Oracle® Cloud

Infrastructure Classic Performance Tuning Best Practices

E83035–04
October 2018

Oracle Cloud Infrastructure Classic Performance Tuning Best Practices,

E83035–04

Primary Author: Sreevalli Setty

Contributors: Nikolai Potapov, Deb Bhattacharjee, Tulasiram Kodali

Contents

Preface

Oracle Cloud Infrastructure Classic Performance Tuning Best Practices includes best practices for migrating or extending existing deployments into Oracle Cloud Infrastructure Classic including Compute systems, Databases, Oracle GoldenGate, Oracle Business Intelligence Cloud Service, Oracle Essbase, Java Virtual Machines, Oracle Traffic Director, and Node.js applications. This document is also applicable to the Oracle Cloud at Customer deployments.

The following best practices and parameter settings are based on our experience with various applications in the lab. We recommend that you thoroughly test these out for your application in your environment before going to production. The best practices stated in this guide pertain to Oracle Cloud Infrastructure Compute Classic.

Topics:

Audience

Migrating to Oracle Cloud Infrastructure Classic: Performance Tuning Best Practices is intended for DevOps Engineers, System Administrators, Storage Administrators, Developers, Database Administrators, Business Intelligence (BI) Administrators, and Essbase Administrators who wants to migrate or extend their existing deployments in Oracle Cloud Infrastructure Classic or want to move their on-premises systems and resources to Oracle Cloud at Customer.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc`.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` or visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` if you are hearing impaired.

Related Resources

For more information, see the Oracle Cloud website at `http://cloud.oracle.com`.

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.
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.

What's New in this Release

This topic lists the changes to the sections in this document.

The following table lists the newly added sections of the previous version (Document ID: E83035-03) released in January 2018.

October 2018

Topic Title	Type of Change
Optimize entropy for TLS/SSL-based application	New
Best Practices for Ubuntu	New
Apply Ring Buffer Fix to Improve IO and Stability for Custom Ubuntu Images	New
Optimize Essbase Dataload In Hybrid Deployment	New
Optimize MTU	New
Update and Use the Latest Paravirtual Driver	Update
OS Oracle Linux Image	Update
Known Issue: Reduced SMB Bandwidth After Installing a Windows Update (KB4038792)	Update
Ensure that RDP Access Through Windows Firewall Is Enabled and Check AD Group Policies	Update
Performance Impact on Oracle Analytics Cloud Service Queries from Network Latency	Update

1

Running Workloads on Oracle Cloud Infrastructure Compute Classic

It is recommended that you follow the performance best practices outlined in this topic when you are creating compute resources like Oracle Cloud Infrastructure Compute Classic instances and block storage on Oracle Cloud. Apply the configurations listed in the following checklist before you start installing or setting up any applications on the Oracle Cloud Infrastructure Compute Classic instances:

Quick Checklist

Platform	Area	Optimizations	Priority	Persona (DevOps/Developer/ Admin)
Linux	Image	• OS Oracle Linux Image • Apply Ring Buffer Fix to Improve IO and Stability for UEK3 and UEK4-Based Images	Mandatory	• DevOps Engineers and System Administrators • DevOps Engineers, System Administrators, and Storage Administrators
Windows	Image	• Update and Use the Latest Paravirtual Driver • Set the Disk Time-Out Value to a Higher Value • Disable SIPS on Older Windows Images • Check Boot Storage • Check the Disk Health After Applying the Changes • Optimize MTU	Mandatory	• DevOps Engineers and System Administrators

Platform	Area	Optimizations	Priority	Persona (DevOps/Developer/ Admin)
		• Ensure that RDP Access Through Windows Firewall Is Enabled and Check AD Group Policies • Create Additional Administrator Accounts • Disable Guest Account Access	Required	
		• Known Issue: Reduced SMB Bandwidth After Installing a Windows Update (KB4038792)	Good to have	
SPARC	Image	• Performing Additional Operating System Tuning When Using SPARC 300	Required	• DevOps Engineers and System Administrators
Ubuntu	Image	• Apply Ring Buffer Fix to Improve IO and Stability for Custom Ubuntu Images	Required	• DevOps Engineers and System Administrators
Compute Generic	Storage and IO	• Using the Default Stripe Size While Creating and Binding a Volume Group • DOS-Compatible Mode Disabled for Storage Volumes • Using SSD-Based Storage Volumes to Optimize IO Bandwidth	Required	• DevOps Engineers, System Administrators, and Storage Administrators

Platform	Area	Optimizations	Priority	Persona (DevOps/Developer/ Admin)
Compute Generic	Network	• Placement of Services to Reduce Network Latency • Disabling IPv6 to Solve Connectivity Issues in Applications	Required	• DevOps Engineers and System Administrators
Compute Generic	Orchestration	• OS Bootable Storage Volume • High Availability and Instance Anti-affinity	Required	• DevOps Engineers, System Administrators, and Storage Administrators • DevOps Engineers and System Administrators
Compute Generic	Other best practices	• SSH Tunneling for Secure Connectivity and Two-Way Communication • Runtime Tools for Monitoring System and Processes • LUKS for File Encryption at Rest • Data Migration	Good to have	• DevOps Engineers and Developers • DevOps Engineers, System Administrators, and Developers • DevOps Engineers and Developers • DevOps Engineers, System Administrators, and Storage Administrators

Best Practices for Linux OS

The following settings are required when running Oracle Linux guest virtual machines in Oracle Cloud Infrastructure Classic.

> **Important:**
>
> Before carrying out the procedures outlined in this section, it is strongly recommended to take the following precautionary actions:
>
> - Take a snapshot of your existing storage boot volume. See Backing Up and Restoring Storage Volumes Using Snapshots in *Using Oracle Cloud Infrastructure Compute Classic*.
> - Reboot the VM.
> - Shutdown all running applications.

- OS Oracle Linux Image
- Apply Ring Buffer Fix to Improve IO and Stability for UEK3 and UEK4-Based Images

> **Note:**
>
> Failure to follow these recommendations could result in instability, intermittent network connectivity issues, or in certain cases, loss of data.

OS Oracle Linux Image

Oracle recommends that you use Oracle Linux (OL) version 6.8 while creating new VMs, since it is the most stable one. OL 6.6 and later contain certain IO optimizations that are critical for IO performance in the Oracle Cloud Infrastructure Compute Classic service. The latest versions of the Unbreakable Enterprise Kernel (UEK), namely UEK3 and UEK4, offer additional optimizations that improve IO and provide stability.

> **Note:**
>
> Kernel upgrades should only be done on IaaS-based services. Do not perform kernel-upgrade for PaaS services to guarantee service integrity. You must follow the PaaS-specific OS patching procedures in a controlled manner to guarantee service integrity. Check for PaaS-specific instructions before updating the kernel version, as they may vary from service to service.

Oracle recommends you to regularly update the kernel to the latest available version during the lifetime of the VM so that you have the latest bug fixes and patches installed for optimum stability. To update the kernel-UEK on an existing OL6 UEK4 Compute VM:

1. Edit the `/etc/yum.repos.d/public-yum-ol6.repo` entry and enable the corresponding UEK4 repository for OL 6; for example, ol6_UEKR4, if it is not enabled already.

   If the file is unavailable on your system, download it from the Oracle Linux Yum Server by following the procedure described in Downloading the Oracle Linux Yum Server Repository Files.

2. Check the UEK version by entering the following command:

   ```
   uname -r
   ```

   Sample output:

   ```
   4.1.12-61.1.14.el6uek.x86_64
   ```

3. Check the available update by entering the following command:

   ```
   yum check-update kernel-uek
   ```

   Sample output:

   ```
   4.1.12-103.9.2.el6uek
   ```

4. Update the UEK by entering the following command:

   ```
   yum update kernel-uek
   ```

5. Check `/etc/grub.conf` has corresponding `initrd/initramfs` line for newly added or upgraded kernel.

   Example: `initrd /initramfs-4.1.12-103.9.2.el6uek.x86_64.img`

6. Reboot the VM by entering the following command:

   ```
   reboot
   ```

7. Check the UEK version by entering the following command:

   ```
   uname -r
   ```

   Sample output:

   ```
   4.1.12-103.9.2.el6uek
   ```

> **Note:**
>
> - It is recommended to use UEK3 version v3.8.13-79 and above.
> - It is recommended to use UEK4 version v4.1.12-112 and above.

See About Oracle-Provided Linux Images in *Using Oracle Cloud Infrastructure Compute Classic*.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Apply Ring Buffer Fix to Improve IO and Stability for UEK3 and UEK4-Based Images

The latest Oracle Linux (OL), UEK3 (kernel version 3.8.13-8.2.2.3.el6uek.x86_64 and higher) and UEK4 (kernel version 4.x and higher) images have a ring buffer fix that significantly improves the IO performance and stability of a system running OL. Your OS images may be missing these critical IO settings and can therefore suffer from issues such as lower IO bandwidth and/or compute instance stability problems including VM hang in rare cases.

