IP list as the source. See Creating a Security IP List in *Using Oracle Cloud Infrastructure Compute Classic*.
  - **Destination**: The security list that you identified (or created) in step 3.

  Example:

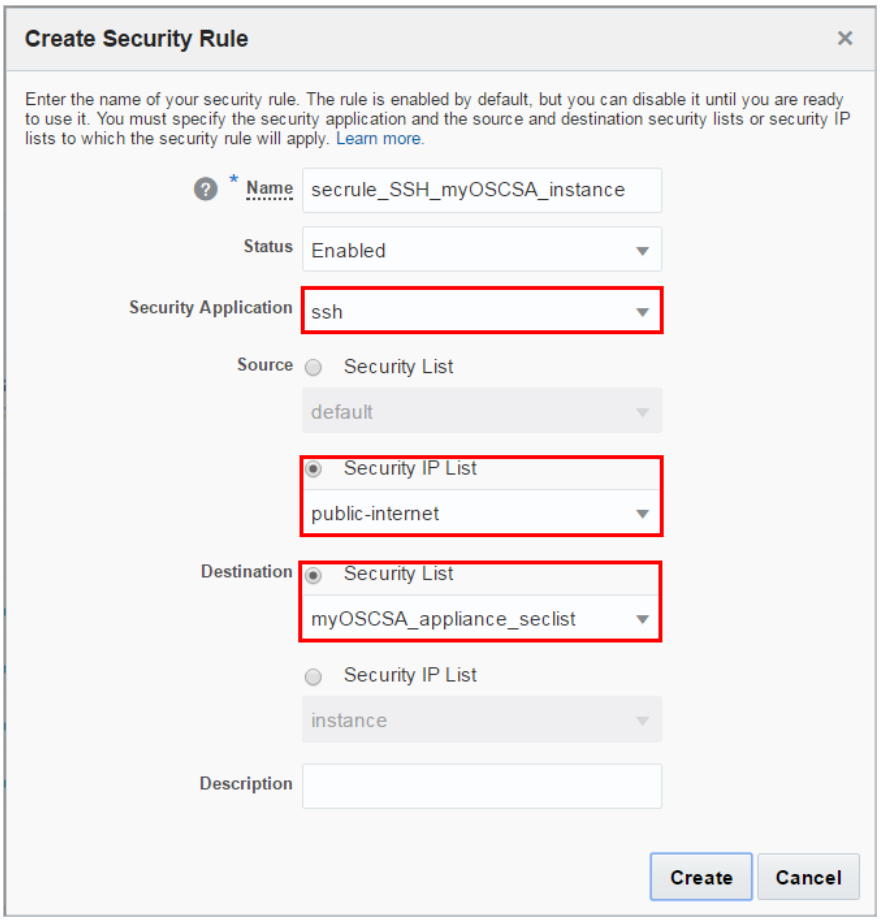

Next Task: Preparing the Appliance Configuration File

4

Preparing the Appliance Configuration File

The appliance configuration file specifies the parameters of the appliance and the Oracle Cloud Infrastructure Compute Classic instance that hosts the appliance.

Before You Begin

- Make sure that you've generated the necessary SSH key pairs, added the public keys to Oracle Cloud Infrastructure Compute Classic, and configured the required network settings in Oracle Cloud Infrastructure Compute Classic. See Preparing the OCI Compute Classic Environment.
- Identify the shape that you'd like to use for the appliance instance. Select **`oc2m`** or a larger shape.

  The shape of an instance determines the number of CPU cores and the amount of memory that's available for the instance. For information about the available shapes, see About Machine Images and Shapes in *Using Oracle Cloud Infrastructure Compute Classic*.
- Decide the number and size of the data disks (storage volumes) that you want to attach to the appliance instance.

  You can attach up to nine data disks, each up to 2 TB in size.
- Keep the following information ready:
  - The REST endpoint for your Oracle Cloud Infrastructure Compute Classic site. To find out the REST endpoint URL for your account, see *REST API for Oracle Cloud Infrastructure Compute Classic*.
  - The identity domain name, user name, and password for your Oracle Cloud Infrastructure Compute Classic site.

    > **Note:**
    >
    > The user must have the `Compute_Operations` role. If this role isn't assigned to the user or you're not sure, then ask your Oracle Cloud Infrastructure Compute Classic administrator to ensure that the role is assigned to the user in Oracle Cloud My Services. See Managing User Roles in *Managing and Monitoring Oracle Cloud*.

    The account-creation email from Oracle would contain the identity domain name and user name as shown in the following example:

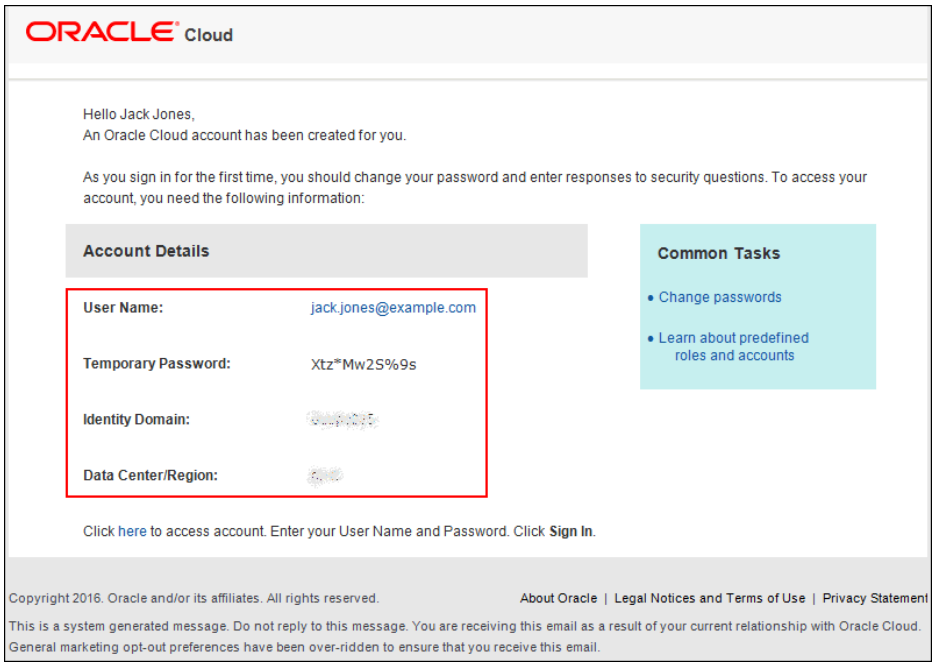

If you don't have this information, contact your service administrator.

Steps

To prepare the appliance configuration file:

1. Make sure that you have the latest version of the configuration file template.

   See Obtaining the Storage Software Appliance Image.
2. Open the configuration file template in a text editor, and specify the required parameters. See Appliance Configuration Parameters.
3. Validate the YAML format of the configuration file by using a tool such as YAML Lint (http://www.yamllint.com).

> **Note:**
>
> Oracle does not support or sponsor any third-party YAML validation tool.

Next Task

Provisioning the Appliance

Appliance Configuration Parameters

The appliance provisioning tool takes a configuration file, in YAML format, as an argument. In the configuration file, you must specify the parameters of the Oracle Cloud Infrastructure Compute Classic instance and Oracle Cloud Infrastructure Storage Software Appliance.

Topics

- Configuration File Template

Configuration File Template

```
oracle-compute:
  endpoint: compute_endpoint
  user:
    name: compute_user_name
  orchestration:
    name: orchestration_name
  vm:
    name: compute_instance_name
    hostname: instance_hostname
    shape: shape
    image: image
    sshkeys: list_of_keys
    storage:
      - datadisk:
          name: storage_volume_name
          size: size
          property: property
      - datadisk:
          name: storage_volume_name
          size: size
          property: property
      - ...<up to nine data disks>
    networking:
      name: ip_reservation_name
      model: NIC_mode
      securitylists: list_of_securitylists
      ippool: ip_pool_name
```

Configuration Parameters

The following table lists the parameters in the appliance configuration template.

> **Note:**
>
> In the **Parameter** column of the table, the parameters are indented to indicate their hierarchy in the YAML format.

Parameter	Description
`oracle-compute`	The parent parameter for the `endpoint` and `user` parameters.

Parameter	Description
`endpoint`	The REST endpoint URL for your Oracle Cloud Infrastructure Compute Classic site. **Example**: `endpoint: https://api-z13.compute.us2.oraclecloud.com` To find out the REST endpoint URL for your account: 1. Sign in to the Oracle Cloud My Services application. See Signing In to the My Services Application in *Managing and Monitoring Oracle Cloud*. The **My Services Dashboard** is displayed. It lists the services that are assigned to your account. 2. Look for **Oracle Cloud Infrastructure Compute Classic**. 3. From the **Actions**  menu, select **View Details**. The **Service Details** page is displayed. 4. Note the URL displayed in the **REST Endpoint** field under **Additional Information**.
`user`	The parent parameter for the `name` parameter.
`name`	The two-part Oracle Cloud Infrastructure Compute Classic user name, in the following format: `/Compute-`*`identity_domain/user_name`* The account-creation email from Oracle would contain the identity domain name and user name. **Note**: The user must have the `Compute_Operations` role. If this role isn't assigned to the user or you're not sure, then ask your Oracle Cloud Infrastructure Compute Classic administrator to ensure that the role is assigned to the user in Oracle Cloud My Services. See Managing User Roles in *Managing and Monitoring Oracle Cloud*. **Example**: `name: /Compute-acme/jack.jones@example.com`
`orchestration`	The parent parameter for the `name` parameter.
`name`	The three-part name of the Oracle Cloud Infrastructure Compute Classic orchestration, in the following format: `/Compute-`*`identity_domain/user_name/name`* **Example**: `name: /Compute-acme/jack.jones@example.com/myOrch` **Tip**: Note down the name that you specify for the orchestration. If you happen to lose the configuration file, you'll need this name to retrieve the configuration parameters and use them to reconstruct the configuration file.

Parameter	Description
`vm`	Parent parameter for all the parameters of the Oracle Cloud Infrastructure Compute Classic instance.

`name`	The three-part name of the instance, in the following format:
	`/Compute-`*`identity_domain`*`/`*`user_name`*`/`*`name`*
	Example:
	`name: /Compute-acme/jack.jones@example.com/myInstance`
	Note that, internally, Oracle Cloud Infrastructure Compute Classic generates an ID for the instance. The full name of an Oracle Cloud Infrastructure Compute Classic instance is the three-part name that you specify followed by the instance ID, as shown in the following example:
	`/Compute-acme/jack.jones@example.com/myInstance/300a7479-ec90-4826-98b9-a725662628f1`

`hostname`	The host name prefix for the instance. The name must not contain periods. The name you specify here is prefixed to the domain name (example: `compute-acme.oraclecloud.com`) to derive the full host name of the instance. **Example**: `hostname: myAppliance`
`shape`	The name of the shape, which defines the number of CPUs and the RAM that will be allocated to the instance. **Example**: `shape: oc2m` For information about the available shapes, see About Machine Images and Shapes in *Using Oracle Cloud Infrastructure Compute Classic*.
`image`	The three-part name of the machine image that must be used to create the Oracle Cloud Infrastructure Compute Classic instance, in the following format: `/Compute-`*`identity_domain`*`/user/`*`image_name`* Here `image_name` is the name that you noted earlier while registering the image. **Example**: `image: /Compute-acme/jack.jones@example.com/OSCSA-CD-image-v1.3` To get a list of the available machine images, send the `GET /imagelist/Compute-`*`identity_domain`*`/user` HTTP request to your Oracle Cloud Infrastructure Compute Classic site. See *REST API for Oracle Cloud Infrastructure Compute Classic*.

