



Oracle Database 19c and Oracle RAC 19c on IBM Z and IBM LinuxONE

Tips and Considerations



Srujan Jagarlamudi, Mike Morgan and David Simpson
IBM Oracle International Competency Center
June 2020



Table of contents

Abstract	1
Introduction	1
Terminology.....	1
Linux on Z and LinuxONE hardware.....	1
IBM Z servers.....	2
IBM zEnterprise® EC12.....	2
IBM z13®	2
IBM z14®	2
IBM z15®	3
IBM LinuxONE servers	3
LinuxONE Emperor.....	3
LinuxONE Emperor II.....	3
LinuxONE Rockhopper II	4
LinuxONE III Model LT1.....	4
LinuxONE III Model LT2.....	4
Hardware support	5
Linux configurations on Z and LinuxONE.....	5
Run in an LPAR or as a z/VM guest.....	6
LPAR configuration	6
IBM z/VM configuration	6
Overview.....	6
Network.....	6
Disk storage	7
Memory allocation for Linux guests	7
Oracle Database 19c (Non RAC)	8
Current certifications.....	9
Oracle Database latest patch set.....	9
Release Update (RU) and Release Update Revisions (RUR)	9
Virtualization for Oracle Database 19c (Non RAC).....	10
Oracle Database 19c with Oracle RAC	11
Current certifications.....	11
Oracle Redundant Interconnect usage feature.....	12
Oracle Clusterware	13
Virtualization for Oracle Database 19c with RAC.....	13
Network interconnects	13
Important Oracle Database fixes or issues specific to Linux on Z releases	14



<i>Tuning tips</i>	14
Automatic Storage Management	14
Managing raw disks on Z and LinuxONE to use with ASM	15
Oracle Sharding	15
Multitenant architecture	15
Networking in Oracle RAC clusters	16
Shared storage configuration in Oracle RAC clusters	16
Oracle RAC IPC	17
IBM Spectrum Scale (formerly General Parallel File System)	17
General and tuning resources	17
<i>Summary</i>	19
<i>Resources</i>	20
<i>About the authors</i>	21
<i>Trademarks and special notices</i>	22



Abstract

This paper consolidates the information necessary for planning and implementing Oracle Database 19c (19.x.x.x) single instance database or Oracle Database 19c Real Application Clusters (RAC) on Linux® on IBM® Z and LinuxONE servers. The material presented can also be applied to the IBM Z® platform when running Linux operating systems.

This paper summarizes the information available at the time of publication. It will be updated as changes (e.g. new certifications) occur. It is meant to be used only as a guide. For official Oracle product certification information, the reader is encouraged to consult Oracle's "My Oracle Support" website at: <https://support.oracle.com/>

This paper is written to a level of detail that assumes readers have an in-depth knowledge of Linux, Linux on LinuxONE or IBM Z servers, Oracle Database 19c, Oracle Database 19c RAC and other related products.

Introduction

There are many technical topics to consider when deploying Oracle Database 19c and Oracle Database 19c RAC running Linux on IBM LinuxONE or Z servers such as: status of Oracle certifications, 19c patch sets, Linux code levels, tools, tuning and related software components. The documentation for these topics is spread across many websites, documents, presentations and forums. This paper consolidates that information for easy reference.

This paper focuses on Red Hat® Enterprise Linux Release 7 (RHEL 7) and SUSE Linux Enterprise Server Release 12 (SLES 12), since these are the Linux versions supported for Oracle Database 19c (on-premises) and Oracle RAC 19c running on IBM Z and LinuxONE servers at the time of this publication.

Note: Some of the "My Oracle Support" notes and web links referenced in this document for Oracle Database 19c on Linux on Z are still being updated to reflect the up-to-date information.

For question or feedback, please send a note to the IBM Oracle International Competency Center at ibmoracle@us.ibm.com.

Terminology

Linux on Z is the collective term for the Linux operating system compiled to run on IBM LinuxONE and Z servers.

In 2015, IBM launched the LinuxONE product line for customers who want to run Linux workloads only.

LinuxONE and Z servers run the identical Linux operating systems and application code as they share the same architecture and I/O systems.

Linux on Z and LinuxONE hardware

"It's pretty startling. 85% of all credit cards, 29 billion ATM transactions per year, 92 of the top 100 banks ... 12 billion passenger flights are booked using IBM Z each year, and if I trust Google and their count of the number of searches they do, we do 5x more transactions on mainframes -- our customers do -- than Google does searches each year," – Philip MacLochlainn, Business Unit Executive, IBM Systems HW Sales.



The Linux operating system benefits from the mainframe's capabilities and strengths:

- Outstanding data security
- Availability and performance
- Efficiency and savings

IBM Z servers

Recent IBM Z servers are described below. The IBM Z product line is released in generations. For example, a z15™ is the 14th and current generation of the Z family of servers. Within generations there are many different model types and features. Each generation of the of IBM Z and LinuxONE servers use the same chip architecture. The current models of these product lines both use the “z15” chip and provide binary compatibility.