Applying Ring Buffer Fix on UEK3

Prerequisite: The kernel version should be 3.8.13-68.2.2.3.el6uek.x86_64 or higher. Use the `uname -r` to determine the kernel version.

Next,

1. Verify whether the active GRUB configuration file (`/etc/grub.conf` for OL 6 or `/boot/grub2/grub.cfg` for OL 7) contains the following kernel parameters:
   - `xen-blkfront.max_ring_page_order=4`
   - `xen-blkfront.max=8`
2. Check whether the following file exists: `/etc/udev/rules.d/13-block-devices-nr-requests.rules`.

   If the file doesn't exist, then create it with the following content:

   ```
   RUN+="/bin/sh -c 'echo 4096 > /sys/$devpath/queue/nr_requests'"
   ```
3. Check whether the following file exists: `/etc/udev/rules.d/99-blksize.rules`.

   If the file doesn't exist, then create it with the following content:

   ```
   KERNEL=="xvd*", RUN+="/bin/sh -c 'echo 32 > /sys$DEVPATH/queue/
   max_sectors_kb'"
   ```
4. Change the access mode for the files that you just created by entering the following command:

   ```
   chmod 644 /etc/udev/rules.d/13-block-devices-nr-requests.rules /etc/udev/
   rules.d/99-blksize.rules
   ```
5. Restart the VM.

Applying Ring Buffer Fix on UEK4

Prerequisite: The kernel version should be 4.x or higher. Use the `uname -r` to determine the kernel version.

Next,

1. Verify whether the active GRUB configuration file (`/etc/grub.conf`) file contains the following kernel parameters:
   - `xen-blkfront.max_ring_page_order=4`
   - `xen-blkfront.max_indirect_segments=8`
   - `xen-blkfront.max_queues=2`

- `xen-netfront.max_queues=2`

2. Check whether the following file exists: `/etc/udev/rules.d/13-block-devices-nr-requests.rules`.

   If the file doesn't exist, then create it with the following content:

   ```
   RUN+="/bin/sh -c 'echo 4096 > /sys/$devpath/queue/nr_requests'"
   ```

3. Check whether the following file exists: `/etc/udev/rules.d/99-blksize.rules`.

   If the file doesn't exist, then create it with the following content:

   ```
   KERNEL=="xvd*", RUN+="/bin/sh -c 'echo 32 > /sys$DEVPATH/queue/
   max_sectors_kb'"
   ```

4. Change the access mode for the files that you just created by entering the following command:

   ```
   chmod 644 /etc/udev/rules.d/13-block-devices-nr-requests.rules /etc/udev/
   rules.d/99-blksize.rules
   ```

5. Restart the VM.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Best Practices for Windows

The following settings are required when running Windows guest virtual machines in Oracle Cloud Infrastructure Classic. Some of these settings have been incorporated into the latest Windows Images from the Oracle Marketplace.

> **Important:**
>
> Before carrying out the procedures outlined in this section, it is strongly recommended to take the following precautionary actions:
>
> - Take a snapshot of your existing storage boot volume. See Backing Up and Restoring Storage Volumes Using Snapshots in *Using Oracle Cloud Infrastructure Compute Classic*.
> - Reboot the VM.
> - Shutdown all running applications.

- Update and Use the Latest Paravirtual Driver
- Set the Disk Time-Out Value to a Higher Value
- Disable SIPS on Older Windows Images
- Check Boot Storage
- Ensure that RDP Access Through Windows Firewall Is Enabled and Check AD Group Policies

- Create Additional Administrator Accounts
- Disable Guest Account Access
- Known Issue: Reduced SMB Bandwidth After Installing a Windows Update (KB4038792)
- Check the Disk Health After Applying the Changes
- Optimize MTU

> **Note:**
>
> Failure to follow these recommendations could result in instability, intermittent network connectivity issues, or in certain cases, loss of data.

Update and Use the Latest Paravirtual Driver

Oracle recommends that you browse the Marketplace to obtain newer OS images for Windows platform. The newer images carry the latest optimization including the latest Windows Paravirtual (PV) driver. The PV drivers are high-performance network and disk drivers that significantly reduce the overhead of the traditional implementation of IO device emulation. These drivers provide improved network and disk throughput to run fully-virtualized Windows guests in an Oracle VM Server for x86 environment. If you are using a custom Windows image or an old copy of the image from the Marketplace, update the Paravirtual (PV) driver and also tune your OS similar to the Marketplace object images, because the current Marketplace Windows images, such as 2012R2 or 2008, are stable and tuned for optimum performance. You can achieve improved stability and performance for your compute instance by using the newer OS images from Oracle Marketplace.

Download the PV driver 3.4.2 or newer version if available from My Oracle Support or Oracle Software Delivery Cloud. On Oracle Software Delivery Cloud, the drivers are bundled with the Oracle VM Server media pack. To upgrade the drivers, see Upgrading the Oracle VM Paravirtual Drivers for Microsoft Windows

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Set the Disk Time-Out Value to a Higher Value

Oracle recommends that you check the disk timeout in the Microsoft Registry and modify the current, default TimeOutValue to 300 seconds (recommended) to prevent your Windows instance from hanging or becoming inaccessible with disk and XenVbd events. The default value of 60 seconds is too low to prevent timeouts.

To change the time-out value, go to `Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\disk\TimeOutValue`. Change the registry key from 60 to 300 and restart the VM.

See DocID: **2309363.1**.

Best practice for:

DevOps Engineers and System Administrators

Disable SIPS on Older Windows Images

If you are using an older version of Windows image in your instance, the SIPS property of the Paravirtual (PV) driver is enabled by default. This can cause the power management to intermittently disable network device and cause network interruption. To improve network stability, disable SIPS in the PV network driver by unchecking **Allow the computer to turn off this device to save power** in the network driver properties.

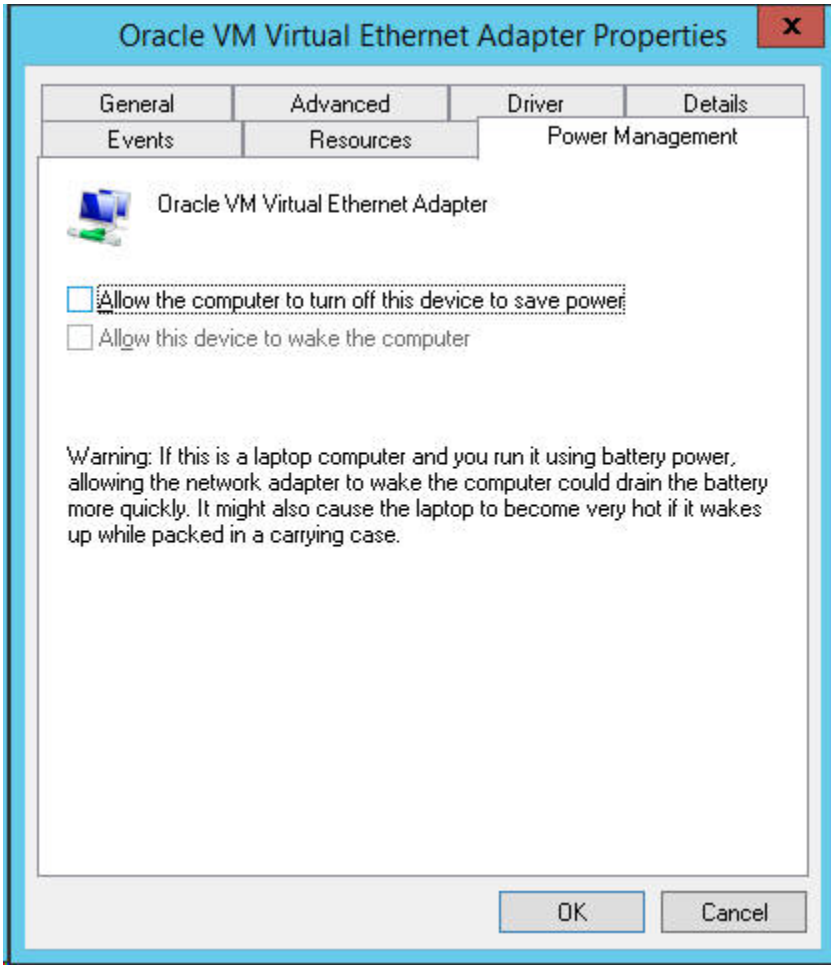

Best practice for:

DevOps Engineers and System Administrators

Check Boot Storage

Oracle recommends that you create the boot storage on a persistent storage volume with adequate free space. Your Windows VM may hang if there is insufficient space on the boot drive.

