

# Oracle Database 19c & Oracle Database 19c RAC on IBM AIX

*Tips and Considerations*



*Ravisankar Shanmugam*  
*IBM Oracle International Competency Center*  
*October 2019*



## Table of contents

<b>Abstract</b> .....	<b>1</b>
<b>Introduction</b> .....	<b>1</b>
<b>Terminology</b> .....	<b>1</b>
<b>IBM AIX</b> .....	<b>2</b>
AIX support .....	3
Recommended Code Levels .....	3
Service Strategy .....	4
C and C++ compilers .....	4
<b>Oracle Database 19c (Non RAC)</b> .....	<b>4</b>
Current certifications .....	5
Oracle Database Latest patch set .....	6
Release Update (RU) and Release Update Revisions (RUR) .....	6
Virtualization for Oracle Database 19c (Non RAC) .....	6
<b>Oracle Database 19c RAC</b> .....	<b>8</b>
Current certifications .....	8
Oracle Redundant Interconnect usage feature .....	9
Oracle Clusterware .....	10
IBM PowerHA (formerly High Availability Cluster Multi-Processing) .....	10
IBM Spectrum Scale (formerly General Parallel File System) .....	10
Virtualization for Oracle Database 19c RAC .....	11
Network interconnects .....	12
Important Oracle fixes or issues specific to AIX releases .....	13
<b>Tuning tips</b> .....	<b>13</b>
Automatic Storage Management .....	14
Managing raw disks in AIX to use with ASM .....	14
Asynchronous I/O .....	15
IBM Power Systems – server-side caching .....	15
RAC IPC .....	15
IBM Spectrum Scale (formerly General Parallel File System) .....	16
Oracle process memory footprint.....	16
Tuning resources .....	16
<b>Summary</b> .....	<b>18</b>
<b>Resources</b> .....	<b>19</b>
<b>About the author</b> .....	<b>20</b>
<b>Trademarks and special notices</b> .....	<b>21</b>

## 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 the IBM AIX® operating system.*

*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, please consult Oracle's "My Oracle Support" website.*

*This paper is written to a level of detail that assumes readers have an in-depth knowledge of AIX, Oracle Database 19c, Oracle Database 19c RAC and related products.*

## Introduction

---

There are many technical topics to consider when running Oracle Database 19c and Oracle Database 19c RAC on AIX such as: status of Oracle certifications, 19c patch sets, AIX code levels, AIX 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 AIX 7.1 and AIX 7.2, since these are the AIX versions supported for Oracle Database 19c (standalone) and Oracle Database 19c RAC.

This is a companion paper to another paper:

- For Oracle Database 12c, *Oracle Database 12c and Oracle RAC 12c on IBM AIX: Tips and Considerations*,  
<https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102425>.

The IBM Oracle International Competency Center (ICC) works closely with the IBM Oracle Center (IOC) in Montpellier, France.

- IBM Oracle Center, Montpellier, France – [ioc@fr.ibm.com](mailto:ioc@fr.ibm.com)

**Note:** Some of the "My Oracle Support" notes and web links referenced in this document for Oracle Database 19c on IBM Power Systems with AIX 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](mailto:ibmoracle@us.ibm.com).

## Terminology

---

In 2008, the IBM System p® and System i™ product families were combined to create the IBM Power® Systems product family. The IBM Power Systems product family includes systems previously referred to as System p, System p5®, eServer™ pSeries, RS/6000™, System i, System i5®, eServer iSeries, AS/400, IBM Flex System®, IBM PureFlex® and IBM PureSystems®.

The brand name "AIX 5L" is no longer used. The term "AIX 5L" will still appear in some places in this document, mostly in document titles which have yet to be updated.



High Availability Cluster Multi-Processing (HACMP) has been renamed to the PowerHA® and PowerHA SystemMirror®. This version of the document will use the name HACMP for versions up to 5.4.1, PowerHA for version 5.5 and PowerHA SystemMirror for versions 6.1 and up.

## IBM AIX

---

IBM AIX is an open standards-based, UNIX® operating system. AIX in combination with IBM's virtualization offerings provides new levels of flexibility and performance to allow you to consolidate workloads on fewer servers, which can increase efficiency and conserve energy. AIX delivers high levels of security, integration, flexibility, scalability and reliability—essential for meeting the demands of today's information technology environments. AIX operates on IBM Systems based on Power Architecture technology. For more information about AIX, see this web page: <https://www.ibm.com/it-infrastructure/power/os/aix>

IBM and Oracle are advancing to focus on supporting AIX 7.1 and AIX 7.2 for Oracle Database 19c. IBM's latest certified AIX version for Oracle Database is AIX 7.2. It contains new features for virtualization, security, availability and manageability. AIX 7.2 is binary compatible with AIX 7.1, AIX 6.1 and AIX 5.x. AIX 7.2 includes support for servers based on the POWER8® and POWER9™ processors.