IBM zEnterprise® EC12

- Launched in 2012
- Up to 120 5.5 GHz CPU cores on 4 internal ‘books’
- Each core can be configured as either CP, IFL, zAAP, zIIP, SAP or ICFs
- Up to 3 TB memory
- Many connectivity options
- These documents provide more detail on the zEC12:
 - IBM zEnterprise EC12 (zEC12) Data Sheet
<https://www.ibm.com/downloads/cas/9ED5WPDQ>
 - IBM System z Introduction July 2014 IBM zEnterprise EC12 (zEC12) and IBM zEnterprise BC12 (zBC12) Frequently Asked Questions
<https://www.ibm.com/downloads/cas/YWR0ZNAK>

IBM z13®

- Launched in 2015
- Up to 141 5.0 GHz CPU cores on 4 internal ‘books’
- Each core can be configured as either CP, IFL, zAAP, zIIP, SAP or ICFs
- Up to 10 TB memory
- Many connectivity options
- These documents provide more detail on the z13:
 - IBM z13 (z13) Data Sheet
<https://www.ibm.com/downloads/cas/K427ZD57>
 - IBM z13 Technical Guide
<https://www.redbooks.ibm.com/redbooks/pdfs/sg248251.pdf>

IBM z14®

- Launched in 2017
- Compact footprint, 40% less physical space than z13
- Up to 170 5.2 GHz CPU cores on 4 internal ‘books’
- Each core can be configured as either CP, IFL, zAAP, zIIP, SAP or ICFs
- Up to 32 TB memory
- Many connectivity options
- These documents provide more detail on the z14:



- IBM z14(z14) Data Sheet
<https://www.ibm.com/downloads/cas/O4VDMBV2>
- IBM z14 Technical Introduction
<http://www.redbooks.ibm.com/redbooks/pdfs/sg248450.pdf>

IBM z15®

- Launched in 2019
- Up to 190 5.2 GHz CPU cores on 4 internal 'books'
- Each core can be configured as either CP, IFL, zAAP, zIIP, SAP or ICFs
- Up to 40 TB memory
- Cloud Native App Development
- Many connectivity options
- These documents provide more detail on the z15:
 - IBM z15(z15) Data Sheet
<https://www.ibm.com/downloads/cas/NN7GBPJ1>
 - IBM z15 Technical Introduction
 - <https://www.redbooks.ibm.com/abstracts/sg248850.html?Open>

IBM LinuxONE servers

Recent IBM LinuxONE servers are described below. The IBM LinuxONE product line is released in generations. For example, a LinuxONE III server is the 3rd generation of the LinuxONE family of servers. Within generations there are many different model types and features. IBM LinuxONE servers are released as a series. Their core technology is based on Z servers, with added functionality and features that provide enhanced support of Linux capabilities. As with Z servers, LinuxONE series servers have many various models and features that can be requested.

LinuxONE Emperor

- Launched in 2015
- Up to 141 5.0 GHz CPU cores (based on z13)
- All cores support CPACF
- Up to 10 TB memory
- Can configure up to 85 LPARs
- Many connectivity options

LinuxONE Rockhopper

- Launched in 2016
- Up to 20 4.3 GHz CPU cores (based on z13)
- All cores support CPACF
- Up to 4 TB memory
- Can configure up to 40 LPARs
- Many connectivity options

LinuxONE Emperor II

- Launched in 2017
- Up to 170 5.2 GHz CPU cores (based on z14)



- All cores support CPACF
- Up to 32 TB memory
- Can configure up to 85 LPARs
- Many connectivity options
- These documents provide more detail on the Emperor II:
 - IBM LinuxONE Emperor II Data Sheet
<https://www.ibm.com/downloads/cas/LVZQYZNB>
 - IBM LinuxONE Emperor II – Specifications
<https://www.ibm.com/products/linuxone-emperor-ii/specifications>

LinuxONE Rockhopper II

- Launched in 2018
- Up to 30 5.2 GHz CPU cores (based on z14)
- All cores support CPACF
- Up to 8 TB memory
- Can configure up to 40 LPARs
- Many connectivity options
- These documents provide more detail on the Rockhopper:
 - IBM LinuxONE Rockhopper II - Announcement Letter
https://www-01.ibm.com/common/ssi/rep_ca/9/897/ENUS118-019/ENUS118-019.PDF

LinuxONE III Model LT1

- Launched in 2019
- Up to 190 5.2 GHz CPU cores (based on z15 architecture)
- All cores support CPACF
- Up to 40 TB memory
- Can configure up to 85 LPARs
- Many connectivity options
- These documents provide more detail on the LinuxONE III:
 - IBM LinuxONE III Mode LT1 – Announcement Letter
https://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/2/897/ENUS119-012/index.html&request_locale=ja
 - Introducing the IBM LinuxONE III
<https://www.ibm.com/downloads/cas/US-ENUS119-012-CA/name/ENUS119-012.PDF>
- These LinuxONE documents are also useful:
 - Practical Migration from x86 to LinuxONE
<http://www.redbooks.ibm.com/abstracts/sg248377.html?Open>
 - Oracle on LinuxONE
<http://www.redbooks.ibm.com/abstracts/sg248384.html?Open>

LinuxONE III Model LT2

- Launched in 2020
- Up to 65 4.5 GHz CPU cores (based on z15 architecture)
- All cores support CPACF
- Up to 16 TB memory
- Can configure up to 40 LPARs



- Many connectivity options
- These documents provide more detail on the LinuxONE III:
 - IBM LinuxONE III Mode LT2 – Announcement Letter
<https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS120-014&appname=USN>
 - Introducing the IBM LinuxONE III
<https://www.ibm.com/products/linuxone-iii>
- These LinuxONE documents are also useful:
 - Practical Migration from x86 to LinuxONE
<http://www.redbooks.ibm.com/abstracts/sg248377.html?Open>
 - Oracle on LinuxONE
<http://www.redbooks.ibm.com/abstracts/sg248384.html?Open>

The IBM Z and LinuxONE architecture is exceptionally good for deploying Oracle data-serving workloads. This architecture can support up to thousands of virtual Linux servers on a single footprint. That means the virtualization capabilities in a single Z or LinuxONE system can result in a less complex Linux infrastructure with fewer components, less management, less space requirements, and optimized core workload efficiency.