For more information about creating a boot storage volume, see Creating a Bootable Storage Volume in *Using Oracle Cloud Infrastructure Compute Classic*.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Ensure that RDP Access Through Windows Firewall Is Enabled and Check AD Group Policies

Oracle recommends that you check the group policy configuration of the Windows Active Directory (AD) and your firewall settings on your Windows VM. AD group policies or updating the firewall can disable Microsoft's Remote Desktop Protocol (RDP). Check the group policy configuration and ensure that the Remote Service Management is marked as an exception in the Windows firewall.

The new windows security update to address the **CREDSSP** encryption and **CVE-2018-0886** restricts the RDP access from the non patched Windows client. It is recommended to apply the Windows CredSSP security update to your Windows Client and Server Operating System.

For more information, see the Microsoft Support knowledge base.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Create Additional Administrator Accounts

Oracle recommends that you create additional administrative users for your Windows server, so that if one account gets locked, you can access it using other accounts.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Disable Guest Account Access

As a standard security procedure, disable the guest accounts.

For more information, see the Microsoft Support knowledge base.

Best practice for:

DevOps Engineers and System Administrators

Known Issue: Reduced SMB Bandwidth After Installing a Windows Update (KB4038792)

Server Message Block (SMB) protocol enables file sharing on your Windows machine. However, you may encounter reduced bandwidth after installing a Windows update on your VM, particulary if you are using the SMB protocol. The drop in the bandwidth is because of the Windows' security patch: KB4038792.

Oracle recommends that you retain the installed update and do not deinstall it from your system, as this security patch is designed to eliminate potential vulnerabilities. It is recommended to use robocopy or xcopy to parallelize the work and to improve the bandwidth.

Best practice for:

DevOps Engineers and System Administrators

Check the Disk Health After Applying the Changes

Oracle strongly recommends that you check the health of the disks after applying the new settings, and treat the disks, if needed. If the disks are already corrupt, then your Windows VM may become unresponsive any time, even if you have tuned your OS image, increased the disk timeout setting, or upgraded the PV drivers.

The following image shows an example of a disk corruption error signature.

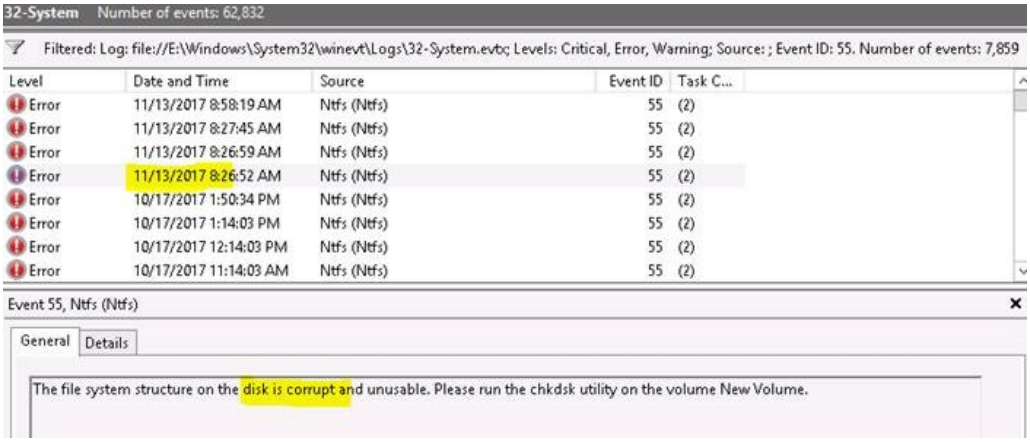

Best practice for:

DevOps Engineers and System Administrators

Optimize MTU

The OPC Infrastructure should be optimized to use MTU 9000.

Oracle recommends you to use the latest OPC Compute images. The latest OPC Compute images use MTU 9000. The older images including the Marketplace and custom images have to be optimized to use MTU 9000. This improves the network bandwidth in the infrastructure components and other Compute VMs.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Best Practices for SPARC

The following settings are required when running SPARC guest virtual machines in Oracle Cloud Infrastructure Classic.

Topics:

Performing Additional Operating System Tuning When Using SPARC 300

If you encounter high network latency in SPARC 300 dedicated compute, you can improve maximum bandwidth by changing the TCP configuration inside the virtual machine.

In addition, you must restart the SSH service every time TCP tuning is applied.

Tuning the TCP send and receive buffers may improve the network throughput.

Example:

The following example lists the sample commands that you can use to perform additional tuning:

```
ipdadm set-prop -p recv_buf=4194304 tcp
ipdadm set-prop -p max_buf=4194304 tcp
ipdadm set-prop -p send_buf=4194304 tcp
svcadm restart ssh
```

In the preceding example, you should evaluate and replace the exact values on a case-to-case basis.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Best Practices for Ubuntu

The following settings are required when running Ubuntu guest virtual machines in Oracle Cloud Infrastructure Classic.

> **Important:**
>
> Before carrying out the procedures outlined in this section, it is strongly recommended to take the following precautionary actions:
>
> - Take a snapshot of your existing storage boot volume. See Backing Up and Restoring Storage Volumes Using Snapshots in *Using Oracle Cloud Infrastructure Compute Classic*.
> - Reboot the VM.
> - Shutdown all running applications.

- Apply Ring Buffer Fix to Improve IO and Stability for Custom Ubuntu Images

> **Note:**
>
> Failure to follow these recommendations could result in instability, intermittent network connectivity issues, or in certain cases, loss of data.

Apply Ring Buffer Fix to Improve IO and Stability for Custom Ubuntu Images

The custom Ubuntu images have a ring buffer fix that significantly improves the IO performance and stability of a system running Ubuntu. Your OS images may be missing these critical IO settings and can therefore suffer from issues such as lower IO bandwidth and/or compute instance stability problems including VM hang in rare cases.

Applying Ring Buffer Fix on Custom Ubuntu Images

Prerequisite: The current recommendation is verified on Ubuntu version should be `vmlinuz-4.4.0-116` release only.

Next,

1. Verify whether the active GRUB configuration file contains the following kernel parameters:
   - `xen-blkfront.max_ring_page_order=4`
   - `xen-blkfront.max=8`
2. Check whether the following file exists: `/etc/udev/rules.d/13-block-devices-nr-requests.rules`.

   If the file doesn't exist, then create it with the following content:

```
RUN+="/bin/sh -c 'echo 4096 > /sys/$devpath/queue/nr_requests'"
```

3. Check whether the following file exists: `/etc/udev/rules.d/99-blksize.rules`.

   If the file doesn't exist, then create it with the following content:

   ```
   KERNEL=="xvd*", RUN+="/bin/sh -c 'echo 32 > /sys$DEVPATH/queue/
   max_sectors_kb'"
   ```

4. Change the access mode for the files that you just created by entering the following command:

   ```
   chmod 644 /etc/udev/rules.d/13-block-devices-nr-requests.rules /etc/udev/
   rules.d/99-blksize.rules
   ```

5. Restart the VM.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Best Practices for Compute Generic

The following settings are required in general with all the guest virtual machines in different operating system used in Oracle Cloud Infrastructure Classic.

Storage and IO:

- Using the Default Stripe Size While Creating and Binding a Volume Group
- DOS-Compatible Mode Disabled for Storage Volumes
- Using SSD-Based Storage Volumes to Optimize IO Bandwidth

Network

- Placement of Services to Reduce Network Latency
- Disabling IPv6 to Solve Connectivity Issues in Applications

Orchestration

- OS Bootable Storage Volume
- High Availability and Instance Anti-affinity

Other Best Practices

- SSH Tunneling for Secure Connectivity and Two-Way Communication
- Runtime Tools for Monitoring System and Processes
- LUKS for File Encryption at Rest
- Data Migration

Using the Default Stripe Size While Creating and Binding a Volume Group

When binding multiple storage volumes into a volume group, changing the default stripe size and making it too small, such as 4KB, may lead to a situation where striped data won't be properly aligned with the rest of the block storage infrastructure.

This could lead to noticeable IO bandwidth degradation, especially for IO-intensive workloads that involve random data access.

Oracle recommends that you use the default stripe size of 64KB when binding multiple storage volumes into a volume group.

> **Best practice for:**
>
> DevOps Engineers, System Administrators, and Storage Administrators

DOS-Compatible Mode Disabled for Storage Volumes

For better IO performance, disable DOS-Compatible mode.