Most of the new features of AIX 7.2 and AIX 7.1 are available on earlier POWER® platforms. However, the best capabilities are delivered on systems based on POWER9 and POWER8 processors. POWER7® and later processors are built for a Smarter Planet® and automatically optimize workload performance and capacity. New intelligent threads technology in POWER9 and POWER8 processors dynamically switch the processor threading mode to deliver optimal performance - either the highest per-thread performance or the maximum system throughput based on application requirements.

POWER9 and POWER8 Hardware Accelerators improve the efficiency of memory expansion even more effectively allowing more memory expansion.

In 2018, IBM introduced the next generation of Power Systems servers, the S924, S922, S914, E950 and E980 servers with POWER9 processor technology.

These Power Systems Scale Out servers (server model names starting with letter “S”) with POWER9 processors come with a large memory footprint of up to 4 TB of DDR4 memory, up to 24 high speed POWER9 processor cores in one or two Single Chip Module (SCM) sockets, PCIe Gen4 to add more performance and flexibility for I/O configurations and Storage backplane options. This system is designed to run commercial, cognitive, and database workloads.

The Power Systems Enterprise servers (server model names starting with letter “E”) with POWER9 processors, the E950 server comes with up to 48 cores in four Single Chip Module sockets and up to 16 TB of DDR4 memory. The E980 server comes with up to 192 POWER9 processor cores, up to 64 TB of DDR4 memory.

Customers can continue to choose IBM POWER8 processor-based servers from one socket scale out server with up to 8 cores to 16 sockets enterprise class servers up to 192 cores.

For information about the IBM Power Systems servers, see this web page: <https://www.ibm.com/it-infrastructure/power>

Refer My Oracle Support Note **1307544.1**, “Certification Information for Oracle Database on IBM AIX on Power Systems” for AIX operating system general checklist and complete operating system requirements.

IBM also produces benchmark reports that demonstrate the performance results of Power Systems running AIX. These benchmarks are run using Oracle E-Business Suite, yet they provide valuable data about Oracle Database and Oracle RAC. The benchmark results are published at

<http://www.oracle.com/us/solutions/benchmark/apps-benchmark/results-166922.html>

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 AIX can be found at:

- IBM Power Systems and AIX Information Center:  
[https://www.ibm.com/support/knowledgecenter/ssw\\_aix](https://www.ibm.com/support/knowledgecenter/ssw_aix)
- IBM developerWorks for AIX:  
<https://www.ibm.com/developerworks/aix/>

## AIX support

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

In general, the IBM Oracle Competency Center recommends that you keep your TL’s and SP’s up to date for your AIX installation. Monitor My Oracle Support notes and IBM Fix Central for the latest issues.

Always perform thorough testing on an OS update before deploying to production.

### Recommended Code Levels

There are two tools to help determine the recommended code levels among AIX and Power Systems related components.

1. The Fix Level Recommendation Tool (FLRT) can determine the recommended code levels among a mixture of AIX, Hardware Management Console (HMC), Server firmware, Virtual IO server (VIOS), Spectrum Scale and PowerHA SystemMirror. The FLRT web page is: <http://www14.software.ibm.com/webapp/set2/flrt/home>. Note, the FLRT recommendation provides a minimum acceptable level of compatibility.
2. The POWER code matrix indicates the recommended code levels for the HMC and Server firmware. The POWER code matrix web page is: [https://www14.software.ibm.com/webapp/set2/flrt/mtm#tab\\_p5](https://www14.software.ibm.com/webapp/set2/flrt/mtm#tab_p5). Note, the POWER code matrix recommendation provides the maximum stability code combinations.
3. IBM Systems to AIX map:  
<https://www-01.ibm.com/support/docview.wss?uid=ssm1platformaix>
4. IBM Systems to VIOS map  
<https://www-01.ibm.com/support/docview.wss?uid=ssm1platformvios>

## Service Strategy

- To review the latest *IBM AIX Operating System Service Strategy Details and Best Practices* document, see this website: <http://www14.software.ibm.com/webapp/set2/sas/f/best/home.html>.

## C and C++ compilers

XL C/C++ is a standards-based, high performance C and C++ compiler with advanced optimizing and debugging features. It gives you the ability to optimize and tune applications for optimal execution on systems using POWER8 and the latest POWER9 processors from IBM. The compiler is focused on IBM Power Systems servers capable of running IBM AIX 7.1 and 7.2.