Parameter	Description
`sshkeys`	A comma-separated list of the SSH public keys that you want to associate with the instance. For each key, specify the three-part name in the `/Compute-`*`identity_domain`*`/`*`user`*`/`*`key`* format. You should have previously generated these keys and added them to Oracle Cloud Infrastructure Compute Classic. See Preparing the OCI Compute Classic Environment. **Example**: `sshkeys: ["/Compute-acme/jack.jones@example.com/ssh_key1","/Compute-acme/jane.williams@example.com/ssh_key2"]`
`storage`	The parent parameter for one or more `datadisk` parameters.
`datadisk`	Each `datadisk` parameter specifies the details of one disk, which is a storage volume in Oracle Cloud Infrastructure Compute Classic. The data disks are attached to the appliance instance. They provide local cache capacity for the appliance. See Best Practices – Configuring Cache Storage and Best Practices – Determining the Cache Size. All the data disks that are attached to the appliance instance are combined into a resizeable logical volume group, named `oracle_data_vg`, which is mounted on the appliance instance at `/opt/oracle/datadisk`. **Note**: At any time, you can add more data disks. See Adding Data Disks to the Appliance Instance. You can attach up to nine data disks, each up to 2 TB in size.
`name`	The three-part name of the storage volume, in the following format: `/Compute-`*`identity_domain`*`/`*`user_name`*`/`*`volume_name`* **Note**: Don't specify the name of a storage volume that already exists in Oracle Cloud Infrastructure Compute Classic. **Example**: `name: /Compute-acme/jack.jones@example.com/myVolume1`

Parameter	Description
`size`	The size of the storage volume. Use one of the following abbreviations for the unit of measurement: • B or b for bytes • K or k for kilobytes • M or m for megabytes • G or g for gigabytes • T or t for terabytes The allowed range is from 1 GB to 2 TB, in increments of 1 GB. **Example**: `size: 10G`
`property`	The three-part name of the storage property to be used for creating the storage volume. • For storage volumes that require low latency and fast I/O, such as for storing database files, specify `/oracle/public/storage/latency`. • For all other storage volumes, select `/oracle/public/storage/default`. **Example**: `property: /oracle/public/storage/default`
`networking`	The parent parameter for the networking parameters for the Oracle Cloud Infrastructure Compute Classic instance.
`name`	The three-part name of the reservation that should be created in Oracle Cloud Infrastructure Compute Classic for the public IP address of the instance, in the following format: `/Compute-`*`identity_domain`*`/`*`user_name`*`/`*`ip_reservation_name`* **Example**: `name: /Compute-acme/jack.jones@example.com/myIPres`
`model`	The type of network interface card (NIC). Specify `e1000` as the value.
`securitylists`	A comma-separated list of the security lists that you want to add the instance to. You must have already identified (or created) these security lists. See Preparing the OCI Compute Classic Environment. **Example**: `securitylist: ["/Compute-acme/jack.jones@example.com/mySeclist1","/Compute-acme/jack.jones@example.com/mySeclist2"]`
`ippool`	The pool of public IP addresses from which an address must be reserved for the Oracle Cloud Infrastructure Compute Classic instance. Specify `/oracle/public/ippool` as the value.

Sample Configuration File

```
oracle-compute:
  endpoint: https://api-z13.compute.us2.oraclecloud.com
  user:
    name: /Compute-acme/jack.jones@example.com
  orchestration:
    name: /Compute-acme/jack.jones@example.com/myOrch
  vm:
    name: /Compute-acme/jack.jones@example.com/myInstance
    hostname: myAppliance
    shape: oc2m
    image: /Compute-acme/jack.jones@example.com/OSCSA-CD-image-v1.3
    sshkeys: ["/Compute-acme/jack.jones@example.com/ssh_key","/Compute-
acme/jane.williams@example.com/ssh_key"]
    storage:
      - datadisk:
          name: /Compute-acme/jack.jones@example.com/myVolume1
          size: 10G
          property: /oracle/public/storage/default
      - datadisk:
          name: /Compute-acme/jack.jones@example.com/myVolume2
          size: 10G
          property: /oracle/public/storage/default
    networking:
      name: /Compute-acme/jack.jones@example.com/myIPres
      model: e1000
      securitylist: ["/Compute-acme/jack.jones@example.com/mySeclist1","/
Compute-acme/jack.jones@example.com/mySeclist2"]
      ippool: /oracle/public/ippool
```

5
Provisioning the Appliance

Topics

About the Appliance Provisioning Tool

Use the appliance provisioning tool to quickly launch a Oracle Cloud Infrastructure Compute Classic instance (virtual machine) with Oracle Cloud Infrastructure Storage Software Appliance pre-installed on the instance, and with all the required networking and storage settings.

The appliance provisioning tool is a Bash shell script (`fscs.sh`). The script takes a configuration file, in YAML format, as an argument. In the configuration file, you specify the parameters of the Oracle Cloud Infrastructure Compute Classic instance and Oracle Cloud Infrastructure Storage Software Appliance. When you run the script, it validates the parameters in the configuration file and uses the parameters to create an Oracle Cloud Infrastructure Compute Classic instance with Oracle Cloud Infrastructure Storage Software Appliance deployed on the instance.

Supported Platforms

The appliance provisioning tool is currently supported on Oracle Linux 6.6.

Download Location

See Obtaining the Storage Software Appliance Image.

Command Syntax

```
./fscs.sh [-c|-r|-u|-d|-dd|-i|-l|-a|-v config_file]
```

Command Parameters

Parameter	Description
`-c`	Creates an Oracle Cloud Infrastructure Compute Classic instance with Oracle Cloud Infrastructure Storage Software Appliance pre-installed on the instance. See Creating the Appliance.
`-r`	Deletes and then re-creates the Oracle Cloud Infrastructure Compute Classic instance with Oracle Cloud Infrastructure Storage Software Appliance pre-installed on the instance. See Re-creating the Appliance.

Parameter	Description
`-u`	Upgrades the operating system of the Oracle Cloud Infrastructure Compute Classic instance. See Upgrading the Appliance.
`-d`	Deletes the Oracle Cloud Infrastructure Compute Classic instance. See Deleting the Appliance.
`-dd`	Deletes the Oracle Cloud Infrastructure Compute Classic instance and the associated resources in Oracle Cloud Infrastructure Compute Classic (that is, the public IP address reservation, all the data disks, and the boot disk). See Deleting the Appliance.
`-i`	Retrieves the public and private IP addresses of the Oracle Cloud Infrastructure Compute Classic instance.
`-l`	Downloads a system dump of the Oracle Cloud Infrastructure Compute Classic instance.
`-a`	Adds one or more data disks. See Adding Data Disks to the Appliance Instance.
`-v`	Returns all the configuration properties of the appliance in a JSON-formatted text file. If you contact Oracle for support, you may be asked to provide this file for diagnostics purposes.
`config_file`	The full path and name of the plain-text file, in YAML format, that specifies the parameters of the Oracle Cloud Infrastructure Compute Classic instance and Oracle Cloud Infrastructure Storage Software Appliance. You must specify this argument with every option of the script. See Appliance Configuration Parameters.

Creating the Appliance

This section describes how to create the appliance using the appliance provisioning tool and the appliance configuration file.

Before You Begin

1. Ensure that the appliance configuration file is ready. See Preparing the Appliance Configuration File.
2. Ensure that you have the latest version of the appliance provisioning tool. See Obtaining the Storage Software Appliance Image.
3. Provide *Execute* permission to `fscs.sh` with `chmod` command.

Steps

1. On the host on which you downloaded the appliance provisioning tool, go to the directory that contains the tool, and run the following command:

   ```
   ./fscs.sh -c config_file
   ```

   In this command, `config_file` is the full path and name of the appliance configuration file.

2. The tool prompts you to enter the password for the Oracle Cloud Infrastructure Compute Classic user specified in the configuration file. Enter the password.

   The tool validates your credentials, and then does the following:

   a. Validates the parameters in the configuration file.
   b. Creates a boot disk.
   c. Creates the data disks specified in the configuration file.
   d. Reserves a fixed public IP address for the appliance instance.
   e. Creates an Oracle Cloud Infrastructure Compute Classic instance with Oracle Cloud Infrastructure Storage Software Appliance deployed on the instance.
   f. Starts Oracle Cloud Infrastructure Storage Software Appliance.

   This process may take a few minutes.

   Wait till you see a message like the following example:

   ```
   Creating Oracle Cloud Infrastructure Storage Software Appliance
   instance ... successful
   Extracting IP addresses ...
   Private IP address: 10.196.35.241
   Public IP address: 203.0.113.49
   ```

3. Note the private and public IP addresses.

   > **Tip:**
   >
   > At any time later, you can retrieve the IP addresses as described in Finding Out the IP Addresses of the Appliance Instance.

4. If the tool prompts you to confirm whether you want to create your first filesystem while creating the appliance, enter `n`. This capability is deprecated and is not supported on Oracle Cloud Infrastructure Storage Software Appliance.

After the installation is complete, you can perform any of the operations for managing the appliance. See Managing the Appliance.

> **Tip:**
>
> Retain a copy of the configuration file. You'll need it to run any further operations on the appliance.

Next Tasks

1. Log in to the management console of the appliance to create, manage, and monitor filesystems. See Logging In to the Management Console of the Appliance.
2. On each of the client Oracle Cloud Infrastructure Compute Classic instances that you want to provide shared storage for, mount the filesystem. See Mounting Appliance FileSystems on Client Instances.

6
Managing FileSystems

Topics

Logging In to the Management Console of the Appliance

Use the management console of the appliance to create, manage, and monitor filesystems.

To access the management console:

1. Using your web browser, go to the following URL:

   `https://appliance_ip`

   Here, `appliance_ip` is the public IP address of the appliance instance. See Finding Out the IP Addresses of the Appliance Instance.

> **Note:**
>
> Your browser may display a warning that the SSL certificate couldn't be verified. This warning is displayed because Oracle Cloud Infrastructure Storage Software Appliance uses a self-signed certificate for the HTTPS connection. If you've entered the correct public IP address of the appliance instance in the browser address bar, then you can safely ignore this warning. The steps to ignore this warning and proceed to the management console vary depending on the browser you use. For example, in Mozilla Firefox v38.4, you can proceed by clicking **Add Exception** on the **This Connection is Untrusted** page. In Google Chrome v46.0, click **Advanced** on the **Your connection is not private** page.

The login page is displayed.

2. Enter the password for the `admin` user.

> **Note:**
>
> If you're logging in to the management console of the appliance for the first time, you'll be prompted to set and confirm the password for the `admin` user of Oracle Cloud Infrastructure Storage Software Appliance. The password can contain from 8 to 32 characters, with at least one special character, one numerical character, one uppercase character, and one lowercase character.

> **Tip:**
>
> Remember the password. You must enter it every time you access the management console of Oracle Cloud Infrastructure Storage Software Appliance.
> If you don't remember the password, you can reset it. See Changing the Admin Password for the Appliance.

Creating Your First Filesystem

A **FileSystem** in Oracle Cloud Infrastructure Storage Software Appliance connects a directory on the Oracle Cloud Infrastructure Compute Classic instance that hosts the appliance to a container on the storage service.

> **Note:**
>
> If you've already created the first filesystem and now want to create additional filesystems, then see Adding a FileSystem.

Before You Begin

- Keep the following information ready:
  - The object storage API endpoint for the storage account to which you want to connect the filesystem.

    > **Note:**
    >
    > * To find the object storage API endpoint for your Oracle Cloud Infrastructure Object Storage Classic account, go to the Oracle Cloud My Services page and look for the **REST Endpoint** field under the **Additional Information** section. (Not available on Oracle Cloud at Customer)

  - The log in credentials for your storage account.

    The account-creation email from Oracle would contain this information.

    > **Note:**
    >
    > The user must have the `Storage Administrator` role. If the user doesn't have the required role or you aren't sure, then ask your service administrator to ensure that the user has the required role in Oracle Cloud My Services. See Managing User Roles in *Managing and Monitoring Oracle Cloud*.

- Ensure that a replication policy has been set. See Selecting a Replication Policy in *Using Oracle Cloud Infrastructure Object Storage Classic*. (Not available on Oracle Cloud at Customer)

Steps

1. Access the management console of the appliance:

   Using your web browser, go to the following URL:

   `https://appliance_ip`

   Here, `appliance_ip` is the public IP address of the appliance instance. See Finding Out the IP Addresses of the Appliance Instance.

2. You're prompted to set the password for the `admin` user of the appliance.

   Enter a password that complies with the password rules displayed on the page, and click **Save**.

   Remember the password. You must enter it every time you access the management console of Oracle Cloud Infrastructure Storage Software Appliance.

   The login page is displayed.

3. Enter the password for the `admin` user, and click **Log In**.

   A wizard to create the first filesystem starts.

4. Enter a name for the filesystem, and click **Next**.

> **Note:**
>
> For the character restrictions applicable in Oracle Cloud Infrastructure Object Storage Classic when you must select a file name or a filesystem name, see Character Restrictions.

5. Enter the object storage API endpoint of your storage account, and click **Next**.
6. Create a filesystem using your Oracle Cloud Infrastructure Object Storage Classic account details:
   - User name of your account (for example: `jack.jones@example.com`).
   - Password of your account.
   - Click **Validate**.
7. Set the following options as required:
   - **Enable Encryption**: If you would like to enable encryption of data in the filesystem before it is uploaded to Oracle Cloud Infrastructure Object Storage Classic, then select this check box. You can provide the encryption keys later.