IBM will also provide sizing estimations to help predict the system resources necessary to support a given workload. To start the sizing process, visit this web page:

<https://www-935.ibm.com/services/oracle/sizing/>

Additional documentation resources for Linux on Z and LinuxONE servers can be found at:

- Linux on IBM Z:
<https://www.ibm.com/it-infrastructure/z/os/linux>
- Linux on IBM LinuxONE:
<https://www.ibm.com/it-infrastructure/linuxone>
- IBM Developer Works for Linux on Z and LinuxONE:
<https://www.ibm.com/developerworks/linux/linux390/index.html>

Hardware support

The latest fixes and updates for your system's hardware can be found at the **Fix Central** web page at: <https://www-933.ibm.com/support/fixcentral/>.

In general, IBM recommends that you keep your operating system updated and SP levels up to date for your Linux installation. Monitor My Oracle Support notes and your Linux distributor's support pages for the latest issues. Always perform thorough testing on an OS update before deploying to production.

Linux configurations on Z and LinuxONE

The following section provides a high-level view of using LPARs (logical partitions) and z/VM®, for more details please see:

- Getting Started with Linux on IBM Z
<https://www.vm.ibm.com/library/710pdfs/71628700.pdf>



- CP Planning and Administration Guide

<https://www.vm.ibm.com/library/710pdfs/71627102.pdf>

Run in an LPAR or as a z/VM guest

A Linux on Z guest operating system image can be configured as a guest virtualized under a z/VM LPAR or it can run directly in the LPAR.

With z/VM 7.1 the user can create a Linux guest having up to 2TB of memory and up to 64 virtual processors. If more capacity is required, then running the Linux guest in an LPAR can be configured. You can also convert a Linux guest from running under z/VM to native LPAR mode, provided the network is converted to a configuration that does not use virtual switches.

z/VM provides excellent monitoring and diagnostics information for Linux guests and provides the ability to overcommit (share) processors, memory and the network configuration easily, and is preferred for small to midsize Linux guests. The number of guests created must be balanced to the workload, and available physical resources.

LPAR configuration

A Linux or z/VM LPAR can run with simultaneous multithreading (SMT) in SMT-1 mode (1 thread per core) or in SMT-2 mode (2 threads per core). Both LPARs and z/VM support SMT. An LPAR for z/VM or a Linux native guest is configured with a given number of CPU cores (Integrated Facility for Linux or IFLs) and memory from the Z or LinuxONE machine.

You should always ensure that you define the minimum number of IFL cores to be at least as large as the maximum number of virtual cores that are defined to any single Linux guest. For example, if the z/VM LPAR has two Linux guests each defined with four virtual processors, the LPAR should be defined with four IFLs assigned to the z/VM LPAR. If SMT-2 is being used, then at least two IFLs (four threads total) should be assigned to the LPAR as each virtual processor assigned to the Linux guest will use one thread.

IBM z/VM configuration

Overview

To configure a new Linux guest running under z/VM a directory entry for the guest OS is needed. Each Linux guest operating system will have a user directory profile entry that will define the network channels, disk devices/channels, virtual memory and virtual CPUs that the Linux guest operating system can utilize.

Network

The network configuration can use several alternatives. The Oracle Database requires a public network interface for both single instance and Oracle RAC environments. Oracle RAC environments also require a private network interface to communicate between Linux systems.

For the public network interface, a dedicated or shared OSA (Open System Adapter) can be utilized. This can be configured with z/VM virtual switches, used directly, or multiple network OSA ports can



be configured with Linux bonding. Most z/VM customers leverage virtual switches which supports redundancy and aggregation of multiple network interfaces for bandwidth considerations.

The private network interface used for Oracle RAC environments can also be configured in several ways. A dedicated or shared 10 Gb, 25 Gb or 25 Gb RoCE (IP mode) can be utilized. A Hipersocket memory to memory network can also be used for the private interconnect, provided there is enough CPU capacity to support the network interfaces and all the Linux guests configured to run Oracle RAC on the same Z or LinuxONE system. The private interconnect can also utilize a separate (different than a public) private virtual switch that is only used for Oracle RAC interconnect network traffic.

The private interconnect must be configured in layer 2 mode for Oracle RAC to function properly.

Disk storage

A Linux guest can use either ECKD DASD (mainframe) or FCP SCSI (open) storage for the operating system and database disk storage. If using ECKD DASD storage for database files, it's recommended to configure HyperPAV (parallel access volumes), in order to provide better throughput for disk I/O.

In general, ECKD uses less CPU as z/VM handles the multipathing, while FCP SCSI storage provides higher throughput, but at a higher CPU cost.

Oracle Database single instance can use local storage for the Linux guest, with the database files on a file system (ext4, xfs, IBM Spectrum Scale) or using Oracle Automatic Storage Management (ASM).

Oracle RAC systems require shared storage using either Oracle ASM or an IBM Spectrum Scale shared file system.

Memory allocation for Linux guests

The amount of virtual memory a Linux guest should be allocated can be calculated based on the following:

- The sum of each Oracle Database instance's System Global Area (SGA)
- Plus the sum of the Program Global Areas (PGA)
- Plus the total number of database connections at 5 MB each
- Plus the Linux page table size (Linux large pages recommended)
- Plus the Linux kernel, monitoring processes, Golden Gate memory (if applicable)
- Plus another 10% should then be allocated to avoid Linux out of memory processes that could impact the system.

In general, for z/VM memory the sum of the Linux guest virtual memory overcommit should be no larger than two times the amount of real memory assigned to all of the Linux guests in the z/VM LPAR for development environments and one to one and a half times for production LPARs. z/VM paging rates should be always be monitored for optimal settings.