To determine the **current certifications** for compilers with Oracle Database 19c, please review My Oracle Support note **43208.1**, “*Certified Compilers*”. Currently, XL C/C++ 13.1.3.3 and later are supported for Oracle Database 19c. If necessary, read My Oracle Support note **335569.1**, “*How to Find C or C++ version on AIX Platform*” to determine your XL C/C++ compiler version.

The XL C/C++ runtime environment is installed during the installation of base AIX. To update to the latest runtime environment, go to the web page: [Fix list for XL C/C++ Runtime for AIX](#).

If the XL C/C++ Enterprise Edition for AIX compiler is installed, to get the latest updates, visit the product’s web page, [Fix list for XL C/C++ for AIX](#).

## Oracle Database 19c (Non RAC)

This section contains the Oracle Database 19c non-RAC technical information that needs to be considered in an AIX 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 Automatic Storage Management (ASM) and Oracle Restart. Oracle Restart is a new feature that provides the ability to monitor, manage, and automatically restart the Oracle components, including the Oracle Database Single Instance, Oracle Net Listener, database services and Oracle ASM. Grid infrastructure is needed for the stand-alone database server if the ASM is selected as a storage option for the database files.

In May 2019, Oracle released Oracle Database 19c IBM AIX. 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.

The Oracle Database 19c new features can be found in a document available at: <https://docs.oracle.com/en/database/oracle/oracle-database/19/newft/database-new-features-guide.pdf>

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

### Oracle Database 19c:

- Oracle Grid Infrastructure Installation and Upgrade Guide 19c for IBM AIX on POWER Systems (64-Bit)

E96274-02, <https://docs.oracle.com/en/database/oracle/oracle-database/19/cwaix/grid-infrastructure-installation-and-upgrade-guide-ibm-aix-power-systems-64-bit.pdf>

- Oracle Database Installation Guide 19c for IBM AIX on POWER Systems (64-bit)

E96437-01: <https://docs.oracle.com/en/database/oracle/oracle-database/19/axdbi/database-installation-guide-ibm-aix-power-systems-64-bit.pdf>

- Oracle Database Release Notes - 19c for IBM AIX on POWER Systems (64-Bit)

E96072-11: <https://docs.oracle.com/en/database/oracle/oracle-database/19/rnrdm/database-release-notes.pdf>

- 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 on AIX:

- 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 “IBM AIX on POWER Systems (64-bit) 7.2” or “IBM AIX on POWER Systems (64-bit) 7.1”, then click the “Search” button. The search result will show the certification status in the link “See Certification Details for Notes and Support information”.

Additionally, refer “My Oracle Support” document **1307544.1** for the latest certification information for Oracle Database on IBM AIX on Power Systems.

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

Operating System Level	Product	Status
AIX 7.1 TL05 SP01 or later AIX 7.2 TL02 SP01 or later	Oracle Database 19c	Certified

Table 1: Oracle Database Enterprise Edition 19c Release 1 certifications of AIX as of publication date

The IBM Power System servers with **POWER9** processor need **minimum recommended AIX levels** as shown below,

- AIX 7.2 TL02 SP02 or later
- AIX 7.1 TL05 SP02 or later

For Oracle Database 19c, download the following APARs from IBM fix central and apply:

- AIX7.2 TL02 : IJ06143 – GETSOCKNAME returns invalid path name for AF\_UNIX sockets
- AIX7.1 TL05 : IJ04311 – GETSOCKNAME returns invalid path name for AF\_UNIX sockets

Here are some certification details to be aware of:

- These products are certified for AIX 7.1, and 7.2 on all Power Systems servers supported by those versions of AIX.

- 64-bit hardware is required for a 64-bit application such as Oracle Database 19c for AIX.
- Servers capable of more than 4 processors are certified only for EE.
- Logical Partitioning (LPARs) and Micro-Partitioning® are supported.
- Virtual IO Server (VIOS) is supported.

## Oracle Database Latest patch set

Oracle Database 19c version 19.0.0.0.0 was released as Standard Edition 2 (SE2), Enterprise Edition (EE), Personal Edition (PE), and other editions for the Cloud.

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 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 Database Release Update (RU) and Release Update Revision (RUR) see My Oracle Support note **2285040.1** for Release Update and Release Update Revisions for Database Proactive Patch Program.

The latest Release Update patch for Oracle Database 19c while writing this paper was **29708769** (19.4.0.0.0). 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 AIX Operating System, select "IBM AIX on POWER Systems (64-bit)". 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 29708769 comes with patches for both Grid Infrastructure and Database.