Creating disk partitions in DOS-Compatible mode can lead to a performance impact. Most disks, including those used in the Oracle Cloud Infrastructure Compute Classic instances, have 4096–byte sectors. When a partition is created in DOS-compatible mode, it attempts to start the partition at what would have been the second cylinder of a disk with 512–byte sectors, which does not fall on an even physical sector boundary for a disk with 4096–byte sectors. Every logical read or write on such a partition ends up being two physical writes, with the first part of the logical sector in part of one physical sector and the second part in the next physical sector. The result is a doubling of physical IO versus non-DOS compatible disks and a corresponding decrease in performance.

See Doc ID 2098513.1.

> **Best practice for:**
>
> DevOps Engineers, System Administrators, and Storage Administrators

Using SSD-Based Storage Volumes to Optimize IO Bandwidth

You can optimize the IO Compute workloads by creating storage volumes on Solid-State Drive (SSD) filers instead of a Hard Disk Drive (HDD) filer.

Some regional sites offer SSD-based filers in addition to HDD-based filers to create storage volumes. In addition to high IO performance, such SSD volumes provide same benefits in regard to data durability, failover, and HA deployment.

If your region doesn't offer an SSD-based filer option, or if you need to deploy your application using a non-persistent storage option, then use the local NVMe SSD disks

available in high IO shapes to optimize the IO workload. If you choose to use non-persistent storage option, then you must devise a proper strategy to have data durability, failover, or HA for your data.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Placement of Services to Reduce Network Latency

Reduced network latency between components can improve performance. If network latency is critical for performance, then place all services in the same data center, if your account allows.

For example, for the Oracle E-Business Suite application, the database and the middle tier should not be in two different data centers.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

Disabling IPv6 to Solve Connectivity Issues in Applications

Some applications may behave unpredictably when both IPv6 and IPv4 are enabled in your operating system.

For example, the curl utility may fail with a timeout, and it works only when you use IPv4 (`curl -4`). Another example may be a custom Java application cluster that is having trouble communicating between its nodes unless you explicitly set the IPv4 usage preference (`-Djava.net.preferIPv4Stack=true`).

IPv6 is enabled by default and it can cause connectivity issues in your application. A suggested workaround is to disable IPv6. Following are the sample commands to temporarily disable IPv6:

```
# sysctl -w net.ipv6.conf.all.disable_ipv6=1
# sysctl -w net.ipv6.conf.default.disable_ipv6=1
```

To make this change persistent, edit the sysctl.conf file located in the `/etc` directory as shown in the following sample and reboot the system.

```
net.ipv6.conf.all.disable_ipv6 = 1
net.ipv6.conf.default.disable_ipv6 = 1
```

> **Best practice for:**
>
> DevOps Engineers and System Administrators

OS Bootable Storage Volume

Use a bootable storage volume for a boot disk instead of local, ephemeral volumes.

While creating an Oracle Cloud Infrastructure Compute Classic instance, you can have a bootable volume either on a local, ephemeral disk or on persistent block storage. Local drives are ephemeral drives. If you use them, you may lose boot volume data on the Oracle Cloud Infrastructure Compute Classic instance restarts, and your orchestration may have to recreate the image, hence your restart time could be slower. An Oracle Cloud Infrastructure Compute Classic instance reboot doesn't guarantee that the same ephemeral disk will be mounted.

See Creating a Bootable Storage Volume in *Using Oracle Cloud Infrastructure Compute Classic*.

> **Best practice for:**
>
> DevOps Engineers, System Administrators, and Storage Administrators

High Availability and Instance Anti-affinity

If you want to configure high availability (HA) against single point of failure, then you can use relationship placement tags. The relationship placement tags can improve your chances against single point of failure by placing the related VMs on two different physical nodes.

- Consider using the anti-affinity placement tag to make sure VMs from the same cluster run on different physical nodes to provide HA in case of physical node failure.
- If you are using Orchestration v1, then you should also use "`ha_policy`" flag to ensure that your VM gets recreated or restarted after the failure. This is a default behavior in Orchestration v2. See Orchestration v1 Attributes for launchplan in *Using Oracle Cloud Infrastructure Compute Classic*.

Example:

```
"relationships": [
    {
        "instances": [
            "A",
            "B"
        ],
        "type": "different_node"
     }
]
```

In this example, two instances A and B are created in a single orchestration. In the relationships section of the orchestration, the relationship type specified is different node, indicating that instance A and instance B should be created on different physical nodes.

> **Note:**
>
> In some rare cases, when you need to achieve the best possible network latency in between two VMs, use a different tag such as "`same_node`" to request VMs placement on the same physical node, if the capacity allows it. This solution works best on Dedicated Compute deployments — where you have better control over capacity.

> **Best practice for:**
>
> DevOps Engineers and System Administrators

SSH Tunneling for Secure Connectivity and Two-Way Communication

Use SSH tunneling to set up ad-hoc and secure connectivity between an on-premises tool, such as SQL Developer and the services running on Oracle Cloud Infrastructure Compute Classic. Use two-way SSH tunneling only when there's a need for a two-way communication or a VPN connection is not ready or if VPN reconfiguration takes a long time. Use an autossh package to establish a two-way channel that is more persistent and self-healing in case of a network failure. Autossh package also monitors the state of the channel and also automatically restarts them if needed.

> **Note:**
>
> SSH tunneling is not a suitable solution for massive, bulk-data transfers.

All applications' traffic is transported through an SSH session channel, and there's no need to open more ports on Oracle Cloud Infrastructure Compute Classic VMs other than the SSH port (`22`).

One-Way SSH Tunneling (for secure communication from an on-premise application to Oracle Cloud Infrastructure Compute Classic instance)

Use one-way SSH tunneling to create an SSH tunnel connecting the SQLDeveloper tool on a client desktop to a database running on Oracle Cloud Infrastructure Compute Classic.

Syntax (Local)

```
 ssh -L
<local_host>:<local_host_port>:<remote_host>:<remote_host_port> <opcuser>@<IPAdress_of_OPC>
```

Example

```
ssh -L localhost:1521:localhost:1521 oracle@x.x.x.x
```

Two-Way SSH Tunneling (interim solution while corporate VPN is not yet set up)

Use two-way SSH tunneling to immediately connect to Oracle Cloud Infrastructure Compute Classic instance while the corporate VPN is still being configured.

Syntax

```
 ssh -L -R
<local_host>:<local_host_port>:<remote_host>:<remote_host_port> <opcuser>@<IPAdress_o
f_OPC>
```

Example

The following example shows how the internal traffic coming to port 9002 on Server A is forwarded to port 9001 on Server B:

```
ssh -L localhost:9002:localhost:9001 oracle@x.x.x.x
```

The following example shows how the internal traffic coming to port 9002 on Server B is forwarded to port 9001 on Server A:

```
ssh -R localhost:9002:localhost:9001 oracle@x.x.x.x
```

Two-Way SSH Tunneling with autossh Package (self-healing in case of network failure)

Use two-way SSH tunneling with autossh package to achieve extra stability to overcome possible network failures and restart SSH channels automatically in case of a network failure, in addition to connecting to an Oracle Cloud Infrastructure Compute Classic instance while the corporate VPN is still being configured.

Example

```
autossh -M 0 -q -f -N -o "ServerAliveInterval 60" -o "ServerAliveCountMax 3" -L
localhost:9002:localhost:9001 oracle@x.x.x.x
autossh -M 0 -q -f -N -o "ServerAliveInterval 60" -o "ServerAliveCountMax 3" -R
localhost:9002:localhost:9001 oracle@x.x.x.x
```

> **Best practice for:**
>
> DevOps Engineers and Developers

Runtime Tools for Monitoring System and Processes

Continuous monitoring of system and application resource usage is critical for any production system. It's important to install some form of monitoring tools inside your Compute Cloud Service instance and ensure that such tools are running in the background and saving logs for certain periods of time.

Oracle recommends that you use at least some basic monitoring tools, such as oswatcher, for monitoring system and process behaviors. See **Doc ID 301137.1**.

> **Best practice for:**
>
> DevOps Engineers, System Administrators, and Developers

LUKS for File Encryption at Rest

If you need to have encryption at rest in a raw Compute Cloud Service instance, you can use the Linux Unified Key Setup (LUKS) technique.

Many PaaS services have encryption at rest already. Sometimes it's required to set up an encrypted data landing page in a raw Compute Cloud Service instance. That's where LUKS can be used.

> **Best practice for:**
>
> DevOps Engineers and Developers

Data Migration

If you plan on moving a large amount of data from an on-premises system to Oracle Cloud Infrastructure Classic, the following tips can help improve the performance of data migration.

- Consider using parallel threads to move data in multiple pipes. For example, use multiple Secure Copy (SCP) processes to send multiple files in parallel as opposed to sending a single file in a single SCP process.

