



ORACLE  
PRESS

# DevSecOps in Oracle Cloud

Securing and Automating  
Oracle Cloud Infrastructure

ORACLE

Erik Benner  
Ahmed Aboulnaga  
Dhrumil Patel

# DevSecOps in Oracle Cloud: Securing and Automating Oracle Cloud Infrastructure

---

Erik Benner, ACED  
Ahmed Aboulnaga, ACE Pro  
Dhrumil Patel



Pearson

*This page intentionally left blank*

# Oracle IaaS—Storage

OCI offers highly reliable storage options that deliver consistently high performance and scalability and are backed by an availability service-level agreement (SLA). Numerous cloud storage options are available. These options are low cost, on demand, and address different workload requirements depending on the purpose. Table 7-1 summarizes the main cloud storage options that OCI offers and the use case they can serve.

**Table 7-1** Cloud Storage Products

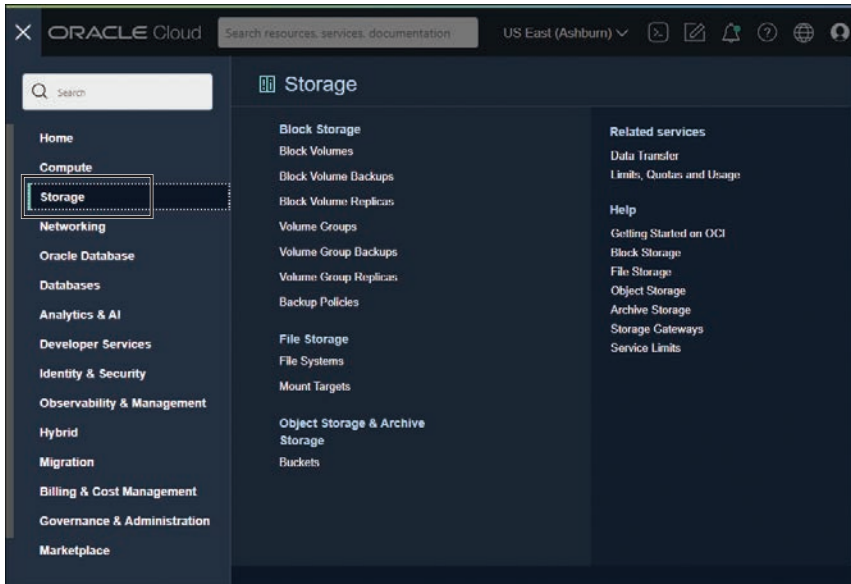
Storage Type	Purpose
Block Volumes	Provides storage to compute instances
Object Storage	Stores unstructured data of any content type
File Storage	Provides an enterprise-grade network file system (NFS)
Archive Storage	Provides storage service for infrequently accessed data
Data Transfer Appliance	Enables you to securely move terabytes or petabytes of data offline between on-premises data centers and the cloud
Storage Gateway	Enables you to store and retrieve data on OCI Object Storage from on-premises data centers

In this chapter, we will discuss the setup and configuration of block volumes, object storage, file storage, and archive storage. *Block volumes* are low latency, SSD-based storage that are attached to compute instances. *Object storage* is intended to store unstructured data and can be accessed through the browser, CLI, REST API, or OCI SDK. This storage is ideal for data lakes, for example. Users have the ability to upload any type of raw data such as images, logs, backups, and text. *File storage* can be accessed from any instance in the VCN and supports the Network File System version 3.0 (NFSv3) protocol, allowing its data to be accessed simultaneously by multiple hosts. *Archive storage* is similar to object storage but is designed to store seldomly accessed data for long periods of time, which can drastically reduce storage cost for archived data.

Although inbound data transfer is free in OCI, as opposed to other cloud service providers, the time to migrate large volumes of data to the cloud may be weeks or months, depending on network bandwidth. This is where the *Data Transfer Appliance* is used to securely migrate extremely large volumes of data from on-premises data

centers to OCI in only days. Lastly, the *Storage Gateway* extends the local file systems in on-premises data centers to access object storage in OCI at no extra charge, essentially extending your storage options for your on-premises systems.

In this chapter, we will cover the first four storage options. They can be accessed from the Storage link in the OCI console (see Figure 7-1). The chapter concludes with information on OCI's security-first architecture as it pertains to cloud storage.