This will list a table of patches and pick a relevant patch to suit your installed database software. The system patch usually comes with patches for both Grid Infrastructure and 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 PowerVM® features such as LPARs, Micro-Partitioning and VIOS are supported. PowerVM is the family of technologies, capabilities and offerings that deliver industry-leading virtualization on IBM POWER processor-based systems.

Power System features	Certified With	Status
Dedicated processor LPAR Shared processor LPAR	Oracle DB 19c	Certified
Virtual I/O server	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.
- **Shared Processor partitions** are LPARs that use Micro-Partitioning in conjunction with a shared processor pool.
  - **Micro-Partitioning** divides a physical processor's computing power into fractions of a processing unit and shares them among logical partitions. Processing capacity can be configured in fractions of 1/100 of a processor. The minimum amount of processing capacity that must be assigned to a partition is 1/10 of a processor.
  - A **shared processor pool** is a group of physical processors that are not dedicated to any LPAR.
- **Virtual I/O Server** allows sharing of physical resources between logical partitions (LPARs) including virtual SCSI and virtual networking. This allows more efficient utilization of physical resources through sharing between LPARs and facilitates server consolidation.

VIOS 2.2.2.3 and above is recommended with AIX 7.2, and 7.1 for Oracle Database 19c.

VIOS 2.2.6.21, or later is recommended for IBM Power Systems with POWER9 processors.

With certified VIOS combinations customers may use virtual SCSI with or without "N Port ID virtualization" (NPIV) to attach disk for data storage and associated voting files or OCR. This may be done for both ASM and GPFS.

The VIOS Support web page is <http://www14.software.ibm.com/webapp/set2/sas/f/vios/home.html>

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

For additional, detailed PowerVM information, see these documents and websites:

<https://www.oracle.com/database/technologies/virtualization-matrix.html>

- PowerVM webpage, <http://www-03.ibm.com/systems/power/software/virtualization/index.html>
- PowerVM Virtualization on IBM System p: Introduction and Configuration Fourth Edition, <http://www.redbooks.ibm.com/abstracts/sg247940.html?Open>
- IBM System p PowerVM Best Practices, <http://www.redbooks.ibm.com/abstracts/redp4194.html?Open>

## Oracle Database 19c RAC

---

This section contains the Oracle Database 19c RAC technical information that needs to be considered in an AIX 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 RAC installation on AIX.

### Oracle Database 19c:

- 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:  
E95728-04, <https://docs.oracle.com/en/database/oracle/oracle-database/19/racad/real-application-clusters-administration-and-deployment-guide.pdf>
- Oracle Grid Infrastructure Installation and Upgrade Guide 19c for IBM AIX on POWER Systems (64-Bit):  
E96274-02, <https://docs.oracle.com/en/database/oracle/oracle-database/19/cwaix/grid-infrastructure-installation-and-upgrade-guide-ibm-aix-power-systems-64-bit.pdf>
  - Oracle Database Installation Guide 19c for IBM AIX on POWER Systems (64-bit):  
E96437-01: <https://docs.oracle.com/en/database/oracle/oracle-database/19/axdbi/database-installation-guide-ibm-aix-power-systems-64-bit.pdf>

Consider the following additional requirement for a successful installation of Oracle 19c Grid infrastructure software.

- Install “bash” shell tool on all the RAC nodes prior to Oracle GI installation.
  - bash-4.2-1 (This can be downloaded from the following link):  
<http://www-03.ibm.com/systems/power/software/aix/linux/toolbox/date.html>)

### I/O Completion Ports:

On AIX on POWER systems, enable I/O completion ports (IOCP) to ensure successful database and grid infrastructure installation.

To check if the IOCP module is enabled, run the following command and look for status “Available” in the output,

```
$ lsdev |grep iocp
```

```
iocp0 Available I/O Completion Ports.
```

If IOCP is in “Defined” state, enable it using “smitty” tool,

```
$ smitty iocp
```

This setting requires reboot of the RAC nodes.

## Current certifications

To determine the **current certifications** for Oracle Database 19c on AIX, 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 “Certification Search” section. In the “Product” field. Type “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 “IBM AIX on POWER Systems (64bit) 7.2”, it actually lists the AIX versions 7.1 and 7.2. Choose one of them. 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.

From February 17, 2015 onwards, IBM General Parallel File System (GPFS) is rebranded as IBM Spectrum Scale.

To find the latest certification information of IBM Spectrum Scale, PowerHA System Mirror with Oracle Database, refer to the Certification section of My Oracle Support. For example, Click “Certification” tab and start to type “IBM PowerHA” in the Product field, it will list out the relevant products and choose “IBM PowerHA SystemMirror(PowerHA)”, then in the Release field select “version”, and click “Search” button.