It is also advisable to setup the Linux swap space needed for the Oracle Database processes to relink executables, and for Oracle Database upgrades. Oracle Database installation guides recommend four to 16 GB of Linux swap depending on the Linux guest memory size. If a Linux system starts using Linux swap for business processes the system will experience degraded

performance. To resolve this the memory of the Linux guest should be increased, or the Oracle Database memory requirements reduced.

Multiple swap devices configured with in-memory Virtual Disk (VDISK) set to a higher priority, as shown in Figure 1, are recommended to help minimize any performance impacts of Linux swapping.

```
# swapon -s
```

Filename	Type	Size	Used	Priority
/dev/dasdo1	partition	131000	0	10
/dev/dasdp1	partition	524216	0	5
/dev/mapper/u603_swap3	partition	6291448	0	1

Figure 1. Typical VDISK swap configuration with priorities

Oracle Database 19c (Non RAC)

This section contains the Oracle Database 19c non-RAC technical information that needs to be considered in a Linux installation. Oracle Database 11g Release 2 introduced the Grid Infrastructure installation for both stand alone and RAC databases. In Oracle Database 19c, Grid Infrastructure is also an option which includes Oracle ASM and Oracle Restart. Oracle Restart, introduced in 11.2, is a feature that provides the ability to monitor, manage, and automatically restart the Oracle Database product components, including a single instance Oracle Database, Oracle Net Listener, database services and Oracle ASM. Oracle Grid infrastructure is needed for the stand-alone database server if Oracle ASM is selected as the storage option for the database files.

In June 2019, Oracle released Oracle Database 19c for Linux on Z. Oracle Database 19c is a long-term support release for Oracle 12.2 family of releases. Beginning with version 12.2, Oracle Database software was released annually with a new version representing the release year. The original version of Oracle Database 12.2.0.2 was released as 18c and 12.2.0.3 was released as 19c. The release 19c is the last release for the 12.2 family.

Refer to the My Oracle Support note **742060.1**: “*Release Schedule of Current Database Releases*” for more information.

Oracle Database 19c new features can be found in a document available at:

<https://docs.oracle.com/en/database/oracle/oracle-database/19/newft/new-features.html#GUID-5490FE65-562B-49DC-9246-661592C630F9>

The following are the basic documents to refer to for an Oracle Database 19c installation on Linux on Z.

Oracle Database 19c:

- Oracle Grid Infrastructure Installation and Upgrade Guide for Linux
<https://docs.oracle.com/en/database/oracle/oracle-database/19/cwlin/grid-infrastructure-installation-and-upgrade-guide-linux.pdf>
- Oracle Database Installation Guide for Linux
<https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/index.html>
- Oracle Database Release Notes - 19c for Linux
<https://docs.oracle.com/en/database/oracle/oracle-database/19/rnrdr/index.html>
- RPM Checker when Installing Oracle Database 19c on IBM Linux on System z (Doc ID 2553465.1)
https://support.oracle.com/knowledge/Oracle%20Database%20Products/2553465_1.html



- Oracle Database Licensing Information User Manual – Oracle Database 19c E94254-06: <https://docs.oracle.com/en/database/oracle/oracle-database/19/dblic/database-licensing-information-user-manual.pdf>

Current certifications

To determine the **current certifications** for Oracle Database 19c with Linux on Z:

- Sign into Oracle's My Oracle Support website (UserID/Password is required) at <https://support.oracle.com/CSP/ui/flash.html> and click on "Certifications" tab, look for "Certification Search" section. In the "Product" field type "Oracle Database", the second field is "Release". Select "19.0.0.0.0" from the list. The next field is "Platform", select either "Linux on IBM Z SLES12" or "Linux on IBM Z Red Hat Enterprise Linux 7" depending on your distribution, then click the "Search" button. The search result will show the certification status in the link "See Certification Details for Notes and Support information".

IBM and Oracle recommend using the following versions of Linux to minimize the number of Linux patches required for Oracle Database 19c.

Operating System Level	Product	Status
SLES 12 SP3 or later Red Hat 7 update 4 or later	Oracle Database 19c	Certified

Table 1: Oracle Database Enterprise Edition 19c Release 1 certifications of Linux on Z as of publication date

Here are some certification details to be aware of:

- Dynamic (through dhcpd servers) GNS IP addresses are not supported. GNS must use static IPs.
- Both ECKD devices and SCSI devices are supported for any type of Oracle Database files.
- Oracle Database can be installed to run in an LPAR or as a virtual z/VM guest.

Oracle Database latest patch set

Oracle Database 19c version 19.0.0.0.0 was released as Enterprise Edition (EE) for Linux on Z. To find the documentation for the latest Oracle Database patch sets, sign into Oracle's My Oracle Support website at <https://support.oracle.com/CSP/ui/flash.html>. In the Quick Find pull down, specify "Document ID" number **161818.1**, "Oracle Server (RDBMS) Releases Support Status Summary". Then look for "19c", this will show support notes for patching and release notes.

Release Update (RU) and Release Update Revisions (RUR)

Starting with Oracle Database 12.2.0.1, the Quarterly Release Updates (RUs) will replace Bundle patches and Quarterly Release Update Revisions (RURs) will replace the PSU's to improve the quality and experience of proactive maintenance. For more information on Oracle Database Release Updates (RU) and Release Update Revisions (RUR) see My Oracle Support note **2285040.1**. The MOS note describes Release Updates and Release Update Revisions for the Oracle Database Proactive Patch Program.