     > **Note:**
     >
     > You can enable encryption for a filesystem only at the time of creating the filesystem, not later.

   - **Enable Archive** (Not available on Oracle Cloud at Customer): Select this check box if you want to connect this filesystem to a container of the `Archive` storage class.
     For more information, see About Archive FileSystems.
8. Click **Show Advanced**, and enter the required information in the advanced configuration fields. For more information, see the table in Adding a FileSystem.
9. Click **Save**.
   The filesystem is created and the details of the filesystem are displayed on the **Dashboard** page.

The filesystem that you just created is mounted on the appliance instance on a directory named `/mnt/gateway/filesystem_name_that_you_specified`.

Next Tasks:

1. If you enabled encryption for the filesystem, then provide the encryption keys. See Encrypting Data for a FileSystem.
2. Connect the filesystem that you just created to a container in your storage service instance in the cloud. See Connecting a FileSystem.
3. Mount the filesystem on the client instances. See Mounting Appliance FileSystems on Client Instances.

> **Note:**
>
> If required, you can create additional filesystems and mount them (see Adding a FileSystem).

About Archive FileSystems


This topic does not apply to Oracle Cloud at Customer.

When you create a filesystem of the `Archive` storage class, if a container by the same name doesn't exist in your account, then the container will be created. In addition, another container named *`filesystem_name`*`-archive` is created.

For example, if you create a filesystem with the name `myFirstArchiveFS` of the `Archive` storage class, then the following two containers are created:

- `myFirstArchiveFS-archive`
  This container is of the `Archive` storage class. All the files that are uploaded in the mounted directory on the NFS client are stored in the `myFirstArchiveFS` filesystem on the appliance host and then asynchronously copied to this container. Metadata backups are also stored in this container.
- `myFirstArchiveFS`
  Metadata synchronization objects are stored in this container. Metadata synchronization objects enable metadata information to be restored if the appliance installation is lost.

  In addition to metadata synchronization objects, the metadata backups are stored in an `Archive` container periodically.

If the filesystem name matches an existing container, and if the container is of the `Standard` storage class, then you cannot mount the filesystem and an error message is displayed.

The appliance automatically performs daily, weekly, and monthly metadata backups.

Adding a FileSystem

You can add one or more filesystems in Oracle Cloud Infrastructure Storage Software Appliance and connect each filesystem to a container in your account.

1. Log in to the management console.
   See Logging In to the Management Console of the Appliance.

   The available filesystems are displayed.
2. Click **Create Filesystem** in the navigation pane on the left.
   The **Create a FileSystem** page is displayed.
3. Enter the required information in the following fields:

Field	Description
FileSystem Name	Name of the filesystem. If a container by the same name doesn't exist in your account, then it will be created. Enter a name that is meaningful to you and unique. **Note**: • For the character restrictions applicable when you enter a filesystem name, see Character Restrictions. If a container with the same name as the filesystem already exists, and if that container isn't empty, then the data cached in the Oracle Cloud Infrastructure Storage Software Appliance filesystem may not be consistent with data stored in the container.
Object Storage API Endpoint	The object storage API endpoint for your service instance. • To find out the object storage API endpoint for your Oracle Cloud Infrastructure Object Storage Classic instance, see About REST URLs for Oracle Cloud Infrastructure Object Storage Classic Resources in *Using Oracle Cloud Infrastructure Object Storage Classic*. (Not available on Oracle Cloud at Customer) Go to Step 4a.
User Name	Oracle Cloud Infrastructure Object Storage Classic user name.
Account Password	Oracle Cloud Infrastructure Object Storage Classic password.

4. a. Click **Validate**.
   If any of the values entered above do not match your Oracle Cloud Infrastructure Object Storage Classic account credentials, then an error message is displayed. Recheck and enter the appropriate values in the respective fields.

   b. Click **Save**.

5. If you would like to enable encryption of data in the filesystem before it is uploaded to your account, then select the **Enable Encryption** check box. You can provide the encryption keys later. For more information, see Encrypting Data for a FileSystem.

> **Note:**
>
> You can enable encryption for a filesystem only at the time of adding the filesystem, not later.

> **Caution:**
>
> Ensure that you enable encryption in the new appliance filesystem when you claim ownership from another appliance filesystem (if encryption was enabled in the other appliance filesystem). Otherwise, the data uploaded from the new appliance filesystem may cause inconsistency in the container, as the container may contain encrypted data from the previous appliance filesystem.
> If a filesystem is deleted in the appliance, ensure that you enable encryption in the re-created filesystem.

6. If you want to create an `Archive` filesystem, then select the **Enable Archive** check box. (Not available on Oracle Cloud at Customer)

   For more information, see About Archive FileSystems .

   If the filesystem name matches an existing container, and if the storage class of the container is `standard`, then you cannot mount the filesystem and an error message is displayed.

7. Click **Show Advanced**, and enter the required information in the following fields:

Field	Description
NFS Allowed Hosts	The hosts allowed to connect to the NFS export. Example: `2001:db8:9:e54::/64`, `192.0.2.0/24`
NFS Export Options	The NFS export options. Example: `rw`, `sync`, `insecure`, `no_subtree_check`, `no_root_squash` Don't specify the `fsid` option.
Maximum Local Cache Size in GiB	The maximum number of bytes that can be cached. When the data in the cache reaches the specified limit or the cache is full, the appliance removes files from the cache based on a least recently used (LRU) algorithm. The files that are yet to be uploaded to the account are not removed from the cache. The preserved files (by cache pinning) are also not removed from the cache. See Configuring the Cache for FileSystems. **Note**: The number of files in cache is limited to 20,000, regardless of the specified cache size in bytes.
Concurrent Uploads	The number of concurrent uploads to the cloud. Allowed range: 1 to 30 This field indicates the maximum number of files that can be concurrently uploaded in the appliance. If the value is 5, the concurrent file uploads can be between 0-5.
Delete Old File Versions	If you select this check box for a filesystem, then every time you start the appliance and once every 24 hours after that, all the older versions of all the objects are deleted in the container that's connected to the filesystem. Only the latest version of each object is retained. For more information, see Enabling File Versions Compaction. This option is disabled by default.
Restore Object Retention	The number of days a file will remain in the restored state. **Note**: This field is displayed only for an `archive` filesystems. Default value: 10

Field	Description
Sync Policy	The metadata operations are flushed to the disk based on the following modes. Select one of the following modes based on your requirement: • **Asynchronous** In this mode, the filesystem operations are time-based and are persisted asynchronously. This mode offers the best performance. **Note**: This mode is not suitable for any filesystem operation that depends on synchronous transactions. • **Posix Standard** This mode is enabled by default. Only the synchronous transactions (like fsync, ODSYNC and OSYNC) are committed to the disk. All the other transactions are handled asynchronously. These sync modes do not affect on-disk consistency.
Cloud Read-ahead	The number of 1-MB blocks to be downloaded and used to *read ahead* when reading files. Use this setting to improve the read performance for large files that aren't cached. Default value: 0 (prefetching is disabled)

8. Click **Save**.

Next Task

Connect the filesystem. See Connecting a FileSystem.

Importing an Existing Filesystem

Prerequisite

Before you import an existing filesystem from another appliance, ensure that all the pending uploads of the files residing on the appliance which last owned the filesystem, are completed and the files are uploaded to your account.

> **Note:**
>
> There may be pending or interrupted file uploads in a failed appliance. If you're importing an existing filesystem from a failed appliance, you must re-create those files in the filesystem on the recovery appliance.

1. Log in to the management console.
   See Logging In to the Management Console of the Appliance.

   The available filesystems are displayed.
2. Click **Create FileSystem** in the navigation pane on the left.
   The **Create a FileSystem** page is displayed.
3. Enter the required information in the **Required** tab.
   For the filesystem name, enter the name of the existing filesystem that you want to import to this appliance.

4. Click **Validate**.
   If any of the values entered do not match your Oracle Cloud Infrastructure Object Storage Classic account credentials, then an error message is displayed. Recheck and enter the appropriate values in the respective fields.
5. Select the options that you'd like to enable in the filesystem.
   For example: **Enable Encryption**.
6. Click **Show Advanced**, and enter the required information.
7. Click **Save**.
   The filesystem is created and displayed on the **Dashboard** tab.
8. Click **Connect** for the filesystem that you want to import.
   If the filesystem that you're importing is connected to another appliance, then **FileSystem: Claim Ownership** window is displayed, prompting you to confirm whether the other appliance must be disconnected.

   If you opt to proceed, then re-enter your Oracle Cloud Infrastructure Object Storage Classic password and select **Claim Ownership**.

   A filesystem may be mounted for read/write on only one appliance at a time.

   If a container with the same name as the filesystem exists in your Oracle Cloud Infrastructure Object Storage Classic account, then the filesystem is connected to that container. If a container by that name doesn't exist, then it's created and the filesystem is connected to the container.

Next Task: Mount the filesystem on the required client instances. See Mounting Appliance FileSystems on Client Instances.

Configuring the Cache for FileSystems

Oracle Cloud Infrastructure Storage Software Appliance caches frequently retrieved data on the local host, minimizing the number of REST API calls to your service instance and enabling faster data retrieval. The appliance uses an upload buffer and a read cache for data storage and retrieval.

> **Note:**
>
> The cache resides on the data disks attached to the appliance instance. At any time, you can increase the disk space available for caching by adding disks. See Adding Data Disks to the Appliance Instance.

Topics

- About the FileSystem Cache
- Guidelines for Sizing and Configuring the Cache
- Configuring the FileSystem Cache

About the FileSystem Cache

The filesystem cache serves two roles: an upload buffer and a read cache. The upload buffer contains data that has been copied to the disk cache and is queued to be stored in your account. The read cache contains frequently retrieved data that's accessible locally for read operations.

When an application transfers a file through an NFS share, the file is queued to be stored in your account. The upload buffer might contain many files. If the host on which Oracle Cloud Infrastructure Storage Software Appliance is installed fails, or if the appliance stops abruptly, the pending upload operations are not lost because they are persisted on the local disk. When the appliance restarts, the pending upload operations resume and the data is stored in your account.

When you retrieve data, the data is stored in the read cache of the appliance. This allows subsequent I/O operations to that file to be done at local disk speed.

When the data in the cache reaches the specified limit or the cache is full, the appliance removes files from the cache based on a least recently used (LRU) algorithm. The files that are yet to be uploaded to the account are not removed from the cache. The preserved files (by cache pinning) are also not removed from the cache.

Guidelines for Sizing and Configuring the Cache

See Best Practices – Configuring Cache Storage and Best Practices – Determining the Cache Size.

Configuring the FileSystem Cache

You can configure the cache for a filesystem while adding the filesystem. See Adding a FileSystem.

Encrypting Data for a FileSystem

To ensure that your data remains protected both when it's stored in your service instance and while in transit, you can enable encryption of the data in Oracle Cloud Infrastructure Storage Software Appliance before it's stored in your service instance.

Topics

About FileSystem Encryption

Data encryption in Oracle Cloud Infrastructure Storage Software Appliance is done using a symmetric key, which is stored in a database on the appliance and is encrypted by using an asymmetric public-private key pair. Administrators can back up and store the asymmetric keys, and use them to recover encrypted data.

When a file is stored in Oracle Cloud Infrastructure Storage Software Appliance, it's first stored in the appliance cache in its original form. The file is encrypted before it's uploaded to your account. When a file is retrieved from your account, the data is decrypted while it's streamed to the appliance cache.

Enabling Encryption for a FileSystem

To enable encryption for a filesystem, you must select the **Enable Encryption** check box when you create the filesystem. See Adding a FileSystem.

> **Note:**
>
> You can enable encryption on a filesystem only when the filesystem is created and before it's connected to your account. You can't enable or disable encryption after a filesystem is created and connected.

The following message is displayed in the management console after you enable encryption for a filesystem:

```
Encryption is enabled
```

After enabling encryption for a filesystem, you can generate encryption keys and specify the keys in the management console. When required, you can change the encryption keys.

Generating Encryption Keys

1. Log in to any UNIX-like host that has the OpenSSL toolkit installed on it. See https://www.openssl.org/.
2. Generate an RSA private key:
   `openssl genrsa -out mykey.pem 2048`
3. Convert the RSA private key to the `pk8` format:
   `openssl pkcs8 -in mykey.pem -out privateKey.key -outform pem -nocrypt -topk8`
4. Generate the public key for the private key:
   `openssl rsa -in privateKey.key -out publicKey.key -pubout -outform pem`

> **Note:**
>
> Ensure that the encryption keys are backed up and saved on another host.