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

By following the navigation instructions in the first paragraph of this sub-section, the RAC Technologies Compatibility Matrix will also be visible. This matrix supplies details about storage technologies, network interconnect technologies and other platform-specific information.

<http://www.oracle.com/technetwork/database/clustering/tech-generic-unix-new-166583.html>

For Oracle Database and Server virtualization features support, look at the information for “IBM AIX Power Systems” 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’s Grid Infrastructure Redundant Interconnect, also known as Cluster High Availability (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 where instances of the same database are not co-located on the same physical frame.

For upgrade customers, it is recommended to maintain their current, typically Etherchannel based, configuration with 11.2.0.4, 12.1.0.2, 12.2.0.1, 18c releases during upgrades. This will allow the Redundant Interconnect Usage to allocate an (HAIP) on top of the Etherchannel 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 Etherchannel in this case; 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, the Etherchannel configuration used for the Oracle interconnect should be removed and Oracle Redundant Interconnect Usage should be enabled directly on the devices formerly managed by Etherchannel. For more information refer to the Oracle Documentation on how to enable Redundant Interconnect Usage.

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 and one VIO server) 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 VIO 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 RAC. For more information on Oracle Clusterware, visit: <http://www.oracle.com/technology/products/database/clustering/index.html>.

## IBM PowerHA (formerly High Availability Cluster Multi-Processing)

Note: High Availability Cluster Multi-Processing (HACMP) has been renamed to the PowerHA SystemMirror for version 7.1. This version of the document will use the term PowerHA.

Oracle no longer supports the use of raw logical volumes with the Oracle Database and with RAC, see My Oracle Support note “Announcement of De-Support of using RAW devices in Oracle Database Version 12.1” (Doc ID **578455.1**). Oracle continues to support the coexistence of PowerHA with Oracle Clusterware.

IBM PowerHA SystemMirror 7.1 is supported for use with Oracle Database 19c RAC.

## IBM Spectrum Scale (formerly General Parallel File System)

If using a file system for your Oracle Database 19c RAC data files (rather than ASM), you'll need to use a cluster file system (CFS). Oracle ACFS allows file system access by all members in a cluster at the same time. That requirement precludes JFS and JFS2 from being used for Oracle Database 19c RAC data files. IBM Spectrum Scale is an Oracle RAC 19c certified cluster file system.

IBM Spectrum Scale is a high-performance shared-disk file system that can provide fast, reliable data access from all nodes in a homogenous or heterogeneous cluster of IBM UNIX servers running either the AIX or the Linux® operating system.

To determine the **current certifications** for IBM Spectrum Scale on Oracle RAC, refer to the “Certification” section on the My Oracle Support website.



IBM Spectrum Scale 4.1 is certified for use with Oracle Database 19c RAC on AIX 7.1 and 7.2.

The following stacks are recommended:

- IBM Spectrum Scale 4.1
- Oracle Database 19c RAC
  - AIX7.2 TL1 SP01
  - AIX7.1 TL05 SP01
- VIOS 2.2.4 or above
- VIOS 2.2.6.21, or later is recommended for IBM Power Systems with POWER9 processors.

Refer to the My Oracle Support note **1376369.1** for more information on GPFS for Oracle RAC.

The GPFS website is at: [https://www.ibm.com/support/knowledgecenter/en/SSFKCN/gpfs\\_welcome.html](https://www.ibm.com/support/knowledgecenter/en/SSFKCN/gpfs_welcome.html).

Be sure to also review the **current GPFS advisories** in the GPFS FAQ available from the GPFS website at:

[https://www.ibm.com/support/knowledgecenter/SSFKCN/com.ibm.cluster.gpfs.doc/gpfs\\_faqs/gpfsclustersfaq.html](https://www.ibm.com/support/knowledgecenter/SSFKCN/com.ibm.cluster.gpfs.doc/gpfs_faqs/gpfsclustersfaq.html).

See the GPFS sub-section in the Tuning Tips section (below) for GPFS tuning information.