**Figure 7-1** Accessing the Storage Options in the Console

## Block Volume

OCI provides the ability to create, attach, resize, and move a *block volume*, also referred to as a *block storage volume*. It is essentially SSD-based storage that can be attached to any available compute instance and generally maintains all software, binaries, and data needed for your application.

Block volumes are the most common storage types that you will use in your cloud journey. The two types of block storage options are block volume and boot volume. Boot volumes are created during the creation of a compute instance and hold the OS binaries. These volumes can be viewed, managed, and terminated in the same way as block volumes and are accessible under the Boot Volumes menu in the service.

## Creating and Attaching

Once a block volume is created and attached to a compute instance, it can be used similarly to a hard drive. When it is detached, the data residing on it remains intact and

can be reattached to the same or another host. When it is detached, charges still accrue on the block volume.

You can attach a volume by navigating to the volume itself and selecting a compute instance, or you can connect to a compute instance and attach the block volume.

## Creating a Block Volume

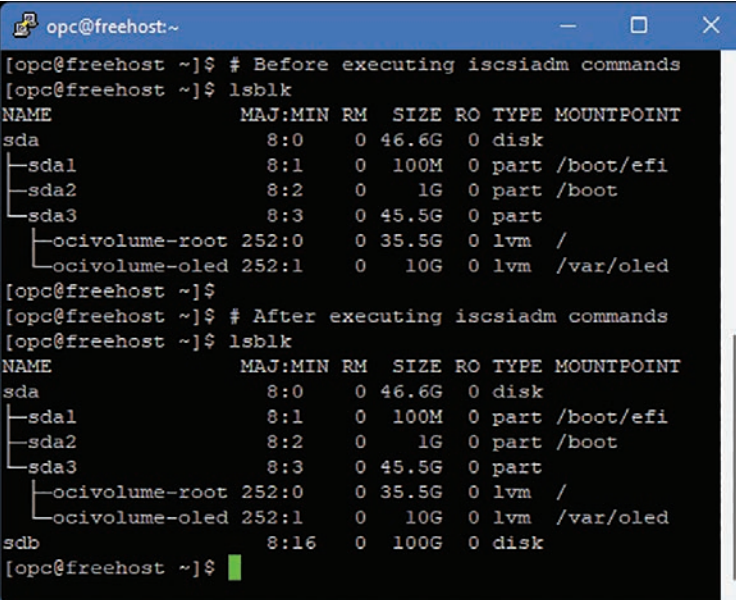
For the quick creation of a 1 terabyte block volume, follow these steps:

- Step 1.** Navigate to the OCI console menu; then click **Storage > Block Volumes**.
- Step 2.** Click **Create Block Volume** to enter the creation wizard.
- Step 3.** Enter the block volume name. All other settings are prefilled with defaults. Customize other settings as needed.
- Step 4.** Click **Create Block Volume** to create the volume.

Compartment, availability domain, size, performance, backup policies, replication, and encryption can be customized as necessary. Block volumes can have a minimum size of 50 GB and a maximum size of 32 TB. All tenancies receive 200 GB of always-free block volume storage.

## Attaching a Block Volume to an Instance

After the block volume is created, its details will immediately appear in the details page, as shown in Figure 7-2.



```

opc@freehost:~$ # Before executing iscsiadm commands
opc@freehost:~$ lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                  8:0      0  46.6G  0 disk 
├─sda1                               8:1      0   100M  0 part /boot/efi
├─sda2                               8:2      0    1G   0 part /boot
├─sda3                               8:3      0  45.5G  0 part 
├─ocivolume-root                    252:0    0  35.5G  0 lvm  /
└─ocivolume-oled                    252:1    0   10G   0 lvm  /var/oled

opc@freehost:~$ # After executing iscsiadm commands
opc@freehost:~$ lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                  8:0      0  46.6G  0 disk 
├─sda1                               8:1      0   100M  0 part /boot/efi
├─sda2                               8:2      0    1G   0 part /boot
├─sda3                               8:3      0  45.5G  0 part 
├─ocivolume-root                    252:0    0  35.5G  0 lvm  /
└─ocivolume-oled                    252:1    0   10G   0 lvm  /var/oled
sdb                                  8:16     0  100G   0 disk 
opc@freehost:~$
  
```

**Figure 7-2** Reviewing the Block Volume Details

The attachment type can either be iSCSI or Paravirtualized. Paravirtualized is simpler to set up because it does not require additional iSCSI configuration commands to be run. iSCSI, however, does not have the virtualization overhead and thus has better input/output operations per second (IOPS) performance than Paravirtualized for larger block volumes, but requires connecting to the host and mounting the volume to make it usable. IOPS is a standard measurement for how many read/write operations a volume can carry out every second. The higher the number, the better the volume performs.