Specifying the Encryption Keys for a FileSystem

1. Log in to the management console.
   See Logging In to the Management Console of the Appliance.
2. Select the filesystem for which you want to specify encryption keys, and go to the **Encryption Keys** tab.
3. Copy the private key that you generated earlier, and paste it in the **Private Key** field in the **Set New Encryption Keys** section.
4. Copy the corresponding public key and paste it in the **Public Key** field in the **Set New Encryption Keys** section.
5. Click **Save New Keys**.

Retrieving the Encryption Keys of a FileSystem

1. Log in to the management console.
   See Logging In to the Management Console of the Appliance.
2. Select the filesystem for which you want to download encryption keys, and go to the **Encryption Keys** tab.
3. Click **Download** in the **Existing Encryption Keys** section.
4. Save the `.tar.gz` file (that contains the encryption keys) in a location of your choice.

Connecting a FileSystem

Connecting a FileSystem

After you create a filesystem, you must connect it to a container in your account before you can store and retrieve data through the filesystem.

> **Note:**
>
> During the appliance provisioning process, if you had opted to create the first filesystem by using the provisioning script, then that filesystem would have been connected automatically to your account. The steps described in this section aren't relevant to that filesystem.

> **Caution:**
>
> If a container with the same name as the filesystem already exists, and if that container isn't empty, then the data cached in the Oracle Cloud Infrastructure Storage Software Appliance filesystem may not be consistent with data stored in the container.

1. Log in to the management console of Oracle Cloud Infrastructure Storage Software Appliance.
   See Logging In to the Management Console of the Appliance.
2. On the **Dashboard** tab, identify the filesystem that you want to connect to your account.
3. Click **Connect**.
   If a container with the same name as the filesystem exists in Oracle Cloud Infrastructure Object Storage Classic, then the filesystem is connected to that container. If a container by that name doesn't exist, then it's created and the filesystem is connected to the container.

> **Note:**
>
> A filesystem may be mounted for read/write on only one appliance at a time.
>
> - If the filesystem name that you've specified matches the name of an existing container in your Oracle Cloud Infrastructure Object Storage Classic account, and if that container is connected to another appliance filesystem, then **FileSystem: Claim Ownership** window is displayed, prompting you to confirm whether the other filesystem must be disconnected. If you opt to proceed, then you must re-enter your Oracle Cloud Infrastructure Object Storage Classic password and select **Claim Ownership**.
>
> This check ensures that you don't inadvertently connect the new filesystem to a container that's already connected to another appliance filesystem.

Next Task: Mount the filesystem on the required client instances. See Mounting Appliance FileSystems on Client Instances.

Disconnecting a FileSystem

To disconnect a filesystem, select it on the dashboard of the Oracle Cloud Infrastructure Storage Software Appliance management console, and click **Disconnect**.

The container to which the filesystem was previously connected and the stored data remain intact even after the filesystem is disconnected.

At any time, you can resume storing and retrieving data through the filesystem by connecting it again. If you no longer need the disconnected filesystem, then you can delete it. See Deleting a FileSystem.

Mounting Appliance FileSystems on Client Instances

Each filesystem in Oracle Cloud Infrastructure Storage Software Appliance maps a directory on the appliance host to a container in your storage service instance in the cloud. To enable your *client* Oracle Cloud Infrastructure Compute Classic instances to access shared storage through the appliance, you must mount the appliance filesystems on the client Oracle Cloud Infrastructure Compute Classic instances.

Supported NFS Clients

- Oracle Linux 6.4, 6.6, and 7.3
- Ubuntu 14.04 and 16.04
- CentOS 7
- Debian 8 Jessie

Before You Begin

- Note the names of the filesystems that you want to mount.

  The appliance management console displays a list of all the available filesystems. See Logging In to the Management Console of the Appliance.

- Make sure that the filesystems that you want to mount are connected to your storage service instance in the cloud. See Connecting a FileSystem.
- Find out the private IP address of the Oracle Cloud Infrastructure Compute Classic instance on which Oracle Cloud Infrastructure Storage Software Appliance is provisioned. See Finding Out the IP Addresses of the Appliance Instance.
- Identify the Oracle Cloud Infrastructure Compute Classic instances (clients) for which you want to provide shared storage, and note the public IP addresses of those instances.

  You can find out the public IP address of an instance by viewing its details in the Oracle Cloud Infrastructure Compute Classic web console. See Monitoring Instances in *Using Oracle Cloud Infrastructure Compute Classic*.

Steps

Complete the following steps on each Oracle Cloud Infrastructure Compute Classic instance for which you want to provide shared storage.

1. Log in to the instance using `ssh`.

   See Accessing an Instance Using SSH in *Using Oracle Cloud Infrastructure Compute Classic*.
2. Create a directory that you'll use as the mount point, as shown in the following example:

   ```
   mkdir /mnt/myFileSystem1
   ```
3. Open the `/etc/fstab` file in a text editor like `vi`.

   > **Note:**
   >
   > You must be the `root` user or a user with `sudo` privileges to be able to edit `/etc/fstab`.

   ```
   sudo vi /etc/fstab
   ```

   Each line in this file defines a local directory on which a particular file system—either local or remote—is mounted.
4. Add a line to mount your appliance filesystem, in the following format:

   `remote-filesystem mountpoint type options`

   Here's an example:

   ```
   10.196.35.246:/myFileSystem1 /mnt/myFileSystem1 nfs vers=4
   ```

   In this example,

   - `10.196.35.246` is the private IP address of the appliance instance.
   - `myFileSystem1` is the filesystem that you want to mount.
   - `/mnt/myFileSystem1` is the directory (on the client instance) on which you want to mount the filesystem.
   - `nfs` indicates that this mount should use the NFS protocol.
   - `vers=4` indicates that the NFS v4 protocol must be used.

Besides `vers=4`, you can specify additional options separated by commas. See http://linux.die.net/man/5/nfs.

5. Save and close `/etc/fstab`.
6. Activate the new mount by running the `mount` command, as shown in the following example:

   ```
   sudo mount /mnt/myFileSystem1
   ```

   Any files that the applications running on the client Oracle Cloud Infrastructure Compute Classic instance write to `/mnt/myFileSystem1` will be cached in the appliance filesystem and uploaded asynchronously to your storage service instance in the cloud.

   At any time, you can unmount and mount the filesystem again by using the `umount` and `mount` commands, respectively, as shown in the following examples:

   ```
   umount /mnt/myFileSystem1
   ```

   ```
   mount /mnt/myFileSystem1
   ```

> **Note:**
>
> - If you're facing difficulty in accessing the filesystem from your NFS client, then run the following command to flush the entries from the kernel's export table:
>
>   ```
>   exportfs -f
>   ```
>
>   Fresh entries will be added to the kernel's export table, when you next mount the filesystems on the NFS client.
> - If your client Oracle Cloud Infrastructure Compute Classic instance isn't set up to boot from a persistent disk, then every time the instance is re-created, you must mount the filesystems again.
>
>   If the client instance boots from a persistent disk, then when the instance is re-created, all the mounts defined in `/etc/fstab` are activated automatically.

Preserving Files in the FileSystem Cache

> **Note:**
>
> Archive support not available on Oracle Cloud at Customer

When you write a file to your filesystem, it's initially stored in the filesystem cache, and then uploaded to the container on your storage service instance in the cloud. Once a file has been uploaded to a container, it may get removed from the filesystem cache by the cache manager. The cache is reclaimed using the *Least Recently Used* (LRU) cache management policy to meet the cache threshold that's specified in the filesystem advanced settings. If you want specific files to be always available in the

cache for quick access, you can preserve them in the filesystem cache by pinning them to the cache. Once pinned, the files are not removed from the filesystem cache, except if you specifically unpin them.

You can pin files to the cache for filesystems connected to any storage class, `Standard` or `Archive`. Files that you write to a filesystem are uploaded to your storage service instance in the cloud, regardless of whether the files are pinned to the cache.

If the file that you want to pin to the filesystem cache is not present in the cache, then it's automatically downloaded to the cache from the container on your storage service. If that file belongs to a filesystem of the `Archive` storage class, then it's first restored, and then downloaded.

Note:

- When selecting the files for cache pinning, consider the overall cache threshold and calculate the residual cache space that would be available for normal cache operations. For example, if your cache threshold is 1 TB, and you estimate files that are pinned to the cache to occupy 300 GB, then you'd have 700 GB usable space on your cache after pinning the files. See Best Practices – Configuring Cache Storage and Best Practices – Determining the Cache Size.
- By default, the cache pinning feature is enabled on all filesystems.
- When you restore a file that belongs to a filesystem of the `Archive` storage class, the file will remain in the corresponding container of the `Standard` storage class for the duration specified in the **Restore Object Retention** field for any filesystem. Its continued availability in the cache will depend on the LRU operation. However, when you pin such a file to the cache, the restored file will remain in the cache, until if you specifically unpin it.

Enabling and Managing Cache Pinning

To perform cache pinning operations for a filesystem, run the following command from the NFS client on which the filesystem is mounted:

```
cat /path/to/mountpoint:::cache:cache_command[:argument]
```

The following table lists the cache pinning operations and the corresponding command and argument for each operation:

Operation	Cache Command	Argument
Enable cache pinning for a filesystem Note that, by default, cache pinning is enabled for all filesystems.	`set-preserve-option`	`true`
Get the cache pinning status for a filesystem	`get-preserve-option`	No argument
Disable cache pinning for a filesystem	`set-preserve-option`	`false`

Operation	Cache Command	Argument
List the files that are pinned to the cache	`list-preserve`	No argument
Remove any files from the preserve list that have been deleted	`list-preserve-update`	No argument
Add a file to the preserve list	`add-preserve`	No argument
Remove a file from the preserve list	`remove-preserve`	No argument
Clear the preserve list	`clear-preserve`	No argument

Example Commands

- To enable cache pinning for the `myFS` filesystem:

  ```
  cat /mnt/gateway/myFS/:::cache:set-preserve-option:true
  ```

- To get the cache pinning status for `myFS`:

  ```
  cat /mnt/gateway/myFS/:::cache:get-preserve-option
  ```

  The output of this command is `true` if cache pinning is enabled for the filesystem. Otherwise, `false`.

- To disable cache pinning for the `myFS` filesystem:

  ```
  cat /mnt/gateway/myFS/:::cache:set-preserve-option:false
  ```

- To add a file `myFile` of the `myFS` filesystem to the preserve list:

  ```
  cat /mnt/gateway/myFS/myFile:::cache:add-preserve
  ```

- To find out which files are added to the preserve list of the `myFS` filesystem:

  ```
  cat /mnt/gateway/myFS/:::cache:list-preserve
  ```

  A sample output of the above command:

  ```
  ["/doNotDelete.txt", "/myFileMetadata", "/myFile"]
  ```

- To remove the file `myFile` from the preserve list

  ```
  cat /mnt/gateway/myFS/myFile:::cache:remove-preserve
  ```

- To update the preserve list when the output of the `cache:list-preserve` command indicates that a pinned file has been removed from the filesystem:

  ```
  cat /mnt/gateway/myFS/:::cache:list-preserve-update
  ```

  A sample of the original preserve list:

  ```
  ["/doNotDelete.txt", "/myFileMetadata"]
  ```

  Output of the `cache:list-preserve` command after the file `myFileMetadata` is removed from the cache:

  ```
  ["/doNotDelete.txt", "Status: 1 files appear to no longer exist. Please run list-preserve-update"]
  ```

  Output of the `cache:list-preserve-update` command:

  ```
  ["/doNotDelete.txt"]
  ```

- To clear the preserve list for a filesystem:

  ```
  cat /mnt/gateway/myFS/:::cache:clear-preserve
  ```

Enabling File Versions Compaction

Oracle Cloud Infrastructure Storage Software Appliance allows file operations with byte-level granularity, such as append, re-write, over-write, and truncate. When a file is modified in an appliance filesystem, it results in a new version of the file being created and uploaded to the account.

When you create a filesystem, you can choose whether older versions of an object stored in the cloud must be retained whenever the corresponding file is updated or deleted in the filesystem.

- If you **select** the **Delete Old File Versions** check box for a filesystem, then a version compaction process runs in the background when you start the appliance. Upon completion of one cycle, the process sleeps for 24 hours before starting the next version compaction cycle. This process removes older versions of all the objects in the container that's connected to the filesystem. Only the latest version of each object is retained.
- If you don't select the **Delete Old File Versions** check box:
  - When a file is updated in the filesystem, a new version of the corresponding object in the cloud is created. Additional capacity is consumed in the cloud after each such update operation.
  - When a file is deleted in the filesystem, all older versions of the object in the cloud are retained. Capacity continues to be used in the cloud for the file that you deleted in the appliance.