The latest Release Update patch for Oracle Database 19c at the time this paper was written is 30116789 (19.5.0.0.0). To obtain the Release Update from My Oracle Support portal, click "Patches & Updates" tab, click on "Product or Family", enter "Oracle Database" in Product field and it will list out the Oracle



Database edition names. Choose one which you have installed. In the next field, click the drop-down button to select a base version of 19c (Oracle Database 19.0.0.0.0). In the next text field, click the drop-down button to select the Linux operating system, select “IBM:Linux on System z”. Then click the “Search” button. The search output will show you a list of patches and pick a relevant patch for your installed database software. The system patch 30116789 comes with patches for both Grid Infrastructure and the database.

Refer to the Patch’s readme file for detailed information.

Also refer the My Oracle Support note **756671.1** for Oracle recommended patches for Oracle Database versions from 11.2.0.4 and later.

Virtualization for Oracle Database 19c (Non RAC)

As mentioned in the “Current certifications” sub-section (see above) for Oracle Database 19c, IBM z/VM features such as LPARs, virtual NICs, virtual disk and virtual memory are supported.

System z feature	Certified With	Status
Dedicated processor LPAR Shared processor LPAR	Oracle DB 19c	Certified
z/VM guest OS	Oracle DB 19c	Certified

- **Logical Partitions** subdivide a computer's processors, memory, and hardware resources into multiple environments so that each environment can be operated independently with its own operating system and applications.
- **Dedicated processor partitions** are LPARs that use dedicated processors.
- **Dedicated processors** are whole physical processors that are assigned to a single LPAR.
- **The z/VM hypervisor** allows sharing of physical resources between multiple guests running in a logical partition (LPAR) including virtual SCSI and virtual networking. This allows more efficient utilization of physical resources through sharing between LPARs and facilitates server consolidation.

Customers should monitor the latest Oracle Certification information to be aware of the ongoing certification of new features.

For additional, detailed z/VM information, see these documents and websites:

- Supported Virtualization and Partitioning Technologies for Oracle Database and RAC Product Releases:
<https://www.oracle.com/database/technologies/virtualization-matrix.html>
- IBM z/VM virtualization technology webpage:
<https://www.ibm.com/it-infrastructure/z/zvm>
- Making z/VM and Linux Guests production... “Best Practices”:
<http://www.vm.ibm.com/education/lvc/lvclbest.pdf>



Oracle Database 19c with Oracle RAC

This section contains the Oracle Database 19c with Oracle RAC technical information that needs to be considered in a Linux on Z installation. Starting with Oracle Database 19c, Oracle has packaged Oracle Clusterware, Automatic Storage Management and the listener as a single package called “Oracle Grid Infrastructure”.

The following are the basic documents to review for an Oracle Database 19c with RAC installation on Linux on Z.

Oracle Database 19c:

- RPM Checker when Installing Oracle Database 19c on IBM Linux on System z (Doc ID 2553465.1)
https://support.oracle.com/knowledge/Oracle%20Database%20Products/2553465_1.html
- Real Application Clusters Installation Guide E96277-03:
<https://docs.oracle.com/en/database/oracle/oracle-database/19/rilin/real-application-clusters-installation-guide-linux-and-unix.pdf>
- Real Application Clusters Administration and Deployment Guide E95728:
<https://docs.oracle.com/en/database/oracle/oracle-database/19/racad/real-application-clusters-administration-and-deployment-guide.pdf>
- Oracle Database Installation Guide 19c for Linux E96432-07:
<https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/database-installation-guide-linux.pdf>
- Oracle Grid Infrastructure Installation and Upgrade Guide 19c for Linux E96272-08:
<https://docs.oracle.com/en/database/oracle/oracle-database/19/cwlin/index.html>

Current certifications

To determine the **current certifications** for Oracle Database 19c with Linux on Z, sign into Oracle’s My Oracle Support website (UserID and Password required) at <https://support.oracle.com/CSP/ui/flash.html> and click on “Certifications” tab, look for the “Certification Search” section. In the “Product” field, enter “Oracle Real Application Clusters”, the second field is “Release”. Select one of them from the list, 19.0.0.0.0, next field is “Platform”, select either “Linux on IBM Z SLES12” or “Linux on IBM Z Red Hat Enterprise Linux 7” depending on your distribution. Then click the “Search” button. The search result will show the certification status in the link “See Certification Details for Notes and Support information” in the search result page. This lists the information with Certification status for the above input.

The same Linux versions mentioned in the section “Oracle Database 19c (Non-RAC)” are applicable to Oracle Database 19c with RAC.

The following link points to the RAC Technologies Compatibility Matrix. This matrix provides details about supported storage technologies, network interconnect technologies and other platform-specific information.

<https://www.oracle.com/database/technologies/tech-generic-linux-new.html>

For Oracle Database and server virtualization features support, look at the information for “IBM System z” in the following link.



<http://www.oracle.com/technetwork/database/virtualizationmatrix-172995.html>

Oracle Redundant Interconnect usage feature

Oracle Database 19c continues to support Oracle Grid Infrastructure Redundant Interconnect, also known as Cluster High Availability IP (HAIP) feature, which provides a Highly Available (HA) IP network functionality for the Oracle RAC node interconnectivity. Oracle's Redundant Interconnect Usage feature will protect production RAC databases when there are private NIC failures or when the subnet associated with a private NIC fails.

For upgrade customers, it is recommended to maintain their current configuration with 11.2.0.4, 12.1.0.2, 12.2.0.1, or 18c releases during upgrades. This will allow the Redundant Interconnect Usage to allocate an (HAIP) on top of the virtual private device but will not enable load balancing or network failover based on the Oracle Redundant Interconnect Usage feature. Load balancing as well as network failover will continue to be managed by link aggregation in a virtual switch or with bonding. No further configuration steps are required.