  If you send data in a single pipe, you may face various limitations, such as single threads read or write bandwidth or single session network cap. Using multiple threads or processes helps to solve it.
- Copying files into a latency volume gives better performance than copying files to a throughput volume.
- If you've configured VPN, consider an alternative, yet secure approach to move big volumes of data, because VPN may significantly slow down data transfer speed.

  VPN works as designed for normal daily application data flows, but it can heavily affect initial, and often, large volume data transmission.

  You can use Oracle Storage Server (OSS) as intermediate storage connecting directly to Oracle Cloud Infrastructure Compute Classic instance, and not through VPN. OSS endpoints are already secured, because OSS uses the HTTPS protocol.

> **Best practice for:**
>
> DevOps Engineers, System Administrators, and Storage Administrators

Optimize entropy for TLS/SSL-based application

Optimizing Entropy helps in improved performance of TLS/SSL during heavy workloads.

Applications working on Cryptography such as TLS/SSL use random numbers taken from the entropy pool. Lack of entropy can have a negative impact on performance and security. This might result in low entropy inside the VM and lead to a delay until there is enough data. This results in slow TLS connections.

In order to reduce the lag and increase the entropy pool, install **rngd rpm**. This helps in improved performance of TLS/SSL during heavy workloads.

Best practice for:

DevOps Engineers and System Administrators

2

Running Java on Oracle Cloud Infrastructure Compute Classic

Some performance best practices should be followed when you are creating JVM resources on Oracle Cloud Classic. Apply the configurations listed in the following checklist before you start installing or setting up any applications on the Oracle Cloud Infrastructure Compute Classic instances:

Quick Checklist

Area	Optimizations	Persona (DevOps/Developer/Admin)
VM Size	• Garbage Collection Tuning Affected by Number of VCPUs in Oracle Cloud Infrastructure Compute Classic Service Instance • Small Compute Shape and Heap Size for WebLogic Administration Server • Maximum Heap Size for Java Process	• Developers • Developers • Developers

Garbage Collection Tuning Affected by Number of VCPUs in Oracle Cloud Infrastructure Compute Classic Service Instance

The number of VCPUs in an Oracle Cloud Infrastructure Compute Classic service instance affects the tuning of garbage collection. You must carefully tune the garbage collection threads in Oracle Private Cloud environment because the VM shapes in Oracle Private Cloud may differ from the on-premise source environment. The process of tuning garbage collection is the same between on-premises environment and Oracle Cloud Infrastructure Compute Classic .

If you have small heaps on large shapes, limiting the number of parallel threads can also help performance. A 2GB heap cannot be effectively parallelized beyond about 8 garbage collection threads.

Best practice for:

Developers

Small Compute Shape and Heap Size for WebLogic Administration Server

You can place your WebLogic Administration Server in a smaller-shaped Compute Cloud Service instance.

Oracle WebLogic Server comes with Managed Servers as well as an Administration Server. In most cases, the Administration Server is not running any large application, so there is no need to set a large maximum heap size for this server. If your topology has a separate Compute Cloud Service instance for the WebLogic Administration Server, then you should consider allocating a smaller memory shape for it.

> **Best practice for:**
>
> Developers

Maximum Heap Size for Java Process

Tune the maximum heap size for your Java process according to the maximum operating system memory available in the shape of your Compute Cloud Service instance.

The actual process memory utilized by the JVM process can be as high as 1.8 times to 2 times that of your maximum heap size defined by the `-Xmx` setting. Also, you need to take into account any other processes running in the same Compute Cloud Service Instance.

A large heap size may end up swapping in the OS, which can significantly degrade performance.

> **Note:**
>
> It is recommended to compress large files or heap dumps before transferring them or when they are not in use. This saves time in transfer and reduces the space on disk as well.

> **Best practice for:**
>
> Developers

3

Running Oracle Databases on Oracle Cloud Infrastructure Compute Classic

Optimizations can help database performance in Oracle Cloud Infrastructure Classic. You can use a database instance as part of a database service or create a database manually in a Oracle Cloud Infrastructure Compute Classic instance. If you create your own database, you need to apply certain optimizations.

Quick Checklist

Area	Optimizations	Priority	Persona (DevOps/Developer/ Admin)
Operating System Kernel	• Oracle Database Kernel Parameter Tuning	Required	Database Administrators
Generic Performance	• Configuration of HugePages for Oracle in Linux • Oracle Row Prefetching • Best Practices for Using Fast Recovery Area	Required	• Database Administrators • Developers • Database Administrators
IO	• Oracle DB Redo Logs and Data Files Mapped to Latency Volumes • Oracle init.ora IO Performance Parameters • Improving Block IO Performance	Mandatory	Database Administrators
Backup and Recovery	• Disaster Recovery • RMAN Backup	Required	Database Administrators
Feature-specific	• Recommended DBaaS Software Edition • Oracle init.ora Diagnostics and Tuning Pack for Enterprise Edition • Best Practices for Using Fast Recovery Area	Required	Database Administrators

Oracle Database Kernel Parameter Tuning

Tune database-related kernel parameters.

See **Configuring Kernel Parameters and Resource Limits** in *Oracle Database Quick Installation Guide*.

Many of these parameter values are the defaults for DBaaS, but you may have to modify the values if you manually install the database on IaaS.

> **Best practice for:**
>
> Database Administrators

Configuration of HugePages for Oracle in Linux

Enable the HugePages feature for a database so the operating system can support memory pages larger than the default (usually 4 KB). Other application databases might also benefit with HugePages.

The HugePages feature is integrated into the Linux kernel 2.6. Using very large page sizes can improve system performance by reducing the amount of system resources required to access page table entries. HugePages is useful for both 32-bit and 64-bit configurations. HugePage sizes vary from 2 MB to 256 MB, depending on the kernel version and the hardware architecture. For Oracle Databases, using HugePages reduces the operating system maintenance of page states and increases the TLB (Translation Lookaside Buffer) hit ratio.

See Overview of HugePages in *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems* or HugePages on Oracle Linux 64-bit (Doc ID 361468.1).

> **Best practice for:**
>
> Database Administrators

Oracle Row Prefetching

Query performance may benefit if you increase the number of rows fetched from the database by each fetch call.

Increasing the Oracle row prefetch size helps in performance if you have relatively large latency between the database and middle tier (more than 20 msec) and your application is fetching more than 100 rows in result sets for your queries.

Increasing the Oracle row prefetch can help for an OLAP application and not necessarily for OLTP applications.

The default value of the JDBC prefetch size is 10. You can increase it to a large value, but increasing beyond 1000 is not recommended because the memory footprint in the

middle tier can grow. See **Class OracleDriver** in the *Oracle Database JDBC Java API Reference*.

> **Best practice for:**
>
> Developers

Oracle DB Redo Logs and Data Files Mapped to Latency Volumes

For improved IO performance, ensure that redo logs and temporary data files are placed in latency volumes.

DBaaS by default uses latency volumes for redo logs and temporary data files. If you install Oracle Database manually on IaaS, place the redo and temporary files in latency volumes. It is also a good practice to keep the redo and temporary (temp) files separate from the data files. See Store Data Files Separate from Redo Log Files in *Oracle Database Administration Guide*.

> **Best practice for:**
>
> Database Administrators

Oracle init.ora IO Performance Parameters

Ensure that the `filesystemio_options=SETALL` parameter is set for better database IO performance.

DBaaS has this parameter setting by default. If you manually create a database on Oracle Cloud Infrastructure Compute Classic, you must ensure that the `filesystemio_options` parameter is set to `SETALL`. IO can severly degrade if you do not set this parameter.

> **Best practice for:**
>
> Database Administrators

Disaster Recovery

Data Guard and Active Data Guard provide disaster recovery (DR) for databases with recovery time objectives (RTOs) that cannot be met by restoring from backup.

See **Disaster Recovery to the Cloud with Data Guard and Active Data Guard**.

RMAN Backup

Use RMAN parallelism and compression to speed up cloud backups and restores.

> **Note:**
>
> RMAN parallel requires Enterprise Edition.

See Doc ID 2078576.1

- Set parallelism (for instance, `10`):

  ```
  CONFIGURE DEVICE TYPE SBT PARALLELISM 10;
  ```

- Do not enable tracing (that is, no `_OPC_TRACE_LEVEL` parameter specified in the Oracle Cloud Infrastructure Compute Classic configuration file).
- If the database size is big in TBs, then use the `SECTION SIZE` option in the backup command to parallelize big file tablespace data files. For instance:

  ```
  BACKUP DEVICE TYPE sbt DATABASE TAG 'MYBACKUP' SECTION SIZE 100g;
  ```

  The preceding command will allocate 10 channels and split big files into 100 GB backup pieces.

See Best Practices to Optimize Cloud Backup and Restore Rates in *Using Oracle Database Backup Cloud Service*.