Viewing the Details of a FileSystem

You can view the configuration details of a filesystem and also monitor the upload activity, through the management console of Oracle Cloud Infrastructure Storage Software Appliance.

To view the details of a filesystem, log in to the management console, and click the name of the filesystem:

- The **Details** tab displays the storage service type.
  For Oracle Cloud Infrastructure Object Storage Classic accounts, the identity domain associated with your account is also displayed.

  For Oracle Cloud Infrastructure Object Storage Classic accounts, you can also view a graphical representation of the amount of cloud storage being used by the filesystem and available free space on the appliance host, with the following details:
  - Current Usage
  - Free Space
- The **Settings** tab displays the following details:
  - Details of the account specified for the filesystem

  - Enabled filesystem properties (such as encryption, `Archive` storage class and deleting old file versions) (Not available on Oracle Cloud at Customer)
  - NFS and cache settings for the filesystem

  You can edit these settings. If you make any changes, remember to click **Save**.

- The **Activity** tab shows the ongoing and pending upload activity.
  If you contact Oracle Support Services about any issue with the filesystem, you may need to provide the filesystem log to help the Oracle Support Services technician diagnose the issue. To view or download the filesystem log, click **View Streaming Logs** near the lower-right corner of the **Details** tab.
- The **Completed Uploads** tab shows the last 100 files that were uploaded to your account during the current browser session. Note that this list doesn't persist across browser sessions. If you refresh the page or if you open the **Completed Uploads** tab in another browser after the files are uploaded, then the list will be empty.
- You can also disconnect the filesystem. See Connecting a FileSystem.

Next Task

You can edit the properties of a filesystem. See Changing the Properties of a FileSystem.

Changing the Properties of a FileSystem

You can change the properties of a filesystem, through the management console of Oracle Cloud Infrastructure Storage Software Appliance.

> **Note:**
>
> You can't enable or disable encryption, and you can't change the storage class of the filesystem.

To change the properties of a filesystem, log in to the management console, and click the name of the filesystem in the **Dashboard** pane:

- You can edit the filesystem properties and advanced settings (such as the cache limits) of the service instance specified for the filesystem in the **Settings** tab.

After updating the filesystem properties, click **Save**.

> **Note:**
>
> For the changes to take effect, you must disconnect and reconnect the filesystem. See Connecting a FileSystem.

Deleting a FileSystem

When you no longer need a filesystem, you can delete it from Oracle Cloud Infrastructure Storage Software Appliance.

To delete a filesystem:

1. Log in to the management console.
2. On the **Dashboard** tab, identify the filesystem that you want to delete.
3. Make sure that the filesystem is disconnected. If it's still connected, then click **Disconnect**.

> **Note:**
>
> The container to which the filesystem was previously connected and the stored data remain intact even after the filesystem is disconnected.

4. After the filesystem is disconnected, click its name.
5. On the page that displays the details of the filesystem, click **Delete**.

The filesystem is deleted from Oracle Cloud Infrastructure Storage Software Appliance. Deleting a filesystem does not automatically delete the objects in the container. If you'd like to remove the objects from the container, all the files should be deleted from the filesystem prior to disconnecting the filesystem, with version compaction enabled.

7

Managing the Appliance

Topics

Changing the Admin Password for the Appliance

The `admin` password is required to log in to the management console of the appliance. You set the password while provisioning the appliance or while creating the first filesystem.

To change the `admin` password:

1. Log in, using `ssh`, to the appliance instance, as the `opc` user:

   ```
   ssh opc@appliance_ip -i private_key
   ```

   - `appliance_ip` is the public IP address of the appliance instance. See Finding Out the IP Addresses of the Appliance Instance.
   - `private_key` is the full path and name of the file that contains the private key corresponding to any of the public keys that you specified in the appliance configuration file.
2. Assume the `root` role: `sudo su`
3. Go the `/opt/oracle/gateway/gateway/admin/bin` directory:

   ```
   cd /opt/oracle/gateway/gateway/admin/bin
   ```
4. Run the following command:

   ```
   gateway password:reset
   ```
5. Set the new password:

   ```
   gateway password:set new_password
   ```

   The password can contain from 8 to 32 characters, with at least one special character, one numerical character, one uppercase character, and one lowercase character.

Finding Out the IP Addresses of the Appliance Instance

The Oracle Cloud Infrastructure Compute Classic instance that hosts the appliance has two IP addresses: a fixed public IP address, and a private IP address that changes every time the instance is re-created or upgraded.

- Use the public IP address to connect to the appliance instance using `ssh`, and to log in to the management console of the appliance.
- Use the private IP address to mount the appliance filesystems on the client Oracle Cloud Infrastructure Compute Classic instances.

To find out the IP addresses of the appliance instance:

1. On the host on which you downloaded the appliance provisioning tool, go to the directory that contains the tool, and run the following command:

   `./fscs.sh -i config_file`

   In this command, `config_file` is the full path and name of the appliance configuration file.

2. The tool prompts you to enter the password for the Oracle Cloud Infrastructure Compute Classic user specified in the configuration file. Enter the password.

   The tool validates the password and displays the IP addresses of the instance, as shown in the following example:

   ```
   Extracting IP addresses ...
   Private IP address: 10.196.35.245
   Public IP address: 203.0.113.48
   ```

Starting, Stopping, and Restarting the Appliance

When you provision the appliance using the provisioning tool, the appliance service and the management console are started automatically. You may need to stop and restart the appliance service or the management console in some situations.

The following are a few examples of the scenarios when such a restart may be necessary:

- Prepare for a graceful shutdown of the appliance host
- Maintain the appliance host
- Suspend appliance operation
- Recover from an issue, such as a hung management console

> **Important:**
>
> If you want to stop or restart the appliance host, then wait for any pending or ongoing write operations from the client instances to complete.

To stop, start, or restart the appliance, complete the following steps:

1. Log in to your client instances using ssh and unmount the appliance filesystems before you stop or restart the appliance. (Example: `umount /mnt/myFileSystem1`)

2. Log in, using `ssh`, to the appliance instance, as the `opc` user:

   `ssh opc@appliance_ip -i private_key`

   - `appliance_ip` is the public IP address of the appliance instance. See Finding Out the IP Addresses of the Appliance Instance.
   - `private_key` is the full path and name of the file that contains the private key corresponding to any of the public keys that you specified in the appliance configuration file.

3. Assume the `root` role: `sudo su`
4. Run the following command:
   - For Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution Release 16.3.1.2 and 16.3.1.3:

     Syntax:

     ```
     supervisorctl [stop | start | restart]  all
     ```

     Examples:

     – To stop the appliance:

     ```
     supervisorctl stop all
     ```

     – To start the appliance:

     ```
     supervisorctl start all
     ```

     – To restart the appliance:

     ```
     supervisorctl restart all
     ```

   - For Oracle Cloud Infrastructure Storage Software Appliance– Cloud Distribution Release 16.3.1.0:

     Syntax:

     ```
     supervisorctl -c /etc/supervisord/supervisord.conf [stop | start | restart]
      all
     ```

     Examples:

     – To stop the appliance:

     ```
     supervisorctl -c /etc/supervisord/supervisord.conf stop all
     ```

     – To start the appliance:

     ```
     supervisorctl -c /etc/supervisord/supervisord.conf start all
     ```

     – To restart the appliance:

     ```
     supervisorctl -c /etc/supervisord/supervisord.conf restart all
     ```

5. Log in to your client instances and mount the appliance filesystems back to your mountpoints. See Mounting Appliance FileSystems on Client Instances.

Re-creating the Appliance

You may need to re-create the appliance instance when the instance is in an error state and you're unable to recover from the error.

> **Note:**
>
> The appliance instance gets a new private IP address. So after the appliance is re-created, you must mount all the filesystems again on the client Oracle Cloud Infrastructure Compute Classic instances by using the new private IP address.
>
> All the data disks are preserved. The appliance configuration settings, filesystem configurations, and data in the cache remain intact.

1. Make sure that you have the appliance configuration file that you originally used to create the appliance.
2. On the host on which you downloaded the appliance provisioning tool, go to the directory that contains the tool, and run the following command:

   ```
   ./fscs.sh -r config_file
   ```

   In this command, `config_file` is the full path and name of the appliance configuration file.
3. The tool prompts you to enter the password for the Oracle Cloud Infrastructure Compute Classic user specified in the configuration file. Enter the password.

   The tool validates your credentials, and then does the following:

   a. Deletes the appliance instance.

   b. Creates the appliance instance.

   After creating the appliance instance, the tool displays the private and public IP addresses of the instances, as shown in the following example:

   ```
   Deleting Oracle Cloud Infrastructure Storage Software Appliance
   instance ... successful
   Creating Oracle Cloud Infrastructure Storage Software Appliance
   instance ... successful
   Extracting IP addresses ...
   Private IP address: 10.196.35.245
   Public IP address: 203.0.113.48
   ```
4. Note the new private IP address.
5. Mount the appliance filesystems on the client Oracle Cloud Infrastructure Compute Classic instances again. See Mounting Appliance FileSystems on Client Instances.

Adding Data Disks to the Appliance Instance

At any time, you can attach additional disks to the appliance instance to provide more local storage capacity that the appliance can use for caching data.

All the data disks that are attached to the appliance instance are combined into a resizeable logical volume group, named `oracle_data_vg`, which is mounted on the appliance instance at `/opt/oracle/datadisk`. When you add data disks, the `oracle_data_vg` volume group on the appliance instance is expanded to include the additional disks.

> **Caution:**
>
> The process of adding disks causes downtime for the appliance. The instance is deleted, the additional disks specified in the configuration file are created, a new appliance instance is created, and all the disks (including the new ones) are attached to the instance.
>
> The appliance instance gets a new private IP address. So after the appliance is re-created, you must mount all the filesystems again on the client Oracle Cloud Infrastructure Compute Classic instances by using the new private IP address.
>
> All the data disks are preserved. The appliance configuration settings, filesystem configurations, and data in the cache remain intact.

1. Determine the additional storage capacity that you need for the appliance cache. See Best Practices – Configuring Cache Storage and Best Practices – Determining the Cache Size.
2. In a text editor, open the appliance configuration file that you used originally to create the appliance.
3. Under the `storage` attribute, add a `datadisk` subattribute for each disk that you want to add, as shown in the following example:

   Example:

```
storage:
  - datadisk:
      name: /Compute-acme/jack.jones@example.com/myVol1
      size: 10G
      property: /oracle/public/storage/default
  - datadisk:
      name: /Compute-acme/jack.jones@example.com/myVol2
      size: 10G
      property: /oracle/public/storage/default
  - datadisk:
      name: /Compute-acme/jack.jones@example.com/myVol3
      size: 10G
      property: /oracle/public/storage/default
```

You can attach up to nine data disks, each up to 2 TB in size.

For more information about the parameters of the `datadisk` attribute, see Appliance Configuration Parameters.

4. Validate the YAML format of the configuration file by using a tool such as YAML Lint (http://www.yamllint.com).

> **Note:**
>
> Oracle does not support or sponsor any third-party YAML validation tool.

5. On the host on which you downloaded the appliance provisioning tool, go to the directory that contains the tool, and run the following command:

   `./fscs.sh -a config_file`

   In this command, `config_file` is the full path and name of the appliance configuration file.

6. The tool prompts you to enter the password for the Oracle Cloud Infrastructure Compute Classic user specified in the configuration file. Enter the password.

   The tool validates your credentials, and then does the following:

   a. Deletes the appliance instance.
   b. Creates the additional data disks that you've specified in the configuration file.
   c. Creates the appliance instance.
   d. Attaches all the data disks (including the new disks) to the appliance instance.
   e. Expands the `oracle_data_vg` volume group on the appliance instance to include the additional disks.

   After creating the appliance instance, the tool displays the private and public IP addresses of the instances, as shown in the following example:

```
Adding new disks ...
Creating Oracle Cloud Infrastructure Storage Software Appliance
instance ... successful
Extracting IP addresses ...
Private IP address: 10.196.35.235
Public IP address: 203.0.113.48
```

7. Note the new private IP address.

> **Tip:**
>
> Retain a copy of the configuration file. You'll need it to run any further operations on the appliance.

8. Mount the appliance filesystems on the client Oracle Cloud Infrastructure Compute Classic instances again. See Mounting Appliance FileSystems on Client Instances.