In order to fully enable Redundant Interconnect Usage to manage load balancing and network failover for the Oracle cluster interconnect, an additional network interface should be defined which provides an additional private NIC on each cluster node. An additional network interface must be created on a new subnet and then defined to the cluster. This is done with the "oifcfg setif" command. Oracle Grid Infrastructure will then provide a HAIP on new private NIC. Please see the section "Redundant Interconnect Usage" in the "Oracle Clusterware Administration and Deployment Guide":

<https://docs.oracle.com/en/database/oracle/oracle-database/19/cwadd/clusterware-administration-and-deployment-guide.pdf>

If one of the defined cluster interconnect interfaces fails or becomes non-communicative, then Oracle Clusterware transparently moves the corresponding HAIP address to one of the remaining functional private interfaces.

Hosting more than one instance of a production Oracle Database with RAC in the same physical environment or frame with a single point of failure (sharing components required for network connectivity, storage access, common Hypervisor, or other critical components) at the same time is generally not recommended by Oracle for a complete high availability solution, as a failure of any of those shared components inevitably affects more than one instance of the production Oracle Database. Under certain circumstances, virtualization solutions and other techniques provided by the hardware or OS vendor may mitigate these negative effects, however, for critical and production deployments, clustering within the same frame, if it has a single point of failure, is discouraged. Furthermore, at this point in time, when enabling Redundant Interconnect Usage, avoid co-location of Oracle RAC instances belonging to the same production database on the same frame as described above, when configured with virtual Ethernet, as certain failures (e.g. the loss of a physical network) in the frame could lead to losing the majority of the Oracle Database instances. Oracle and IBM are working to integrate the Redundant Interconnect Usage feature so that optimized high availability can be ensured. Alternatively, physical devices (as opposed to virtual or z/VM based devices) can be used and managed by the Redundant Interconnect Usage feature directly to avoid such scenarios.



Oracle Clusterware

In a RAC environment Oracle Clusterware provides the high availability functionality. This includes monitoring or restarting the nodes of the cluster for the database instances, for the listeners and for the database services. Oracle Clusterware is required for Oracle Database 19c with RAC. For more information on Oracle Clusterware, visit:

<http://www.oracle.com/technology/products/database/clustering/index.html>.

Virtualization for Oracle Database 19c with RAC

As mentioned in the “Current certifications” sub-section (see above) for Oracle Database 19c with Oracle RAC, LPARs and Micro-Partitioning are supported in Oracle Database 19 with RAC environments. LPARs and z/VM virtualization are part of IBM z/VM. IBM z/VM is the certified hypervisor available on all Z and LinuxONE servers. The z/VM capabilities supported in Oracle Database 19c are:

System z features	Certified With	Status
Dedicated processor LPAR Shared processor LPAR	Oracle DB 19c RAC	Certified
z/VM	Oracle DB 19c RAC	Certified

- **Logical Partitions** subdivide a computer's processors, memory, and hardware resources into multiple environments so that each environment can be operated independently with its own operating system and applications.
- **Dedicated processor partitions** are LPARs that use dedicated processors.
- **Dedicated processors** are whole physical processors that are assigned to a single LPAR.
- **The z/VM hypervisor** allows sharing of physical resources between multiple guests running in a logical partition (LPAR) including virtual SCSI and virtual networking. This allows more efficient utilization of physical resources through sharing between LPARs and facilitates server consolidation. These z/VM features are specifically supported in Oracle Database 19c RAC:
 - Virtual LAN (IBM virtual switches) for public and private interconnects
 - Virtual disks (ECKD or SCSI) for shared and local disk storage
 - Virtual CPU for cluster nodes
 - Virtual memory for cluster nodes
 - Create section on z/VM

Network interconnects

To determine the **current certifications** for network interconnects on Oracle RAC, browse to the “General Notes for RAC for Linux on System z based Systems (RAC only)” web page, as described in the beginning of the “Current certifications” sub-section (see above) for Oracle Database 19c with Oracle RAC. Then select the link for the RAC Technologies Compatibility Matrix (RTCM). Within RTCM, Network Interconnect certifications are listed. Currently, the following are supported and certified with RHEL 7.4+ and SLES 12 SP1+:

1 and 10 Gigabit Ethernet
System z Hipersockets



For supported Interconnect Network Technologies for Oracle RAC on UNIX, refer the following link:

<https://www.oracle.com/database/technologies/tech-generic-linux-new.html>

Important Oracle Database fixes or issues specific to Linux on Z releases

Database Patch 30116789 – for both GI and RDBMS

Tuning tips

Most of the Linux on Z tuning tips recommended for lower versions of Oracle Database are applicable to Oracle Database 19c. Please note that Oracle's Linux documentation pertains to Linux on Z except where noted.

- VM workshop – Oracle on IBM Z Performance Tips/Update
<http://www.vmworkshop.org/2019/present/oraperft.pdf>
- Performance Tuning Tips for Oracle workloads with Linux on IBM Z
<https://events.share.org/Summer2018/Public/SessionDetails.aspx?FromPage=Sessions.aspx&SessionID=5501&nav=true&Role=U%27>
- Oracle Database 12c Release 2 New Features (OVERVIEW) from zSig meeting
https://www.zseriesoraclesig.org/2017presentations/12_2_Feature_SIG2017_FINAL_june_13.pdf

The tuning items chosen for inclusion in this section are those where the defaults are sub-optimal, or settings that require special consideration. Instead of repeating the existing documentation, this section will consolidate references to the existing documentation. The URLs of the documents referenced in this section are listed at the end of the section.

Automatic Storage Management

Automatic Storage Management (ASM) is a feature in Oracle Database 19c that provides the database administrator with a simple storage management interface that is consistent across all server and storage platforms.