The most common access type is Read/Write, but other options such as Read/Write – Shareable and Read-only – Shareable are available.

From here, it is possible to attach this volume to an instance (the instance must be running) by following these steps:

- Step 1.** Click **Attached Instances**.
- Step 2.** Click **Attach to Instance**.
- Step 3.** Though optional, preferably select a Device Path.
- Step 4.** Select the running instance from the drop-down.
- Step 5.** Click **Attach**.

Even though the volume is attached, since the default attachment type is set to iSCSI, a few extra steps are required:

- a. On the details page of the block volume, click the menu icon and select **iSCSI Commands & Information**, as demonstrated in Figure 7-3.

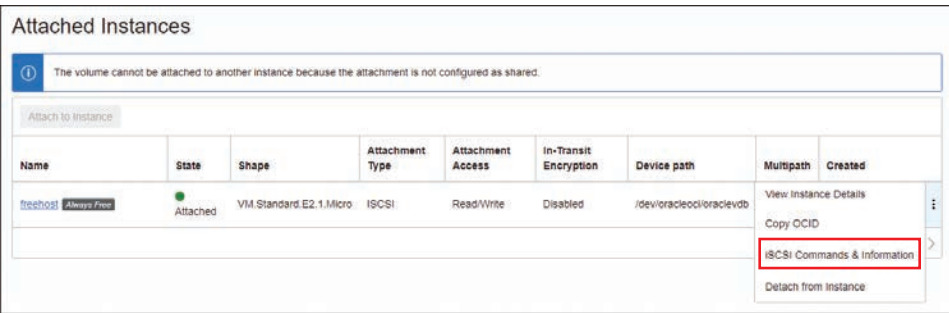


Figure 7-3 Capturing the iSCSI Commands

- b. Copy the commands under the Attach Commands section.
- c. Log in to your host and execute the three **iscsiadm** commands. There is no need to make changes to these commands because they already take into account your system details. Then format and mount the filesystem volume.

Example 7-1 walks through not only an example of executing the iSCSI commands but also the remaining steps required to view, format, and mount the file system.

**Example 7-1** Mounting the iSCSI Volume as root

---

```
# Run the iscsiadm commands below on the compute instance
# The parameters are specific to the volume and instance
# These commands should be copied and executed from the OCI console as is
sudo iscsiadm -m node -o new -T iqn.2015-12.com.oracleiaas:13c56eb9-639a-4fa1-
b8ea-5378b9e0c7f0 -p 169.254.2.5:3260
sudo iscsiadm -m node -o update -T iqn.2015-12.com.oracleiaas:13c56eb9-639a-4fa1-
b8ea-5378b9e0c7f0 -n node.startup -v automatic
sudo iscsiadm -m node -T iqn.2015-12.com.oracleiaas:13c56eb9-639a-4fa1-b8ea-
5378b9e0c7f0 -p 169.254.2.5:3260 -l

# View available disks
# The volume will appear as a 'disk' with the correct size shown
lsblk

# Confirm that the file system is a 'data' volume
file -s /dev/sdb

# One time creation of the file system as ext4; all data will be lost
mkfs.ext4 /dev/sdb

# Create a local directory in which the file system will be mounted to
mkdir /u01

# Mount the filesystem
mount /dev/sdb /u01

# Confirm that file system is mounted and available
df -h
```

---

## Configuring Performance

The performance levels of a block volume can be changed at any time, but it must be detached and reattached to the instance. Oracle delivers consistent, low-latency performance of up to 200 IOPS/GB to a maximum of 300,000 IOPS and 2,680 MBps of throughput per volume.

By default, block volumes are created with a “balanced” performance level. The balanced level operates at 10 VPUs. *VPU* stands for volume performance unit and indicates how much performance resources are allocated to a volume. At 10 VPUs, the block volume is expected to deliver 6,000 IOPS with a throughput of 48 MBps. At 100 VPUs, it will deliver 19,500 IOPS with a throughput of 156 MBps.

How does changing the VPU fare in the real world? For example, moving from a 10 VPU volume to a 20 VPU volume can yield a 16 percent improvement in I/O performance. The balanced performance level is generally a good level to start at unless you start experiencing or have high performance I/O needs.