> **Best practice for:**
>
> Database Administrators

Recommended DBaaS Software Edition

Use the DBaaS High Performance version for production because it comes with a Diagnostics and Tuning pack that helps in debugging performance issues.

While creating a DBaaS instance, for **Software Edition** choose the **Enterprise Edition - High Performance** version, or above, for production installations. Standard Edition lacks AWR and other diagnostic and tuning features.

> **Best practice for:**
>
> Database Administrators

Oracle init.ora Diagnostics and Tuning Pack for Enterprise Edition

Verify that `control_management_pack_access = DIAGNOSTIC+TUNING`, to enable the Diagnostics and Tuning pack in Enterprise Edition.

If you create the database manually, you may have to change this setting.

> **Best practice for:**
>
> Database Administrators

Improving Block IO Performance

If you notice poor IO performance on the Compute storage volumes of your Database Cloud Service deployment, you can download and run a script to enable the multi-page ring support feature of Linux to improve performance and OS stability. Some of the latest DBaaS images may not have the multi-page ring support feature enabled.

Configuration settings of certain Linux OS features on existing and newly created database deployments can cause noticeably slower IO, system instability and, in certain extreme cases, system hang. Oracle strongly recommends that you perform the steps provided in Steps needed to stabilize Linux OS on existing and newly created deployments to alter these settings, both on the compute nodes of your existing database deployment and on the compute nodes of any database deployments you create.

> **Best practice for:**
>
> Database Administrators

Best Practices for Using Fast Recovery Area

By default, DBaaS has a retry timeout of 60 seconds. However, if you create a database on a compute VM, you must manually set the timeout value to avoid excessive IO usage when the Fast Recovery Area disk runs low on free space. When you configure the archive logs to use Fast Recovery Area, make sure that it never runs out of disk space and also increase the log archiver retry timeout.

Set the database initialization parameter to 60 seconds.

```
alter system set log_archive_dest_1="LOCATION=USE_DB_RECOVERY_FILE_DEST REOPEN=60";
```

> **Best practice for:**
>
> Database Administrators

4

Running MySQL Databases on Oracle Cloud Infrastructure Compute Classic

Optimizations can help database performance in the Cloud. You can use a database instance as part of a database service or create a database manually in a Oracle Cloud Infrastructure Compute Classic instance. If you create your own database, you need to apply certain optimizations.

Quick Checklist

Area	Optimizations	Persona (DevOps/Developer/Admin)
Generic Performance	• Database Performance Bottleneck with UTF8 Character Set • Configure Data and Transaction Logging on Latency LUNs • Increase Flush Thread Count to Accelerate the Flushing of Dirty Pages • Tune the `innodb-io-capacity` and `innodb-io-capacity-max` Parameters for Maximum IO • Use jemalloc Memory Allocator for CPU-Intensive Workloads • Disable Binary Logging	Database Administrators

Database Performance Bottleneck with UTF8 Character Set

Use a different character set than the UTF8 character set, unless you need it, because UTF8 has a performance bottleneck.

Depending on your workload and how CPU-intensive it is, the overhead of using a multi-byte character set like UTF8 can be significant compared to using other character sets.

For example, in internal tests conducted by Oracle, Latin1 performed much better than UTF8. However, if you need to use UTF8, then consider upgrading to MySQL version 8.0 and above. MySQL version 8.0 and above has some advantages working with UTF8 over older versions, such as MySQL 5.7 or older.

> **Best practice for:**
>
> Database Administrators

Configure Data and Transaction Logging on Latency LUNs

In Oracle MySQL Cloud Service, data files and transaction logs are stored on latency volumes by default. Latency LUNs provide better scalability and performance-gain over the throughput LUNs when running IO-intensive workloads. Oracle recommends that you use latency volumes for data and transaction logging if you have manually set up MySQL database on a compute VM.

> **Best practice for:**
>
> Database Administrators

Increase Flush Thread Count to Accelerate the Flushing of Dirty Pages

For heavy-write scenario, consider increasing the flush thread count to accelerate the flushing of the dirty pages by changing the value of the `innodb_page_cleaners` system variable. Using this variable, you can specify the number of dirty pages to be flushed from the buffer pool instances.

For example, you can specify `innodb_page_cleaners=8`.

> **Note:**
>
> Number of cleaner threads should be in the following ratio with the BP instances: 1:1, 1:2, or 1:3.

> **Best practice for:**
>
> Database Administrators

Tune the `innodb-io-capacity` and `innodb-io-capacity-max` Parameters for Maximum IO

If you system is running a large write-workload, then tune the `innodb-io-capacity` and `innodb-io-capacity-max` parameters to maximize IO.

Set `innodb-io-capacity=4000` and `innodb-io-capacity-max=10000`. If you are unsure about the maximum IO capacity, then you can set `innodb-io-capacity-max=innodb-io-capacity * 2` (twice the value of `innodb-io-capacity`).

Best practice for:

Database Administrators

Use jemalloc Memory Allocator for CPU-Intensive Workloads

If you are running CPU-intensive, highly-concurrent workloads, use jemalloc to optimize the scalability of your system.

Best practice for:

Database Administrators

Disable Binary Logging

If you are running write-heavy workloads, then consider disabling the binary logging if not required.

To disable binary logging, open the `/etc/mysql/my.cnf` file for editing, and comment out all lines starting with the `log-bin` parameter.

Note:

Disabling the binary logging affects the incremental backup operation. If you choose to disable binary logging, then make sure that you have a proper strategy for backup and restore operations.

Best practice for:

Database Administrators

5

Replicating Data Between On-Premises and Oracle Cloud Infrastructure Compute Classic Using Oracle GoldenGate

Oracle GoldenGate is a comprehensive software package for real-time data integration and replication in heterogeneous IT environments. The product set enables high-availability solutions, real-time data integration, transactional change data capture, data replication, transformations, and verification between operational and analytical enterprise systems.

Quick Checklist

Area	Optimizations	Persona (DevOps/Developer/Admin)
Generic Performance	Configuring Oracle GoldenGate for Performance, Manageability, and Database Stability	DevOps Engineers
Feature-specific	Best Practices for Replication between Cloud and On-Premises Environments	DevOps Engineers

Configuring Oracle GoldenGate for Performance, Manageability, and Database Stability

Configure Oracle GoldenGate for the best performance, simple manageability, and stability for Oracle databases.

See **Oracle GoldenGate Performance Best Practices**.

> **Best practices for:**
>
> DevOps Engineers

Best Practices for Replication between Cloud and On-Premises Environments

Follow best practices for replication with Oracle GoldenGate.

See Doc ID 1996653.1.

6

Running Analytics in Oracle Analytics Cloud

Oracle Analytics Cloud comprise of Oracle BI Cloud Service and Oracle Essbase Cloud. Some performance best practices should be followed for running analytics in Oracle Analytics Cloud.

Quick Checklist

Area	Optimizations	Persona (DevOps/Developer/Admin)
Generic Performance	Performance Impact on Oracle Analytics Cloud Service Queries from Network Latency Validating Datasources and initblock During RPD Lift-and-Shift Maintaining Minimum Number of Executors for Spark Dynamic Allocation Tuning Spark Configuration to Work with Analytics Optimizing the Buffer Size and Executor Memory for Spark SQL Queries Optimizing Resource Usage by Disabling Spark History Best Practices for Downloading Large BI Reports	Developers and BI Administrators Developers and BI Administrators BI Administrators Big Data Cloud Service Administrators and BI Administrators Big Data Cloud Service Administrators and BI Administrators Big Data Cloud Service Administrators and BI Administrators Developers and BI Administrators
IO	• Latency Volumes for Oracle Analytics Cloud – Essbase Data Files • High-Memory Shapes for Oracle Analytics Cloud – Essbase Service Instances	Developers and Essbase Administrators
Feature-specific	HTTP Layer Cache Control Pragmas for Static Files	Developers

Best Practices for Downloading Large BI Reports

Downloading or exporting large reports in .xls or .pdf format is a resource-intensive process. Therefore, Oracle recommends that you use the CSV format to export or download the BI reports that span over 25k (twenty-five thousand) rows, and avoid multiple concurrent downloads.

The default number of rows that you can display in views or that you can download is specified by the following parameters in the instanceconfig.xml file:

Element	Description	Default Value
DefaultRowsDisplayedInDownload	Specifies the default number of rows that can be included in the view when it is downloaded, such as to a PDF file.	65000 (6500 for Simple Trellis; 10000 for Advanced Trellis)
DefaultRowsDisplayedInDownloadCSV	Specifies the default number of rows that can be included in the view when it is downloaded to a CSV file.	65000

The instanceconfig.xml file is located at: *BI_DOMAIN*`/config/fmwconfig/biconfig/OBIPS`.