Upgrading the Appliance

You can upgrade your appliance to use the specific image by downloading it from the Oracle Cloud Downloads page at http://www.oracle.com/technetwork/topics/cloud/downloads/index.html. For the workflow to upgrade your appliance to the target release, see Upgrade Workflow .

When you upgrade the appliance instance, the binaries of the appliance are updated.

Caution:

- The upgrade process causes downtime for the appliance. The appliance instance and its boot disk are deleted, a new boot disk is created (by using the image to which you want to upgrade the appliance), and the appliance instance is re-created using the new boot disk. Note that any OS-level changes you may have made on the boot disk of the appliance instance will be lost after the appliance is upgraded.

  The appliance instance gets a new private IP address. So after the appliance is re-created, you must mount all the filesystems again on the client Oracle Cloud Infrastructure Compute Classic instances by using the new private IP address.

  All the data disks are preserved. The appliance configuration settings, filesystem configurations, and data in the cache remain intact.

- Ensure that you plan the downtime appropriately, as the upgrade may take some time. The downtime varies, depending on the system resources and if there are any filesystems to be restored from the cloud. If there are no filesystems to be restored, approximately 1 million records per minute may be transferred during migration, depending on the available system resources. If the filesystem needs to be restored from the cloud, then additional time would be required to download the metadata information and prepare the metadata database for migration.

  To minimize downtime, configure and connect the filesystems on the current appliance version before you upgrade the appliance.

Note:

Releases 16.3.1.2 and 16.3.1.3 include file system changes which require migration of filesystem internal data (metadata).

- The migration takes place when your filesystems are reconnected for the first time to your storage service instance in the cloud during or after the upgradation.
- Do not delete any filesystem or change the properties of a filesystem during the migration.
- During the migration, all the filesystems are in the `read-only` state. After the filesystems are reconnected, the filesystems that had `read / write` permissions before the migration will return to the `read / write` state.
- The duration of the migration process is dependent on the size of the filesystem and can range from a few minutes to an hour or more.

Upgrade Workflow

Current Release of Your Appliance	The Release that You Want to Upgrade Your Appliance To	Steps Involved
16.3.1.0.13	16.3.1.3	Download the target image from the provided link and upgrade your appliance. See Steps to Upgrade the Appliance.
16.3.1.2	16.3.1.3	
16.3.1.0	16.3.1.3	When you're upgrading your appliance from release 16.3.1.0 to release 16.3.1.3, you must first upgrade it to release 16.3.1.0.13, and then upgrade it to your target release. 1. Download the release 16.3.1.0.13 image from the provided link and upgrade your appliance. See Steps to Upgrade the Appliance. Ignore step 8. 2. Download the target image from the provided link and upgrade your appliance. See Steps to Upgrade the Appliance.

Steps to Upgrade the Appliance

1. Edit the appliance configuration file that you used to create the appliance, and update the `image` attribute to point to the new image that you obtained earlier.

Example:

```
image: /Compute-acme/jack.smith@example.com/OSCSA-CD-image-v16.3.x.x
```

To get a list of the available machine images, send the `GET /imagelist/Compute-identity_domain/user` HTTP request to your Oracle Cloud Infrastructure Compute Classic site. See *REST API for Oracle Cloud Infrastructure Compute Classic*.

See Appliance Configuration Parameters.

2. Validate the YAML format of the configuration file by using a tool such as YAML Lint (http://www.yamllint.com).

> **Note:**
>
> Oracle does not support or sponsor any third-party YAML validation tool.

3. Log in to your appliance host, assume the `root` role, and stop the appliance.
   - If you want to upgrade your appliance from release 16.3.1.2:

     ```
     sudo su
     supervisorctl stop all
     ```

   - If you want to upgrade your appliance from release 16.3.1.0.x:

     ```
     sudo su
     supervisorctl -c /etc/supervisord/supervisord.conf stop all
     ```

4. On the host on which you downloaded the appliance provisioning tool, go to the directory that contains the tool, and run the following command:

   `./fscs.sh -u config_file`
   In this command, `config_file` is the full path and name of the appliance configuration file.

5. The tool prompts you to enter the password for the Oracle Cloud Infrastructure Compute Classic user specified in the configuration file. Enter the password.

   The tool validates your credentials, and then does the following:

   a. Validates the name of the image that you've specified in the configuration file.
   b. Deletes the appliance instance.
   c. Deletes the boot disk that was created by using the old image.
   d. Creates a new boot disk by using the image that you've specified in the configuration file.
   e. Creates the appliance instance.

   This process may take a few minutes.

After creating the appliance instance, the tool displays the private and public IP addresses of the instances, as shown in the following example:

```
Creating Oracle Cloud Infrastructure Storage Software Appliance
instance ... successful
Extracting IP addresses ...
Private IP address: 10.196.35.241
Public IP address: 203.0.113.48
```

> **Tip:**
>
> Retain a copy of the configuration file. You'll need it to run any further operations on the appliance.

6. Note the new private IP address.
7. Log in to the management console of your appliance using the public IP address.
   - If you're upgrading your appliance from release 16.3.1.0.13 to release 16.3.1.3:

     After the metadata migration is complete, the following message appears on the dashboard for the respective filesystem: `Migration completed. Please reconnect the filesystem.`

     If your filesystem is connected to your storage service instance in the cloud, but you don't see the message about the completion of metadata migration, then wait for the migration to complete.

     If your filesystem isn't connected to your storage service instance in the cloud, click **Connect**. Wait for the migration to complete.

     Disconnect and reconnect the filesystems after the migration is complete. See Connecting a FileSystem.
   - Otherwise:

     Connect your filesystems to the cloud service.
8. Mount the appliance filesystems on the client Oracle Cloud Infrastructure Compute Classic instances again. See Mounting Appliance FileSystems on Client Instances.

Deleting the Appliance

When you no longer need the appliance or if you want to set up the appliance in a different Oracle Cloud Infrastructure Compute Classic site, you can delete the appliance and the Oracle Cloud Infrastructure Compute Classic instance hosting it.

You can also choose to delete all the Oracle Cloud Infrastructure Compute Classic resources that are associated with the appliance instance.

> **Note:**
>
> Deleting the appliance has no effect on the data that was previously uploaded to your storage service instance in the cloud through the appliance. The containers and objects in your account remain intact.

Prerequisites

Wait for any pending or ongoing write operations from the client instances to complete.

Stop the appliance service and the management console from the appliance host:

```
sudo su
supervisorctl stop all
```

Steps for Deleting the Appliance

1. Make sure that you have the appliance configuration file that you used to create the appliance.
2. On the host on which you downloaded the appliance provisioning tool, go to the directory that contains the tool, and run the following command:
   - To delete *only the appliance instance* and retain its boot disk, data disks, and public IP address reservation:
     `./fscs.sh` **`-d`** *`config_file`*

     You can use the same boot disk, data disks, and public IP address later when you want to create the appliance.
   - To delete the appliance instance as well as its boot disk, data disks, and public IP address reservation:
     `./fscs.sh` **`-dd`** *`config_file`*

   In this command, `config_file` is the full path and name of the appliance configuration file.
3. The tool prompts you to enter the password for the Oracle Cloud Infrastructure Compute Classic user specified in the configuration file. Enter the password.

   The tool validates the credentials and then deletes the appliance instance.

   After deleting the Oracle Cloud Infrastructure Compute Classic instance, the tool displays the following message:

   ```
   Deleting Oracle Cloud Infrastructure Storage Software Appliance
   instance ... successful
   ```

   If you used the `-dd` option, then the tool displays the following additional messages:

   ```
   Deleting boot disk ...  successful
   Deleting data disks and IP reservation ...  successful
   ```

Monitoring the Appliance

Topics

Monitoring Upload Activity

The **Activity** tab shows the ongoing and pending upload activity in a filesystem.

When you upload a file to a filesystem, you can view the status of the upload activity.

1. Log in to the management console.
2. Select the filesystem.
3. Click the **Activity** tab.
   You can see the upload progress of the file in the **Uploading** pane.

Downloading Support Bundle

If you contact Oracle Support Services about any issue with the appliance, then you may need to provide the support bundle to help the Oracle Support Services technician diagnose the issue.

1. Log in to the management console.
2. Select the **System** tab on the upper-right side of the management console.
3. Select the **Help** tab.
4. Click **Download Support Bundle** in the **System Logs** pane.
   You can download and save the support bundle.

Support Bundle

The support bundle contains the following information:

- All necessary logs for diagnostics
- Local storage usage information
- Basic system information such as memory size, appliance version etc.
- List of filesystems

Monitoring System Status Health Check

Health Check

You can monitor the overall system status through the **System Status** pane in the right side of the management console.

The appliance health check service is an automated process run on the system to monitor the status of the following:

- Services – Databases used by the appliance, management console and other auxiliary processes
- Disk space – Local storage
  For example: If the available local storage is lesser than 10 GB, the health check service reports this as an alert.

Depending on the appliance health check analysis, the following status is displayed in the **System Status** pane:

- Healthy
- Unhealthy

The appliance health check service also displays the following details in the **System Status** pane and highlights potential issues:

- Throughput
- Available local cache
- Pending uploads

Local Cache Modes

The **Local I/O** mode might display any one of the following values based on the local disk usage:

- **Normal**
  The free space is higher than 10 GB in the appliance. You can upload files in the appliance and store them in your account.
- **Rejecting I/O**
  The free space is lower than 10 GB in the appliance. The appliance is running on protection mode and will not allow any writes to its local disk. All read operations will work as normal. All metadata operations will fail in the appliance except for deletion and truncation.

  To return to **Normal** mode, you must wait until all the ongoing upload activities are complete and the files are removed from the local cache.

> **Note:**
>
> For optimum local storage configuration, see Best Practices – Configuring Cache Storage.

You can also view the system status details and track any issues using the support bundle. See Viewing System Notifications.

Monitoring Appliance Storage Usage

The **System Stats** tab enables you to track the storage usage and availability.

1. Log in to the management console.
2. Select the **System** tab on the upper-right side of the management console.
3. Click the **System Stats** tab in the **System** pane.
   The system data is displayed in the following three panes:
   - Local Storage
   - Local I/O
   - Local Resources

Local Storage

In this pane, you can view a graphical representation of the amount of storage being used and available free storage on the appliance host, with the following details:

- Available local storage
- Storage used for pending uploads and preserved cache files
- Storage used for metadata
- Storage used for logging
- Storage used for other applications

Local I/O

This pane displays the local I/O mode of the appliance based on the local disk space usage in the appliance host.

Local Resources

In this pane, you can view the overall memory usage and memory availability for the appliance from the following fields:

- Available Cores -The number of CPUs being used by the appliance
- Maximum Memory Available to the Appliance - The total RAM available for the appliance
- Memory Used by the Appliance - The amount of memory being used by the filesystems in the appliance
- Free Memory - The amount of free RAM available in the appliance host

Viewing System Notifications

The **System Notifications** tab allows you to view the system notifications and track the overall system performance.

1. Log in to the management console.
2. Select the **System** tab on the upper-right side of the management console.
3. Click the **System Notifications** tab in the **System** pane.
   You can view the list of warnings or critical system notifications.

8

Using the Appliance to Store and Retrieve Data

Topics:

Deleting Files

Remove the files that you no longer need from the NFS client by deleting them from the directory on which the filesystem is mounted.

> **Caution:**
>
> If you've enabled the check box to delete old file versions in the filesystem, then depending on the configuration settings, the older versions of the objects may be automatically removed from your account.

To delete the older versions of the objects and the corresponding metadata files, log in to the management console, click the filesystem name, and select the check box **Delete Old File Versions** in the **Advanced** section of the **Settings** tab.

The files in the filesystem are not deleted.

For more information, see Enabling File Versions Compaction.

> **Note:**
>
> Deletion of archived objects may result in an early-deletion fee. For more information, go to https://cloud.oracle.com/storage and see the **Pricing** tab.

> **Caution:**
>
> Don't use the REST API, Java library, or any other client to retrieve, create, update, or delete objects in a container that's mapped to a filesystem in Oracle Cloud Infrastructure Storage Software Appliance. Doing so will cause the data in the appliance to become inconsistent with data in your storage service. You can't recover from this inconsistency.
>
> To prevent unauthorized users from retrieving, creating, updating, or deleting objects in a container that's connected to a filesystem in Oracle Cloud Infrastructure Storage Software Appliance, define custom roles, assign them to the appropriate container, assign the roles to only the users that should have access to the container, and specify only one of these users when defining the filesystem to be connected to the container.