## Virtualization for Oracle Database 19c RAC

As mentioned in the “Current certifications” sub-section (see above) for Oracle Database 19c RAC, LPARs and Micro-Partitioning are supported in Oracle Database 19 RAC environments. LPARs, Micro-Partitioning and VIOS are part of IBM PowerVM. PowerVM is the family of technologies, capabilities and offerings that deliver industry-leading virtualization on IBM POWER processor-based systems. The PowerVM capabilities supported in Oracle Database 19c are:

Power System features	Certified With	Status
Dedicated processor LPAR Shared processor LPAR Micro-Partitioning	Oracle DB 19c RAC	Certified
Virtual I/O server	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 processors that are assigned to a single LPAR.
  - **Shared Processor partitions** are LPARs that use Micro-Partitioning in conjunction with a shared processor pool.
- **Micro-Partitioning** divides a physical processor's computing power into fractions of a processing unit and shares them among logical partitions. Processing capacity can be configured in fractions of 1/100 of a processor. The minimum amount of processing capacity that must be assigned to a partition is 1/10 of a processor.
- A **shared processor pool** is a group of physical processors that are not dedicated to any LPAR.

- For dedicated processor partitions, it is only possible to dynamically add, move or remove whole processors. When a processor is removed from a dedicated processor partition, it is then assigned to the shared processor pool.
- For a shared processor partition, it is also possible to dynamically change the shared processor capacity, the weight of the uncapped attribute, virtual processors and capped/uncapped mode.
- **Virtual I/O Server** allows sharing of physical resources between logical partitions (LPARs) including virtual SCSI and virtual networking. This allows more efficient utilization of physical resources through sharing between LPARs and facilitates server consolidation. These VIOS features are specifically supported in Oracle Database 19c RAC:
  - Virtual LAN for public and private interconnects and all supported data storage options.
  - Two VIO servers are necessary to improve serviceability and provide the availability required for an Oracle Database 19c RAC implementation.
  - VIOS can be used for non-RAC functions such as Virtual SCSI based root volume.
  - groups (rootvg), and Virtual SCSI & NPIV (N port ID Virtualization) for ASM data storage.
  - VIOS with ASM and GPFS in Oracle Database 12c Release 1 & 2 RAC is supported.

In VIOS environment, NPIV feature allows multiple N\_Port IDs to share a single Physical N\_Port, which means multiple client LPARs can access external storage LUNs through the same Fibre channel adapter assigned to VIOS LPAR.

Note: With certified VIOS combinations customers may use either Virtual SCSI (vSCSI) or N Port Id Virtualization (NPIV) to attach disk for data storage and OCR. This may be done for both ASM and GPFS 4.1 and above. Customers must confirm IBM support of the configuration and install any required AIX and Oracle updates before using.

VIOS 2.2.4 and above are supported with Oracle Database 19c RAC with AIX 7.1 and 7.2 respectively.

VIOS 2.2.6.21, or later is recommended for IBM Power Systems with POWER9 processors.

The VIOS Support website is: <http://www14.software.ibm.com/webapp/set2/sas/f/vios/home.html>.

- PowerVM website:  
<http://www-03.ibm.com/systems/power/software/virtualization/index.html>
- IBM PowerVM Virtualization Introduction and Configuration, Fourth Edition:  
<http://www.redbooks.ibm.com/abstracts/sg247940.html?Open>

## Network interconnects

To determine the **current certifications** for network interconnects on Oracle RAC, browse to the “General Notes for RAC for Unix on IBM AIX based Systems (RAC only)” web page, as described in the beginning of the “Current certifications” sub-section (see above) for Oracle Database 19c 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 AIX 7.1, and 7.2:

- 1 Gigabit
- 10 Gigabit Ethernet
- 40 Gigabit Ethernet

For supported Interconnect Network Technologies for Oracle RAC on UNIX, refer the following link, <http://www.oracle.com/technetwork/database/clustering/tech-generic-unix-new-166583.html>

Some additional consideration as part of a 10 GigE setup are as follows:

- LACP timeout: Use the “long timeout” switch setting for the amount of time to wait before sending LACPDU.
- Flow control: Enable flow control at the switch port and on the server-side ports (using HMC) for the 10GE adapter or 10GE HEA configuration.
- UDP tuning: The RAC interconnect uses UDP for interprocess communications. Tune the `udp_sendspace` and `udp_recvspace` parameters until the “`netstat -s`” command indicates there are no “socket buffer overflows”. See the Network Tuning section of the “Oracle Architecture and Tuning on AIX” document:  
<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100883>
- Jumbo frames: enable Jumbo frames on the RAC interconnect: Configure Jumbo frames at the switch port. In the certification project we set Jumbo frames to 9252 at the switch.

Configure Jumbo frames on the server-side ports (using HMC) for the 10GE adapter configuration. Configure Jumbo frames in the Etherchannel (network interface) definition.

Note: When using Gigabit Ethernet, 10 Gigabit Ethernet and IP over InfiniBand customers may configure the network routing using either EtherChannel or AIX VIPA based on their requirements. Customers must confirm IBM support of the configuration and install any required AIX and Oracle updates before using.