Oracle Clusterware OCR and voting disk files can be stored in an Oracle ASM disk group. In addition to OCR and voting files, ASM can be used for database data and recovery files.

For Oracle ASM block devices (DASD or SCSI) asynchronous I/O is the default Oracle setting. If using a Linux file system (ECKD DASD or SCSI) then it's recommended to utilize Oracle's `filesystemio_options = setall` parameter with ext4 or xfs file systems.

ASM Dynamic Volume Manager (DVM), ASM Intelligent Data Placement, ASM Storage Management Configuration Assistant (ASMCA), ASM File Access Control and ASMCMD are some of the extended functions of ASM.

In Oracle Database 12c, Oracle introduced several new capabilities in ASM. Oracle Flex ASM is one of the major features which protects Oracle Database 12c instances that were relying on a failing ASM instance by reconnecting to another surviving ASM instance on a different server.

In Oracle Database 18c, a new ASM feature “ASM Database Clones” was introduced for Multitenant database.



In Oracle Database 19c, parity protection is available for write-once files such as archive logs and backup sets.

For more information on ASM new features , refer to the Oracle document “A Technical Overview of New Features for Automatic Storage Management in Oracle Database 19c” at:

<https://www.oracle.com/technetwork/database/database-technologies/cloud-storage/asm/overview/asm-new-features-5303288.pdf>

and refer to “Oracle Database 19c New Features Guide” at:

<https://docs.oracle.com/en/database/oracle/oracle-database/19/newft/new-features.html#GUID-5490FE65-562B-49DC-9246-661592C630F9>

Managing raw disks on Z and LinuxONE to use with ASM

These documents, while written for Oracle Database 12c, are applicable for 18c and 19c and can help with configuring ECKD or SCSI devices for use with ASM:

- How to Manually Configure ECKD/DASD devices for use with Oracle ASM on IBM: Linux on System z for versions 12c Database (Doc ID 2407588.1)
<https://support.oracle.com/epmos/faces/DocumentDisplay?id=2407588.1>
- How to Manually Configure multipath fcp/scsi devices for use with Oracle ASM on IBM: Linux on System z with 12c Database (Doc ID 2407589.1)
<https://support.oracle.com/epmos/faces/DocumentDisplay?id=2407589.1>

Oracle Sharding

Sharding is a data tier architecture in which data is horizontally partitioned across independent databases. Sharding divides a database into a farm of independent databases and thereby avoiding the scalability or availability edge cases associated with a single database.

Each database is hosted on a dedicated server with its own local resources - CPU, memory, flash, or disk. Each database in such configuration is called a shard. All the shards together make up a single logical database, which is referred to as a sharded database (SDB).

The documents below provide a good start in implementing and configuring an SDB:

- OOW2019 – Achieving Massive Scalability and Total Fault Isolation through Oracle Sharding
https://static.rainfocus.com/oracle/oow19/sess/1552362603972001eq8Q/PF/Yu_Tao_oracelsharding_OOW19_1568736537357001klJA.pdf
- Oracle Database – Using Oracle Sharding
<https://docs.oracle.com/en/database/oracle/oracle-database/19/shard/index.html>

Multitenant architecture

The multitenant architecture enables an Oracle Database to function as a multitenant container database (CDB). A CDB includes zero, one, or many customer-created pluggable databases (PDBs). A PDB is a portable collection of schemas, schema objects, and non-schema objects that appears to an Oracle Net client as a non-CDB. All Oracle Databases before Oracle Database 12c were non-CDBs.

Please refer to the following guide for implementing CDBs with PDBs:

- Oracle Multitenant Administrator's Guide

<https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/multitenant-administrators-guide.pdf>

Networking in Oracle RAC clusters

Please refer to the Chapter 5 of the Grid Infrastructure Installation and Upgrade guide for network configuration information:

- Oracle Grid Infrastructure - Grid Infrastructure Installation and Upgrade guide
<https://docs.oracle.com/en/database/oracle/oracle-database/19/cwlin/grid-infra-structure-installation-and-upgrade-guide-linux.pdf>

Additionally, we recommend:

- Use Jumbo frames for private NICs (MTU='8992')
- Use a different private NIC interface on each node for vender cluster file systems than those used for Oracle Clusterware and CRS internode communication.
- Create at least two private NICs per cluster node for redundancy.
- Use bonded NICs to create the public interface on each node.

If all your nodes are on the same physical machine, use virtual switches for private NICs. Here is some useful documentation:

- IBM Knowledge Center – Virtual Switch
https://www.ibm.com/support/knowledgecenter/en/SSB27U_6.4.0/com.ibm.zvm.v640.hcpa6/hcpa691.htm
- IBM Systems Technical University – z/VM Virtual Switch (Part 1)
<https://www.vm.ibm.com/devpages/altmarka/vswitch1.pdf>

Shared storage configuration in Oracle RAC clusters

These documents (referenced earlier) can help with configuring ECKD or SCSI devices:

- How to Manually Configure ECKD/DASD devices for use with Oracle ASM on IBM: Linux on System z for versions 12c Database (Doc ID 2407588.1)
<https://support.oracle.com/epmos/faces/DocContentDisplay?id=2407588.1>
- How to Manually Configure multipath fcp/scsi devices for use with Oracle ASM on IBM: Linux on System z with 12c Database (Doc ID 2407589.1)
<https://support.oracle.com/epmos/faces/DocumentDisplay?id=2407589.1>

Additionally, we recommend:

Have sufficient paths from the virtual disk devices to the real devices.