Uploading Files

Topics

Uploading Files to Standard Containers

> **Caution:**
>
> Don't use the REST API, Java library, or any other client to retrieve, create, update, or delete objects in a container that's mapped to a filesystem in Oracle Cloud Infrastructure Storage Software Appliance. Doing so will cause the data in the appliance to become inconsistent with data in your storage service. You can't recover from this inconsistency.
>
> To prevent unauthorized users from retrieving, creating, updating, or deleting objects in a container that's connected to a filesystem in Oracle Cloud Infrastructure Storage Software Appliance, define custom roles, assign them to the appropriate container, assign the roles to only the users that should have access to the container, and specify only one of these users when defining the filesystem to be connected to the container.

Prerequisite

Ensure that the filesystem in Oracle Cloud Infrastructure Storage Software Appliance is connected to the appliance host. See Connecting a FileSystem.

Procedure

Copy the files to the mounted directory on the NFS client host. Oracle Cloud Infrastructure Storage Software Appliance writes the files to the disk cache. The files

are queued and then uploaded asynchronously to Oracle Cloud Infrastructure Object Storage Classic.

> **Note:**
>
> For the character restrictions applicable when you must select a file name, see Character Restrictions.

You can check the status of the files being uploaded in the management console. See Viewing the Details of a FileSystem.

Uploading Files to Archive Containers

 This topic does not apply to Oracle Cloud at Customer.

> **Caution:**
>
> Don't use the REST API, Java library, or any other client to retrieve, create, update, or delete objects in a container that's mapped to a filesystem in Oracle Cloud Infrastructure Storage Software Appliance. Doing so will cause the data in the appliance to become inconsistent with data in your storage service. You can't recover from this inconsistency.
>
> To prevent unauthorized users from retrieving, creating, updating, or deleting objects in a container that's connected to a filesystem in Oracle Cloud Infrastructure Storage Software Appliance, define custom roles, assign them to the appropriate container, assign the roles to only the users that should have access to the container, and specify only one of these users when defining the filesystem to be connected to the container.

Prerequisite

Before you connect the filesystem to the container, ensure that the container is of the `Archive` storage class. To find out the storage class of the container, see Getting Container Metadata.

Ensure that the filesystem in Oracle Cloud Infrastructure Storage Software Appliance is connected to the container in Oracle Cloud Infrastructure Object Storage Classic. See Connecting a FileSystem.

Procedure

Mount the appliance filesystem on the NFS client. Copy the files to the mount point. The appliance caches the files while they are queued and asynchronously uploads them to the corresponding `Archive` container in Oracle Cloud Infrastructure Object Storage Classic.

> **Note:**
>
> For the character restrictions applicable when you must select a file name, see Character Restrictions.

Retrieving Files

Topics

Reading Files

When a file is written to an appliance filesystem, it is stored in the local disk cache, and you can read the file directly from the mounted directory. The file is asynchronously copied to the corresponding container in your account. To retrieve the data from the container in your account by using the appliance, read the required files from the mounted directory. The appliance will automatically place the files in the local cache, if space is available.

> **Note:**
>
> (Not available on Oracle Cloud at Customer)
>
> When a file is copied to an `archive` filesystem, it is stored in the local disk cache. After the file is asynchronously copied to the corresponding `Archive` container in Oracle Cloud Infrastructure Object Storage Classic, it is stored as an archived object. If the file is in the local disk cache, then you can retrieve the file immediately. However, if the file is not available in the local disk cache and stored in the `Archive` container, then you must first restore the archived object. For more information, see Restoring Files From Archive Filesystems.
>
> If you try to download a file which does not exist in the local cache and is stored as an archived object, then an error message is displayed.

Reading the Checksum for a File

To read the checksum for a file in a filesystem, run the following command from the NFS client on which the filesystem is mounted:

```
cat /path/to/mountpoint/filename:::meta:csm
```

Restoring Files From Archive Filesystems


This topic does not apply to Oracle Cloud at Customer.

To download a file from an `Archive` filesystem, you must first restore the corresponding archived object in the container. The restored object is then downloaded and stored as a file in the appliance cache. Restoring archived files is an asynchronous operation.

To restore a file from an `Archive` filesystem, run the following command on the NFS client:

```
cat path_to_filename:::archive:restore
```

For example, a file `myFirstFile` is copied to a mounted directory `myArchiveDir` on the NFS client and is uploaded to an `Archive` filesystem `myFirstArchiveFS`. The file is asynchronously stored as an archived object in the `myFirstArchiveFS-archive` container in your Oracle Cloud Infrastructure Object Storage Classic account. To restore the archived object, enter the command:

```
cat /path_on_NFS_client/myArchiveDir/myFirstFile:::archive:restore
```

Sample Response:

```
{"path":"/
myFirstFile","restoreStatus":"inprogress","restoreObjectPercent":
{"13456760 1079-11-v1":2},"additionalInfo":""}
```

> **Note:**
>
> You can restore an archived object in an `archive` container. If you try to restore an object in a `standard` container, then the following error message is displayed:
> `archive is not a valid command class`

> **Note:**
>
> If the filesystem is deleted and if you restore an object in the `Archive` container at the same time, the object restoration is not affected.

Tracking Restoration of a File in an Archive Filesystem


This topic does not apply to Oracle Cloud at Customer.

To track the restoration progress of the file in the `Archive` filesystem, run the following command on the NFS client:

```
cat path_to_filename:::archive:restore-status
```

For example, a file `myFirstFile` is copied to a mounted directory `myArchiveDir` on the NFS client and is uploaded to `myFirstArchiveFS`. The file is asynchronously stored as an archived object in the `myFirstArchiveFS-archive` container in your Oracle Cloud Infrastructure Object Storage Classicaccount and you've run the command to restore the object. To track the restoration status, enter the command:

```
cat /path_on_NFS_client/myArchiveDir/myFirstFile:::archive:restore-status
```

Sample Response:

For Oracle Cloud Infrastructure Object Storage Classic accounts:

```
{"path":"/myFirstFile","restoreStatus":"restored","restoreObjectPercent":
{},"additionalInfo":""}
```

By default, a restored object will be downloaded and stored as a file in the `Archive` filesystem for one day. You can now read the file from the `Archive` filesystem before the restoration expires. For more information, see Reading Files.

Tracking Restoration of All Files in an Archive FileSystem


This topic does not apply to Oracle Cloud at Customer.

To track the restoration status of all the files in an `Archive` filesystem, run the following command:

```
cat /path_on_NFS_client_to_mounted_directory:::archive:jobs
```

The following is a sample response:

```
{"/myFirstFile":"restored","/mySecondFile":"inprogress", "/
myThirdFile":"inprogress"}
```

Example:

The following is an example to show the restoration and tracking the restoration status of the files in an `Archive` filesystem `myArchiveDir`:

1. Restoring the file `myFirstFile`:
   ```
   cat /mnt/dir1/myArchiveDir/myFirstFile:::archive:restore
   ```
   Output:
   ```
   {"path":"/
   myFirstFile","restoreStatus":"inprogress","restoreObjectPercent":
   {"1464707450825-13-v1":0},"additionalInfo":""}
   ```
2. Restoring the file `mySecondFile`:
   ```
   cat /mnt/dir1/myArchiveDir/mySecondFile:::archive:restore
   ```
   Output:
   ```
   {"path":"/
   mySecondFile","restoreStatus":"inprogress","restoreObjectPercent":
   {"1353643456634-13-v1":0},"additionalInfo":""}
   ```
3. Tracking the restoration progress of all the files:

```
cat /mnt/dir1/myArchiveDir/:::archive:jobs
```

Output:

```
{"/myFirstFile":"restored","/mySecondFile":"inprogress", "/
myThirdFile":"inprogress”}
```

9

Best Practices for Using Storage Software Appliance

Follow the best practices described here to get maximum benefit from Oracle Cloud Infrastructure Storage Software Appliance in terms of manageability, performance, reliability, and security.

Topics

Best Practices – Configuring Cache Storage

Oracle Cloud Infrastructure Storage Software Appliance uses local storage attached to the server (or virtual server) for hosting the filesystems and cache. Files written to a filesystem in the appliance are uploaded to the associated container, with a portion of the file set maintained locally in the filesystem as a warm cache.

For optimal performance, reliability, and fault tolerance, consider the following guidelines when configuring the local appliance storage:

- Allocate a dedicated volume for the appliance filesystem and metadata.
- Enable read-ahead on the volume.
- Provision a volume that can accommodate the local cache **and** ingest new files (upload buffer) without ever becoming more than 80% full.
  A general guideline is to use a volume that is at least 1.5 times the size of the data set that you want to hold in local cache. For example, if the expected size of the entire file set is 50 TB and if 10% (5 TB) of that file set will be accessed frequently, then the cache storage volume should have at least 7.5 TB of usable capacity.

> **Note:**
>
> – If the cache size reaches a near-full threshold, any data ingest will result in **out of space** error
>
> in the appliance.
>
> – You can increase the cache size by adding data disks to the appliance instance. See Adding Data Disks to the Appliance Instance.

Best Practices – Determining the Cache Size

The local cache of Oracle Cloud Infrastructure Storage Software Appliance serves two roles: ingest cache (upload/write buffer) and read cache. You can specify the maximum size for the read cache. The write buffer will use any remaining available space on the local storage volume and does not have a cache size setting.

The maximum size of the write buffer is an important criterion to determine the cache size. The write buffer size increases when data is uploaded in the appliance. And the write buffer size decreases after the data is transferred to cloud. Write buffer cannot be removed from local cache. When the write buffer uses all the available local cache space, any data ingest will result in **out of space** error in the appliance.

Use the following guidelines to determine the appropriate setting for the ingest cache:

- Identify the amount of data to be uploaded in the appliance. If a large amount of data must be uploaded, the appliance write buffer may reach its maximum. This will lead to I/O failure as the local cache has no space. If the data transfer can be regulated, for example, by pausing after a certain amount of data is transferred or allow the uploads to complete periodically, the local cache space can be increased and I/O failure can be avoided. You can also follow this approach for backup/cron jobs if the local cache space is lesser than the amount of data to be uploaded.
- Calculate the amount of data that would be uploaded on any typical day or a week in the appliance. Also, calculate the amount of data that can be uploaded over a time period, based on the available bandwidth or historical data. The difference between former and latter data quantities should not exceed the write buffer size.
- If the application can handle I/O failure and resume the data transfer, set the write buffer size with the amount of data that you'd like to upload before the cache size decreases.

Configuring the read cache size is necessary only when the appliance must retrieve a significant amount of data from the cloud. Alternatively, you can preserve frequently accessed files in the local cache. Setting aside a large portion of the local cache for read cache might impact the amount of data that can be uploaded before the cache size decreases. Configuring the read cache is optional and depends on the appliance workload.

Use the following guidelines to determine the appropriate setting for the read cache:

- The default limit of the read cache size is the lower of 300 GB or the storage volume size.
- Do not set the read cache maximum to the size of the local storage volume. Doing so would allocate 100% of the volume for read cache and would not leave available capacity for ingest. If there is no available space for new file ingest, then the appliance might stop the data ingest and begin evicting files from the read cache to create space. This severely degrades ingest performance.
- Start with a read cache setting that is 50% of the size of the local storage volume (leaving 50% for ingest). Monitor the available capacity on the local storage volume over time, especially after periods of very high or sustained ingest activity. If the available capacity remains above 30% consistently, consider increasing the read cache size. If the available capacity is consistently below 20%, then consider decreasing the read cache size.

- The general strategy is to set the read cache size to equal the amount of data that you anticipate to be accessed frequently, while leaving enough capacity on the volume for the ingest cache (write buffer).

After you size the cache, you can choose to configure the read cache either while creating the filesystem or later. See Adding a FileSystem and Changing the Properties of a FileSystem.

> **Note:**
>
> You can increase the cache size by adding data disks to the appliance instance. See Adding Data Disks to the Appliance Instance.

Best Practices – Encrypting Data Using Keys

You can provide your own RSA asymmetric keys if you've enabled encryption for a filesystem. The symmetric key converts the data to a readable form called *cleartext*. If you lose the keys, you lose the data.

Asymmetric keys: There's a single key pair for every instance of the appliance. The same key pair is used to encrypt information related to local configuration. If you provide an asymmetric key pair, then the key pair is used to encrypt or decrypt the specified filesystem database configuration items. Ensure that the asymmetric keys are backed up.

Symmetric keys: The symmetric key is stored within the local filesystem database. Each filesystem can have its own unique symmetric encryption key. The symmetric key is encrypted using the asymmetric key that's stored locally on the disk.

At any time, you can download a `tar.gz` file containing the details of all the keys stored on the disk.

Key rotation enables data recovery if the appliance fails at any time.