Performance Impact on Oracle Analytics Cloud Service Queries from Network Latency

If your Oracle Analytics Cloud service instance points to a database that is on a different data center and has longer latency, then increase the result set array size inside the repository to 1000 (**DEFAULT_BULK_FETCH_ROW_COUNT**), to reduce the impact of long network latency, and ensure **MAX_BULK_FETCH_BUFFER_SIZE** is large enough to accommodate the rows.

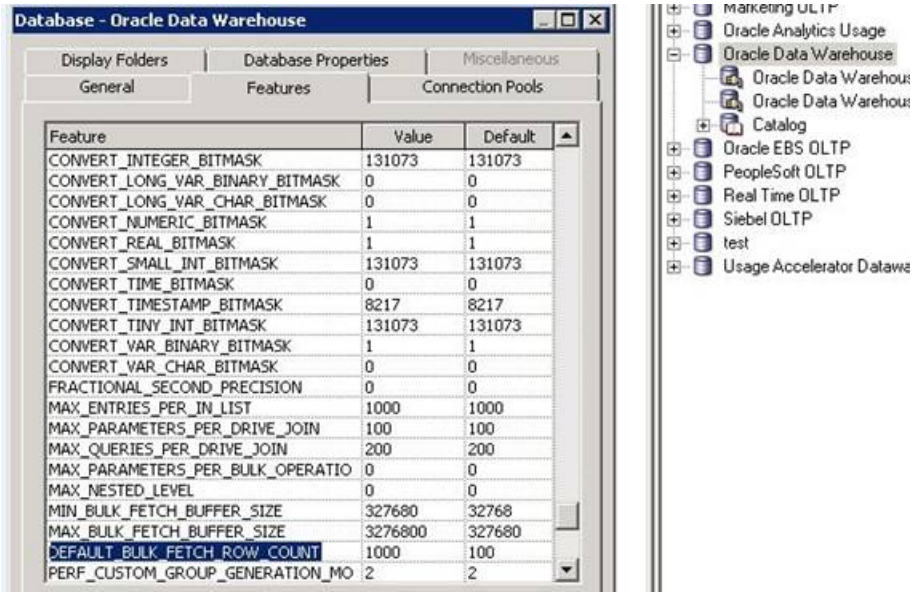

Your Oracle Analytics Cloud service instance pointing to a database on a different data center with longer latency can have a performance impact on the queries. To minimize the performance impact, you can reduce the network round-trips between the Oracle BI server and the data source, which reduces the response time of the queries.

Increase the bulk fetch row count and buffer size to reduce negative effect from network latency in retrieving bulk data. These settings are in the BI RPD. The number of rows in a bulk fetch is limited by the bulk fetcher buffer size. The parameter **MAX_BULK_FETCH_BUFFER_SIZE** must be increased to accommodate rows in a bulk fetch.

If the buffer size is not high, then each bulk will have fewer number of rows than expected by the **DEFAULT_BULK_FETCH_ROW_COUNT**. This leads to degradation when the network latency is high.

> **Best practice for:**
>
> BI RPD Developers and BI Administrators

Latency Volumes for Oracle Analytics Cloud – Essbase Data Files

To improve IO latency and performance for Oracle Analytics Cloud – Essbase queries, use latency-optimized volumes for your data files.

Oracle Cloud Infrastructure Compute Classic provides two types of block storage options: one storage is optimized for throughput, and the other one is optimized for latency. Because Oracle Analytics Cloud – Essbase is very IO-sensitive, Oracle recommends that you use the latency-optimized option.

If you have manually installed Essbase on a compute VM, then make sure that you use latency volumes for Essbase cubes data, or SSD volumes if your account is entitled to it.

> **Best practice for:**
>
> Developers and Essbase Administrators

High-Memory Shapes for Oracle Analytics Cloud – Essbase Service Instances

Oracle Analytics Cloud – Essbase would greatly benefit from a high-memory shape in its Oracle Cloud Infrastructure Compute Classic instance, so use high-memory shape Oracle Cloud Infrastructure Compute Classic instances if possible.

Oracle Analytics Cloud – Essbase uses buffered IO and not direct IO, so a bigger OS file buffer cache can improve IO latency and performance for Oracle Analytics Cloud – Essbase queries. A high-memory shape for an Oracle Cloud Infrastructure Compute Classic instance allows more memory, hence more room for the buffer cache.

> **Best practice for:**
>
> Developers and Essbase Administrators

HTTP Layer Cache Control Pragmas for Static Files

If you manage an instance of Oracle Analytics Cloud service, ensure that you have cache control pragmas set for your static files to be cached on the browser. Caching of static files on the client browser help in reducing the network latency to render the user interface.

Typically the images, CSS files, and JavaScript files are considered to be static files.

These cache control pragmas are not required in a managed PaaS environment.

> **Best practice for:**
>
> Developers

Validating Datasources and initblock During RPD Lift-and-Shift

During the lift-and-shift of Oracle Business Intelligence (BI) repository (RPD), some datasources specified in an on-premises RPD may not be reachable or may be irrelevant in an Oracle Cloud Infrastructure Compute Classic setup, and this can slow down performance significantly.

During an RPD lift-and-shift, validate datasource and corresponding initblocks.

For example, in `nqserver.log`, look for the following error:

```
"[13011] Query for Initialization Block 'Loglevel' has failed. [["
```

> **Best practice for:**
>
> Developers and BI Administrators

Maintaining Minimum Number of Executors for Spark Dynamic Allocation

When using Oracle Analytics Cloud service with Oracle Big Data Cloud Service, by default Oracle Big Data Cloud Service allocates 0 executors.

The executors are created on demand that adds more than 5 seconds to the time that the first SQL statement takes to execute when running SparkSQL with no executor allocated. So, maintaining minimum number of executors helps the SQL runtime.

Define the initial or minimum number of executors for spark dynamic allocation in the **Advanced spark-thrift-sparkconf** section in Ambari UI.

```
spark.dynamicAllocation.initialExecutors = <# of nodes in the Spark cluster>
spark.dynamicAllocation.minExecutors = <# of nodes in the Spark cluster>
```

Best practice for:

BI Administrators

Tuning Spark Configuration to Work with Analytics

When you use Oracle Business Intelligence Cloud Service with Oracle Big Data Cloud Service, the Oracle Business Intelligence Cloud Service datasets are typically not as large as other Spark scenarios.

Reducing the shuffle partitions helps the performance of the Spark queries.

For configuring Spark in Ambari UI, specify the shuffle partitions in the **Custom spark-thrift-sparkconf** section.

```
spark.sql.shuffle.partitions = 4
```

Best practice for:

Big Data Cloud Service Administrators and BI Administrators

Optimizing the Buffer Size and Executor Memory for Spark SQL Queries

Spark SQL queries may fail because of insufficient buffer size and executor memory, insufficient capacity of Spark Kryo serializer, and short timeout setting for long queries.

To optimize the buffer size and executor memory for Spark SQL:

Set the following in the *Custom spark-thrift-sparkconf* section of Spark configuration in the Ambari UI:

```
spark.kryoserializer.buffer=128m
spark.kryoserializer.buffer.max=2047m
spark.network.timeout=600
```

Configure a proper Spark executor JVM heap size to ensure that sufficient memory is being allocated. You can increase this if the executor nodes have highly available memory.

```
spark.executor.memory=8g
```

Best practice for:

Big Data Cloud Service Administrators and BI Administrators

Optimizing Resource Usage by Disabling Spark History

If you don't need to maintain Spark history data through event logging, consider disabling it to save disk and CPU usage.

For example, when Spark SQL request rate is high, Spark history data can pile up in the disk, and Spark history can occupy at least one core of the CPU to process the data. Disable the event logging to save the disk space and to optimize the processing power of the CPU.

To disable event logging, set the following in the *Advanced spark-defaults* and *Advanced spark-thrift-sparkconf* sections of Spark configuration in the Ambari UI:

```
spark.eventLog.enabled=false
```

Best practice for:

Big Data Cloud Service Administrators and BI Administrators

Optimize Essbase Dataload In Hybrid Deployment

This explains optimizing Essbase Dataload in hybrid deployment.

Oracle recommends that you use **CLI Push** method in a hybrid deployment between OAC/Essbase in cloud and a remote database on on-premise. CLI compresses data before transferring it into Essbase. This improves the network transfer speeds in deployments with high network latency.

Best practice for:

Administrators

7

Running Oracle Traffic Director on Oracle Cloud Infrastructure Compute Classic

Oracle Traffic Director (OTD) is an on-board, highly available Application Delivery Controller (ADC) to optimize application-to-application communication. Several performance best practices should be followed for OTD on Oracle Cloud Infrastructure Compute Classic .