### Important Oracle fixes or issues specific to AIX releases

While creating an ASM diskgroup with more than one disk or adding a disk to an existing diskgroup ORA-15020 is encountered in Oracle Database 12c Release 2 (12.2.0.1).

Solution: Add the 'name' clause to the 'alter diskgroup' command.

Example: ALTER DISKGROUP DATA ADD DISK '/dev/ASM\_hdisk2' NAME DATA\_02;

Refer My Oracle Support note **2355979.1** for more details.

**Grant DBA to user can cause latch row cache objects:** granting “DBA” role to all application users can result in huge waits with an event “latch: row cache objects” which affects performance of the database. To resolve this issue, refer My Oracle support note **1639293.1**. If this issue occurs, explicitly grant access to the underlying object to the user (not through a Role).

For more issues and tuning recommendations, refer “Oracle 18c, 12c and 11.2.0.4 Database Performance Considerations with AIX on Power Systems including POWER9” paper found at: <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102608>

## Tuning tips

---

Most of the AIX OS tuning tips recommended for lower version of Oracle Database are applicable to Oracle Database 19c.

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.

Starting with Oracle Database 11g Release 2, Oracle Clusterware OCR and voting disk files can be stored in an Oracle ASM disk group.

ASM becomes a complete storage management solution for both Oracle Database and non-database files and has many extended functions for not only storing database files, but also storing binary files, report files, trace files, alert logs and other application data files.

ASM Cluster File Systems (ACFS) extends ASM by providing cluster file system scaled to a large number of nodes and uses extend-based storage allocation for improved performance. ACFS can be exported to remote clients through NFS and CIFS.

ASM Dynamic Volume Manager (DVM), ASM FS Snapshot, 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 has 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” at:

<https://docs.oracle.com/en/database/oracle/oracle-database/19/newft/database-new-features-guide.pdf>.

### Managing raw disks in AIX to use with ASM

To prevent the accidental use of disks which are assigned to ASM disk group in AIX for some other purpose, there are two new AIX commands introduced in AIX 6.1, 7.1 and AIX 7.2.

The two commands “lkdev” and “rendev” are introduced to effectively use the disks for ASM disk groups. The “lkdev” command is used to lock the disk to prevent the device from inadvertently being altered by a system administrator later. It locks the device so that any attempt to modify the device attributes (chdev, chpath) or remove the device or one of its paths (rmdev, rmpath) will be denied.

The “rendev” command can be used to assign a meaningful name to the disks used by Oracle Database, ASM, Cluster Ready Services and Voting files. This is useful in identifying disk usage

because there is no indication in the output from AIX disk commands indicating that a disk is being used by the Oracle Database.

For more details of using these commands, refer the “My Oracle Support” note **1445870.1** – “Managing Raw Disks in AIX to use with Oracle ASM”.

For information about ASM, see Oracle’s ASM web page:  
<http://docs.oracle.com/database/121/OSTMG/toc.htm>

## Asynchronous I/O

Asynchronous I/O (AIO) allows a program to initiate an I/O operation then continue with other work in parallel to the I/O operation. Oracle Database 19c often requires multiple server and user processes running at the same time. Therefore Oracle Database 19c takes full advantage of AIO services provided by AIX. AIO is implemented with AIO server processes. The configuration values of: minservers, maxservers and maxreqs control the AIO server configuration of AIX. Refer to the “Tuning resources” section of this document for links to many whitepapers. The IBM whitepaper *Tuning IBM AIX 5L for an Oracle Database* has an “Asynchronous I/O” section that describes recommendations for the configuration values. GPFS configuration recommendations are also supplied. There is also a “Using Asynchronous I/O” section in the *Oracle Architecture and Tuning on AIX* whitepaper.

Note, if you use the AIX filemon command, the AIO API calls of aio\_read, aio\_write and lio\_listio are not included in the filemon report.

### IBM Power Systems – server-side caching

Server-side caching is a new feature introduced in AIX 7.1 TL04 SP02 and AIX 7.2. This feature is supported for use with the Oracle Database to improve the performance of read I/O intensive workloads on AIX and Power Systems environments. Server-side caching provides the capability to cache the application data stored on a SAN to Solid State Devices (SSD) or Flash Storage LUNs or Virtual Disks provided by VIOS on the AIX server. After server-side caching is enabled in AIX, all the read I/O requests are first redirected to the caching area created with fast SSDs, Flash Storage or VIOS virtual disk on the server. This feature can be enabled or disabled dynamically, no reboot is required, and changes are transparent to the running application or workload.