For example, using the Linux 'vmcp' command, for virtual disk '6F4A' (an ECKD device), we can see that it has 4 paths to real device 6F4A:

```
# vmcp q paths to 6f4a
Device 6F4A, Status ONLINE
CHPIDs to Device 6F4A (PIM) 58 59 5A 5B
```



Physically Available (PAM)	+ + + +
Online (LPM) :	+ + + +
Transport Mode supported	+ + + +
Legend	+ Yes - No

For SCSI devices, the following command will show virtual disks and their mapping info:

```
### Check attached fcp devices:
node2:~ # vmcp q fcp | grep "ON FCP"
FCP 0300 ON FCP D207 CHPID A5 SUBCHANNEL = 0000
FCP 0301 ON FCP D307 CHPID A6 SUBCHANNEL = 0001
### Two attached virtual fcp devices..300 and 301
```

```
### Check available SCSI disks with their size:
node2:~ # lsscsi -s
[0:0:0:0] storage IBM 2810XIV-LUN-0 0000 - -
[0:0:0:1] disk IBM 2810XIV 0000 /dev/sda 206GB
[0:0:0:2] disk IBM 2810XIV 0000 /dev/sdb 34.4GB
[1:0:0:0] storage IBM 2810XIV-LUN-0 0000 - -
[1:0:0:1] disk IBM 2810XIV 0000 /dev/sdc 206GB
[1:0:0:2] disk IBM 2810XIV 0000 /dev/sde 34.4GB
```

The first entry on each line (the tuple) is the hostadppter, channel, target_number, LUN. We can see that /dev/sda and /dev/sdc are the same disk. This can be verified with the multipath command:

```
node1:~ # multipath -l
volc (2001738004e2603c5) dm-11 IBM ,2810XIV
size=32G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=0 status=active
  |- 0:0:0:1 sda 8:96 active undef running
  `- 1:0:0:1 sdc 8:112 active undef running
...
volo (2001738004e2603c1) dm-4 IBM ,2810XIV
size=32G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=0 status=active
  |- 0:0:0:2 sdb 66:0 active undef running
  `- 1:0:0:2 sde 65:208 active undef running
```

Oracle RAC IPC

Oracle RAC with Oracle Database 19c uses the User Datagram Protocol (UDP) for inter-process communication (IPC) between nodes. The Grid Infrastructure Information and Upgrade guide has information in Appendix B about setting UDP and TCP kernel parameters.

IBM Spectrum Scale (formerly General Parallel File System)

Useful IBM Spectrum Scale documents can be found at the following link:
https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale_welcome.html.

General and tuning resources

These are the documents referenced throughout this section. The same tuning recommendations for Linux on Z with Oracle Database 12c apply to Oracle Database 19c.



- VM workshop – Oracle on IBM Z Performance Tips/Update
<http://www.vmworkshop.org/2019/present/oraperft.pdf>
- Performance Tuning Tips for Oracle workloads with Linux on IBM Z
<https://events.share.org/Summer2018/Public/SessionDetails.aspx?FromPage=Sessions.aspx&SessionID=5501&nav=true&Role=U%27>
- Oracle Database 12c Release 2 New Features (OVERVIEW) from zSig
https://www.zseriesoraclesig.org/2017presentations/12_2_Feature_SIG2017_FINAL_june_13.pdf
- A Technical Overview of New Features for Automatic Storage Management in Oracle Database 19c
<https://www.oracle.com/technetwork/database/database-technologies/cloud-storage/asm/overview/asm-new-features-5303288.pdf>
- Oracle Database 19c New Features Guide
<https://docs.oracle.com/en/database/oracle/oracle-database/19/newft/new-features.html#GUID-5490FE65-562B-49DC-9246-661592C630F9>
- How to Manually Configure ECKD/DASD devices for use with Oracle ASM on IBM: Linux on System z for versions 12c Database (Doc ID 2407588.1)
<https://support.oracle.com/epmos/faces/DocContentDisplay?id=2407588.1>
- How to Manually Configure multipath fcp/scsi devices for use with Oracle ASM on IBM: Linux on System z with 12c Database (Doc ID 2407589.1)
<https://support.oracle.com/epmos/faces/DocumentDisplay?id=2407589.1>
- IBM Spectrum Scale documents
https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale_welcome.html



Summary

This document gathers together the key technical topics that need to be considered in planning or implementing Oracle Database 19c with and without Oracle RAC when running under Linux for IBM Z and LinuxONE servers. In almost all cases, there are reinforcing references included in each section. The “Resources” appendix below contains pointers to general documentation and additional supporting documents.



Resources

These websites and documents provide useful references to supplement the information contained in this document:

- Oracle Database 19c Documentation Library:
<https://docs.oracle.com/en/database/oracle/oracle-database/19/index.html>



About the authors

Michael Morgan is a Senior Oracle Consultant currently working on certifications of the Oracle Database running Linux on IBM LinuxONE and Z servers in the Oracle Redwood Shores office. Before joining IBM as a consultant in 2015, he worked at Oracle Corporation in various system administration, management, and development roles.

Srujan D. Jagarlamudi is a Senior Oracle Consultant currently working with Linux running on IBM LinuxONE and Z server certifications in the Oracle Redwood Shores office. Before joining IBM as a consultant in 2015, he worked in several IT companies as a Senior Oracle DBA and Applications DBA and played a major role in upgrade projects / migration to Oracle Databases and applications.

David J. Simpson is part of the IBM North America Channels and ISV Solutions technical sales team with almost 30 years of experience in the IT Industry. He is an Oracle certified Oracle specialist working with Oracle Database since version 5 of this product.



Trademarks and special notices

© Copyright. IBM Corporation 1994-2020. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, IBM Z, z13, z15 and z/VM are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.