Rotating Keys in the Appliance

1. Log in to the management console and select the filesystem.
2. Provide your asymmetric key pair. The appliance ensures that the keys are valid by encrypting and decrypting randomly generated sample data.
3. If the keys are valid, then they are saved in a temporary location on the disk. The old keys are moved to a backup location on the disk.
4. The database's encrypted configuration items are encrypted again in the filesystem by using the asymmetric key pair.
5. The new keys are saved to a permanent location on the disk.
6. Download the compressed key archive of the encryption keys. The compressed archive includes the new key details as well as the backup keys.

Best Practices – Scalability Recommendations

- Ensure that the number of objects stored in an appliance filesystem doesn't exceed 10 million (10000000). For data sets that consist of more than 10 million objects, ensure that the objects are distributed across multiple filesystems.
- The minimum amount of memory required for any appliance filesystem is 16 GB.

  - For filesystems with the number of files up to 5 million, the required amount of memory is 32 GB.
  - For large filesystems with the number of files up to 10 million, the required amount of memory is 64 GB
- To improve the efficiency of file ingest and cloud upload operations, and to reduce the number of objects in the namespace, bin-pack or zip small files before writing them to the appliance.
- Multiple filesystems can be created on a single appliance. However, for optimal performance, ensure that each filesystem is hosted on a dedicated appliance.

Best Practices – Recommended Workloads and Use Cases

Oracle Cloud Infrastructure Storage Software Appliance is an effective cloud gateway for many workloads. Use the following guidelines to determine whether the appliance is appropriate for your specific use cases and workloads:

- The appliance supports NFSv4 in asynchronous mode and POSIX Sync mode. The POSIX Sync mode is enabled in the appliance by default.
  In the asynchronous mode, there is scope for data loss in the event of a sudden server failure. Avoid using the appliance for workloads and use cases that require synchronous write behavior.
- The appliance is ideal for backup and archive use cases that require the replication of infrequently accessed data to cloud containers. (Not available on Oracle Cloud at Customer)
- Carefully consider use cases that involve frequent changes to existing files. Each time a file is modified and closed, the appliance creates a new version of the file, which is then uploaded to the container in your service instance, replacing the previous version. The appliance will be less efficient and may not perform optimally for this type of workload.
- Don't run applications and executables directly from the appliance mount points, particularly if the appliance cache is not large enough for all the files that the applications will access. Applications typically create temporary files and modify them often, affecting the operational efficiency of the appliance.

10

Frequently Asked Questions

Go to https://cloud.oracle.com/storage-classic/storage-appliance/faq.

11
Troubleshooting

This section provides solutions for problems you may encounter while using Oracle Cloud Infrastructure Storage Software Appliance.

Topics:

The provisioning tool keeps failing with an error that the Oracle Cloud Infrastructure Compute Classic endpoint is not valid. I'm sure that I've specified the correct endpoint in the configuration file.

Validating the Oracle Cloud Infrastructure Compute Classic endpoint is one of the first tasks that the provisioning tool does. If it fails at that point, and if you're sure that the endpoint that you've specified in the configuration file is correct, then the issue may be caused by formatting errors in the configuration file.

For example, if you edited the configuration file on a Windows computer and then did a binary transfer to the Oracle Linux host on which you're running the provisioning tool, then the file may contain extra line-ending characters.

To solve this problem, use the `dos2unix file` command to convert the file to the UNIX format. Alternatively, open the file on Linux using a text editor like `vi`, and run the `set ff=unix` command. Note that even if the line endings are fine in your configuration file, you can run the `dos2unix file` and `set ff=unix` commands safely.

If the invalid endpoint error continues, download the configuration file template afresh and edit it on your Oracle Linux host. If you prefer editing the configuration file in Windows, then when you transfer the file to your Linux host, do an *ASCII* (not binary) transfer. The steps to do an ASCII transfer depend on the file transfer client that you use. In FileZilla, for example, you should set **Transfer type** (in the **Transfer** menu) to ASCII.

I get a permission error like the following while creating the appliance.

```
{"message": "User /Compute-acme/jack.jones@example.com is not permitted to
perform \"orchestration.add\" on orchestration:/Compute-acme/
jack.jones@example.com/oscsa-boot"}
```

This error occurs if the Oracle Cloud Infrastructure Compute Classic user that you've specified in the configuration file doesn't have the `Compute_Operations` role.

- Ask your Oracle Cloud Infrastructure Compute Classic administrator to assign the `Compute_Operations` role to the user in Oracle Cloud My Services. See Managing User Roles in *Managing and Monitoring Oracle Cloud*.
- Alternatively, update the configuration file to specify a user that has the `Compute_Operations` role. Then, try creating the appliance by using the updated configuration file.

I interrupted the provisioning tool before it finished running. What should I do now?

Run the tool again, with the same option that you specified earlier.

"Cannot satisfy both the placement and resource requirements" error during provisioning.

At times, you may get a long error message with the text `Cannot satisfy both the placement and resource requirements`, as highlighted in the following example, which has been truncated for readability:

```
"errors":{"0":"re-launching instances: /Compute-acme/jack.smith@examp
le.com/vm"}},"obj_type":"launchplan","ha_policy":"active","label":...
..."state":"error","error_reason":"Cannot satisfy both the placement
and resource requirements.",...,"info":{"errors":{"/Compute-acme/jack
.smith@example.com/gateway/vm":"error"}},"status_timestamp":"2016-01-
27T20:18:53Z","name":"/Compute-acme/jack.smith@example.com/gateway"}
```

This error occurs when the Oracle Cloud Infrastructure Compute Classic site doesn't have sufficient CPU or memory resources to provision the appliance instance. Try specifying a smaller `shape` attribute in the configuration file. If the provisioning continues to fail, change the Oracle Cloud Infrastructure Compute Classic endpoint in the configuration file to a different site and try again. If the error persists, inform your Oracle Cloud Infrastructure Compute Classic administrator about the quota problem.

When I tried re-creating the appliance (by using the -r option), the provisioning tool exited with the following error.

```
{"message": "Orchestration \"/Compute-identity_domain/user/
orchestration_name\" not found"}
```

This error occurs if you try re-creating the appliance after deleting it. It occurs because, when you delete the appliance (using the `-d` option), the orchestration that controls the instance is deleted from Oracle Cloud Infrastructure Compute Classic. If you want to create the appliance again after deleting it, then run provisioning tool with the `-c` option. See Creating the Appliance.

Error: Oracle Cloud Infrastructure Storage Software Appliance configuration file not found!

Make sure that the name and path of the configuration file that you specified when running the provisioning tool are correct, and run the tool again.

I forgot the appliance `admin` password. How do I reset or change it?

See Changing the Admin Password for the Appliance.

I'm unable to `ssh` to the Oracle Cloud Infrastructure Compute Classic instance that hosts my appliance.

See **Can't connect to an instance using SSH** in *Using Oracle Cloud Infrastructure Compute Classic*.

The mount points on my client Oracle Cloud Infrastructure Compute Classic instances are unresponsive or unusable.

This problem occurs if the appliance instance is re-created. The appliance instance may have been re-created either manually (using the `-r`, `-d`, `-u` options of the provisioning script) or automatically by Oracle Cloud Infrastructure Compute Classic to recover from a crash of the instance or the physical node on which the appliance is provisioned.

1. Verify that the appliance instance is running.

   Note that automatic re-creation of the appliance instance by Oracle Cloud Infrastructure Compute Classic may take a few minutes.

   If you can log in, using `ssh`, to the public IP address of the appliance instance as the `opc` user, or if you can access the login page of the management console (`https://`***`instance_public_ip_address`***), then the appliance instance is running.

2. Find out the current private IP address of the appliance instance. See Finding Out the IP Addresses of the Appliance Instance.

3. Do the following on each of the client Oracle Cloud Infrastructure Compute Classic instances on which you had previously mounted filesystems:

   a. Unmount each of the unresponsive mount points:

      `sudo umount` *`mount_point`*

   b. Mount the filesystem again.

      See Mounting Appliance FileSystems on Client Instances.

“Stale file handle” error on the NFS clients

If the `Stale file handle` error occurs and if you get the `mount.nfs: Connection timed out` error when you attempt to remount the appliance filesystem, then the appliance instance may have unexpectedly rebooted while under a heavy load.

To solve this problem:

1. Stop the read and write operations of the filesystem. Disconnect the filesystem, and then connect it again. See Connecting a FileSystem.
2. Stop the read and write operations of the filesystem. Unmount the filesystem, and then mount it again. See Mounting Appliance FileSystems on Client Instances.

If the problem persists, then run the following command to flush the entries from the kernel’s export table:

```
exportfs -f
```

Fresh entries will be added to the kernel’s export table, when you next mount the filesystems on the NFS client.

My application can’t write any more data to the appliance.

The filesystem cache may be full. Consider adding more data disks. See Adding Data Disks to the Appliance Instance.

The appliance instance isn’t responding and may have failed. How do I recover from this failure without losing data?

All the data that you’ve written to the appliance is safe. Data that’s been uploaded to your storage service instance is not affected by a failure of the appliance instance. Data that’s in the appliance cache is safe as well, because the block storage volumes that hold the cached data are intact.

Re-create the appliance using the same configuration file that you used to provision the appliance originally. See Re-creating the Appliance.

My appliance disks have crashed and are irrecoverable. What should I do?

1. If you can log in to the management console, log in, and make a note of all the filesystem names, and their caching and encryption settings.
2. Delete the appliance. See Deleting the Appliance.
3. Provision the appliance again. See Creating the Appliance.
4. Create the required filesystems, **with the same names as in the old appliance**. See Adding a FileSystem.
5. Connect the filesystems to your storage service instance in the cloud. See Connecting a FileSystem.

   As long as the filesystem names are the same as they were in the old appliance, the filesystems will be connected to the correct containers that contain the data that was uploaded previously using the old appliance.
6. Mount the filesystems on the client Oracle Cloud Infrastructure Compute Classic instances. See Mounting Appliance FileSystems on Client Instances.

Contacting Oracle for Support

1. Go to https://support.oracle.com.
2. In the **Sign In** pane, select **Cloud Portal** as the portal and click **Sign In**.
3. On the **Dashboard** page, click **Create Service Request**.
4. In the **Create Service Request** wizard, do the following:
    a. In the **Service Type** field, select your storage service instance in the cloud.
    b. In the **Problem Type** field, Select **Issues accessing a Storage Software Appliance** or **Issues installing or upgrading a Storage Software Appliance**, and select the appropriate problem subtype.
5. Follow the prompts in the wizard to complete the service request.

A
Character Restrictions

This section lists the character restrictions when creating and updating resources in your storage service instance in the cloud.

Input Restrictions for FileSystem Name

Container Operation	Input Parameter	Input Restrictions	Unsupported Characters (If any)
• Create your first filesystem • Add a filesystem	Filesystem name	• Only UTF-8 characters • Maximum of 256 bytes • Can start with any character • Cannot contain a slash (/) character because this character delimits the filesystem name	• **Characters**: ', ", `, <, and >. For example: jack's_container, "Future_Use"_files, ToUs<Clients **Strings**: When the name contains `/./` or `/../`; When the name ends with `/.` or `/..`; For example: mymachine/./etc, current/../folder, download_directory/., root/misc/..
• Create a container • Update container metadata	Container custom metadata name (`X-Container-Meta-{name}`)	• Only ASCII characters • Maximum of 128 bytes	The following US-ASCII characters are not supported for use in the container metadata names: • Control characters (octet 0-31) and DEL (127) • Separators (,), <, >, @, , , ;, :, \, ", /, [,], ?, =, {, }, space, horizontal-tab

Input Restrictions for File Name

Object Operation	Input Parameter	Input Restrictions	Unsupported Characters (If any)
• Upload files to `standard` containers • Upload files to `archive` containers	File name	• Only UTF-8 characters • Maximum of 1061 bytes • Can start with any character	• **Characters**: ', ", `, <, and >. For example: jane's_file, "Hello_World".txt, Send>Customers.pdf • **Strings**: When the name contains `/./` or `/../`; When the name ends with `/.` or `/..`; For example: mymachine/./etc, current/../file, php/., object/..
• Create or replace an object • Update object metadata	Object custom metadata name (`X-Object-Meta-{name}`)	• Only ASCII characters • Maximum of 128 bytes	The following US-ASCII characters are not supported for use in the object metadata names: • Control characters (octet 0-31) and `DEL` (127) • Separators `(`, `)`, <, >, `@`, `,`, `;`, `:`, \, ", /, `[`, `]`, ?, =, {, }, space, horizontal-tab