To take advantage of this feature for Oracle Database, no changes are needed in the database configuration side. This works only with Oracle Database in a non-RAC environment.

## RAC IPC

Oracle RAC with Oracle Database 19c uses the User Datagram Protocol (UDP) for inter-process communication (IPC) between nodes. The *Oracle Architecture and Tuning on AIX* whitepaper <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100883> provides UDP kernel settings tuning information in the “Network Tuning” section.

## IBM Spectrum Scale (formerly General Parallel File System)

The IBM Spectrum Scale documents can be found at the following link:

[https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale\\_welcome.html](https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale_welcome.html). My Oracle Support note **1376369.1** contains IBM GPFS and Oracle RAC certification information.

## Oracle process memory footprint

The AIXTHREAD\_SCOPE environment variable can be used to control if an AIX process runs with process-wide contention scope (the default) or with system-wide contention scope. System-wide contention scope significantly reduces the memory required for each database process. AIX operates most effectively with Oracle Database 19c and Oracle RAC when using system-wide contention scope (AIXTHREAD\_SCOPE=S).

## Tuning resources

These are the documents referenced throughout this section. The same tuning of AIX for Oracle Database 12c applies to Oracle Database 19c.

1. Oracle Database 11g and 12c on IBM Power Systems S924, S922 and S914 with POWER9 processors:  
<https://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102750>
2. Oracle 18c, 12c and 11.2.0.4 Database Performance Considerations with AIX on Power Systems including POWER9:  
<https://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102608>
3. Oracle RAC on IBM AIX best practices in memory tuning and configuring for system stability:  
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101513>
4. Oracle Architecture and Tuning on AIX (whitepaper):  
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP100883>
5. IBM Knowledge Center, IBM Spectrum Scale:  
[https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale\\_welcome.html](https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale_welcome.html)
6. GPFS: Concepts, Planning and Installation Guide (GPFS document):  
<http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.gpfs.doc/gpfsbooks.html>
7. Managing AIX Devices used by Oracle Automatic Storage Management (ASM):  
<https://www-01.ibm.com/support/docview.wss?uid=tss1wp102158>
8. Oracle DB 11g and 12c on IBM Power Systems built with POWER8 technology and IBM Flash System 840:  
<https://www-01.ibm.com/support/docview.wss?uid=tss1wp102440>

These are supplemental tuning resources.

- IBM AIX 7.2 Live Update - Experiences testing AIX 7.2 Live Update feature with Oracle Database 12c Release 2:  
<https://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102794>
- Optimizing AIX 7 memory performance: Part 1, Memory overview and tuning memory parameters:  
<http://www.ibm.com/developerworks/aix/library/au-aix7memoryoptimize1/>

- Performance Management for Power Systems:  
<http://www-01.ibm.com/support/docview.wss?uid=isg3T1026749>
- Diagnosing Oracle Database Performance on AIX Using IBM NMON and Oracle Statspack Reports (whitepaper):  
<http://www-01.ibm.com/support/docview.wss?uid=tss1wp101720>

## Summary

---

This document gathers together the key technical topics that need to be considered in planning or implementing Oracle Database 19c with and without RAC when running under AIX. 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:  
<http://docs.oracle.com/en/database/>
- Oracle My Oracle Support note **341507.1**: Oracle Products on Linux on IBM POWER:  
[https://support.oracle.com/knowledge/Oracle%20Database%20Products/341507\\_1.html](https://support.oracle.com/knowledge/Oracle%20Database%20Products/341507_1.html)
- Oracle and IBM System Storage:  
<http://www-03.ibm.com/systems/storage/solutions/compatibility/oracle/index.html>



## About the author

---

**Ravisankar Shanmugam** is an Advisory Software Engineer with the IBM Systems and Oracle ISV Enablement team at Oracle Headquarters, Redwood Shores, California. He writes technical collateral for Oracle products on IBM Power Systems with AIX. He leads Power Systems projects and provides technical support for AIX and Power Systems at the IBM Oracle ICC and onsite at Oracle Headquarters in Redwood Shores.



## Trademarks and special notices

---

© Copyright. IBM Corporation 1994-2019. 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, AIX, AIX 5L, AS/400, eServer, GPFS, PowerHA, iSeries, Micro-Partitioning, Power, PowerHA, PowerVM, Power Architecture, POWER, POWER3, POWER4, POWER5, POWER6, POWER7, POWER7+, POWER8, POWER9, PowerPC, Power Systems, pSeries, RS6000, System i, System i5, System p, System p5, and IBM Spectrum Scale are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

UNIX is a registered trademark of The Open Group in the United States and other countries.

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.