Quick Checklist

Area	Optimizations	Persona (DevOps/Developer/Admin)
Generic Performance	Oracle Traffic Director Tuning for Performance	DevOps Engineers
High Availability	Placement of Oracle Cloud Infrastructure Compute Classic Instances on Different Nodes for High Availability	DevOps Engineers
Feature-specific	• HTTP Keep-Alive Timeout Parameter • Thread Pool Settings • Always Use Keep Alive Setting for Connections to Origin Servers • Static Files Cached at Oracle Traffic Director • Cache Headers for Static Files in the Web Tier • Compression of Text for Faster Request and Response Times	DevOps Engineers

Oracle Traffic Director Tuning for Performance

For best performance, do the generic Oracle Traffic Director tuning.

See Tuning Oracle Traffic Director for Performance in *Oracle Traffic Director Administrator's Guide*.

> **Best practice for:**
>
> DevOps Engineers

Placement of Oracle Cloud Infrastructure Compute Classic Instances on Different Nodes for High Availability

Placement of Oracle Traffic Director Infrastructure Compute Classic Cloud Service instances on different physical nodes using the orchestration flags ensures high availability during node failures.

Such placement provides true HA capabilities.

While provisioning Oracle Traffic Director node Oracle Cloud Infrastructure Compute Classic Service instances, make sure that the type attribute under relationships in a launch plan is used to increase high availability.

> **Best practice for:**
>
> DevOps Engineers

HTTP Keep-Alive Timeout Parameter

Make sure HTTP keep-alive is enabled to remove the cost of establishing connections in the high-latency network path.

In the Administration Console, select **Advanced Settings > HTTP > KeepAlive** and make sure HTTP keep-alive is enabled. Set the **KeepAlive** time as follows:

- 61 seconds if OTD supports a web-based UI that is not opening a large number of connections per user
- 5 seconds for most REST invocations and most JET-based UIs

```
<keep-alive>
    <timeout>5</timeout>
</keep-alive>
```

Establishing a connection is an expensive process that involves round trips between the client and server. The `keep-alive timeout` parameter sets the client-server connections to be preserved for the traffic to flow without connection cost. Too many keep-alive connections with high `timeout` values would let new users beyond threshold to starve for resources. You need to tune this for the workload requirements.

> **Best practice for:**
>
> DevOps Engineers

Thread Pool Settings

Tune the thread pool to allow the maximum number of threads that can run in parallel to avoid exhaustion of resources.

Queues are used to avoid errors during burst requests (that is, spikes).The queue size is defined to let the requests wait until a free thread is available to serve. This is done to avoid errors when all the threads are busy, avoid exhaustion of resources (throttling), and make pre-initialized threads available to serve requests.

In the Administration Console, select **Advanced Settings > Thread Pool** and tune as follows:

- Enable the thread pool.
- Change `max-threads` to `20480` (make sure that the OS File Descriptor number is tuned).
- Change `queue-size` to `3000`.

In the instance `server.xml` file, this looks like:

```
<thread-pool>
    <queue-size>3000</queue-size>
    <max-threads>20480</max-threads>
</thread-pool>
```

> **Best practice for:**
>
> DevOps Engineers

Always Use Keep Alive Setting for Connections to Origin Servers

The `keep_alive` setting reduces response times and improves throughput by removing the cost of establishing a new connection for every request.

Connection establishment and SSH handshakes are costly operations. Maintaining a pool of connections between OTD and origin servers reduces the overall cost of serving client requests. Connections can be reused across requests from the same and different clients.

When the origin server is WebLogic Server, the connection between OTD and WLS will always be persisted and kept alive.

1. In the Administration Console, select **Virtual Servers > Config > Routes > Default Route**.
2. Set **Always Use Keep Alive** to **true**.

> **Best practice for:**
>
> DevOps Engineers

Static Files Cached at Oracle Traffic Director

Use static file caching to reduce the network and server processing cost of transferring static artifacts from the origin server to Oracle Traffic Director (OTD).

Static file caching allows OTD to cache static artifacts from origin servers on the web server to avoid a round trip to the origin server for every such request. Because the caching happens in-memory, the sizing has to be appropriate to avoid exhaustion of resources for other functions.

You can use OTD to cache static artifacts. Adjust the cache size to be able to cache the objects you want:

1. In the Administration Console, select **Advanced Settings > Cache**.
2. Make sure the cache is enabled and set the following values:
   - **Heap Size** to `104857600`
   - **Heap Object** size for caching to `4194304`
   - **Replacement** algorithm LRU
   - **Max Entries** to `4096`

> **Best practice for:**
>
> DevOps Engineers

Cache Headers for Static Files in the Web Tier

Static file headers improve processing time and reduce network congestion between the web server and origin servers.

Caching static artifacts from the origin servers in the web tier helps with improved request processing. Configuring the right artifacts and the expiry of those artifacts is key.

You can enable caching and add cache headers like this:

```
<If $uri =~ '.(js|css|woff)$'
    ObjectType fn="proxy-cache-config" id="cachecompressjaandcss" enabled="true"
cache-https-response="true" compression="true"
    ObjectType fn="set-cache-control" control="public,max-age=31536000"
</If>
<If $uri =~ '.(?:gif|jpe?g|png|bmp|ico)$'
    ObjectType fn="proxy-cache-config" id="cacheimages" enabled="true" cache-https-
response="true"
    ObjectType fn="set-cache-control" control="public,max-age=31536000"
</If>
```

> **Best practice for:**
>
> DevOps Engineers

Compression of Text for Faster Request and Response Times

You can use compression to reduce the size of text contents, resulting in reduced network transfer time for requests and responses.

Typically, the network path between the client and the web server is the slowest in the topology because it represents clients residing in the public Internet or WAN. You need to reduce the number of round trips between these two components. Compression works well for text content and results in reduction of content size. This shouldn't be configured for images and binary files, which are already compressed, or would not benefit from compression. Because compression uses CPU cycles, it is a trade-off between CPU versus network resources.

OTD should do response compression (when the client indicates that it can receive a compressed response).

You can enable compression of resources (except for images, PDf files, and SWF) like this:

```
<If $uri !~ '\.(?:gif|jpe?g|png|bmp|ico)$' and $uri !~
'\.(pdf)$' and $uri !~ '\.swf$'
      Output fn="insert-filter" id="response-compression-rule"
filter="http-compression"
</If>
```

> **Note:**
>
> If the origin server adds the cache control header and compresses the contents, then the aforementioned change is not needed in OTD.

> **Best practice for:**
>
> DevOps Engineers

8

Running Node.JS Applications on Application Container Cloud Service

Some best practices apply for Node.js applications on Application Container Cloud Service in Oracle Cloud Infrastructure Classic.

Quick Checklist

Area	Optimizations	Persona (DevOps/Developer/Admin)
Generic Performance	Single-Threaded Node for Performance	Developers
High Availability	• Applications of Node.js on Application Container Cloud Service for High Availability • Application Container Cloud Service AntiAffinity Setting for High Availability • Application Container Cloud Service Rolling Restart Mode for High Availability during Upgrade	Developers

Single-Threaded Node for Performance

For best performance, Node.js applications should never make blocking, sync calls.

For example, never do sync file system calls like **readFileSync**. Instead, do **readFile**.

> **Best practice for:**
>
> Developers

Applications of Node.js on Application Container Cloud Service for High Availability

To support high availability for the Node.js part of a service, run applications of Node.js on the Application Container Cloud Service.

See **Create Sample Node.js Applications** in the Oracle Application Container Cloud Service documentation.

> **Best practice for:**
>
> Developers

Application Container Cloud Service AntiAffinity Setting for High Availability

To ensure high availability against node failure, use the `requiresAntiAffinity` setting provided by ACCS.

The following example is from `manifest.json`:

```
{
  "runtime": {
    "majorVersion": "6.3"
  },
  "command": "node server.js",
  "startupTime": 30,
  "shutdownTime": 30,
  "isClustered": false,
  "requiresAntiAffinity": true,
  "release": {},
  "notes": "NodeJS Sample Application"
  "mode": "rolling"
}
```

> **Best practice for:**
>
> Developers

Application Container Cloud Service Rolling Restart Mode for High Availability during Upgrade

To support high availability during upgrade, use the rolling restart mode provided by ACCS.

Make sure instances are not stuck in start or stop. Use appropriate start and stop times in `manifest.json`.

Example from `manifest.json`:

```
{
  "runtime": {
    "majorVersion": "6.3"
  },
  "command": "node server.js",
  "startupTime": 30,
  "shutdownTime": 30,
  "isClustered": false,
  "requiresAntiAffinity": true,
  "release": {},
  "notes": "NodeJS Sample Application"
```

```
  "mode": "rolling"
}
```

Best practice for:

Developers