Oracle Database 11*g* R2
Oracle Database 11*g* R2 RAC on IBM AIX

Tips and Considerations

Ravisankar Shanmugam
IBM Oracle International Competency Center
July, 2018

© Copyright 2018, IBM Corporation

Table of contents

© Copyright 2018, IBM Corporation

Abstract

This paper consolidates the information necessary for planning and implementing Oracle Database 11g Release 2 (11.2) single node database or Oracle Database 11g Release 2 Real Application Cluster (RAC) on the IBM AIX® operating system.

This paper summarizes the information available at the time of publication. It will be updated as changes (eg. new certifications) occur. It is meant to be used only as a guide. For any official Oracle 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 11g Release2, Oracle Database 11g Release 2 RAC and the related products.

Introduction

There are many technical topics to consider when running Oracle Database 11*g* Release2 and Oracle Database 11*g* Release 2 RAC on AIX such as: status of Oracle certifications, 11*g* Release 2 patch sets, AIX code levels, tuning and related software components - just to name a few. 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 5.3, AIX 6.1, AIX 7.1, and AIX 7.2 since these are the AIX versions certified on Oracle Database 11*g* Release 2 and Oracle Database11*g* Release 2 RAC.

This is a companion paper to three other papers.

- For 9*i* and 10*g*R1, *Oracle 9i & 10gR1 on IBM AIX5L: Tips and Considerations*, http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100556
- For 10*g*R2, *Oracle Database 10gR2 and Oracle RAC 10gR2 on IBM AIX: Tips and Considerations*, http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101089
- For 11*g*, *Oracle Database 11g and Oracle RAC 11g on IBM AIX: Tips and Considerations*, http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101753

Feedback is important; please send any comments about this paper to the IBM Oracle International Competency Center at ibmoracl@us.ibm.com.

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

- IBM Oracle Center, Montpellier, France – ioc@fr.ibm.com
- IBM Oracle Competency Center, Tokyo, Japan – ibmoracc@jp.ibm.com

Terminology

In 2008, the 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®, RS6000®, System i™, System i5™, eServer iSeries®, and AS/400®.

© Copyright 2018, IBM Corporation

The brand name of “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 version up to 5.4.1, PowerHA for version 5.5 and PowerHA SystemMirror for the version 6.1 and up.

IBM General Purpose File System (GPFS) is rebranded to IBM Spectrum Scale from version 4.1.1 onward.

IBM AIX

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 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: http://www-03.ibm.com/systems/power/software/aix/

The latest certified AIX version with Oracle Database is AIX7.2. 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 and AIX 7.1 include supporting high-end POWER7 server with 256 CPUs, which can provide up to 1024 logical processors with use of SMT4. The IBM latest POWER8 server supports up to 192 physical cores with maximum of 1536 Logical processors with use of SMT8.

For additional information about AIX 7.2, and AIX 7.1, visit this web page:

https://www.ibm.com/it-infrastructure/power/os/aix

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™, POWER8™, POWER7+, POWER7™, and ,POWER6™ processors.

On March 20, 2018 IBM introduced the next generation of new Power Systems servers, S924, S922, and S914 with POWER9 processor technology.

These new Power Systems servers with POWER9 processors come with large memory footprint 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.

Customers can continue to choose IBM POWER8 processor based servers from one socket scalable server with up to 8 cores to 16 sockets enterprise class servers up to 192 cores. For information about the POWER9, POWER8, POWER7 and POWER6 processors, see this web page:

http://www-03.ibm.com/systems/power/advantages/power.html.

The IBM POWER9 and POWER8 processor chip provides up to eight hardware threads (SMT8) per core. Each SMT hardware thread is represented as a logical processor in AIX, IBM i, or Linux. To gain the full

© Copyright 2018, IBM Corporation

benefit from the functionality of SMT8, the application must use all of the SMT threads of the processor cores.

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

The IBM latest Power System Server with POWER8 processors can have up to 192 cores and up to 16 TB of memory.

Refer the document “Oracle Database 11g and 12c on IBM Power Systems S924, S922 and S914 with POWER9 processors” posted at the following link for POWER9, AIX and Oracle Database compatibility information.

https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102750

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/index.html

Additional documentation resources for AIX can be found at:

- IBM Power Systems and AIX Information Center, http://publib16.boulder.ibm.com/pseries/index.htm
- IBM developerWorks 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: https://www-945.ibm.com/support/fixcentral

Be sure to review My Oracle Support note **282036.1**, “Minimum Software Versions and Patches Required to Support Oracle Products on IBM pSeries”, for the latest, up-to-date issues regarding AIX and Oracle Database 11*g*R2. This My Oracle Support note is the primary mechanism used to broadcast any breaking news, such as PTF’s or bug fixes, about AIX and Oracle Database 11*g*R2.

In general, the IBM Competency Center recommends that you keep your TL’s and SP’s up to date for your AIX installation. Monitor My Oracle Support note **282036.1** and Fix Central for the latest issues. Always perform thorough testing on an OS update before deploying to production.

Refer My Oracle Support note **169706.1** “Oracle Database (RDBMS) on Unix AIX, HP-UX, Linux,Mac OS X,Solaris,Tru64 Unix Operating Systems Installation and Configuration Requirements Quick Reference (8.0.5 to 11.2)” also for AIX pre-requisite requirement for Oracle Database 11g Release 2.

Note: AIX 6.1 TL5 SP1 is recommended for POWER7 for the issue -Threads are not being utilized efficiently in SMT4 and Performance impact observed.

Refer to the document “*Oracle 12.1.x and 11.2.0.4 Database Performance Considerations with AIX on POWER8*” for tuning Oracle DB and potential list of bugs from both AIX and Oracle DB.

© Copyright 2018, IBM Corporation

https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102608

The AIX APARs mentioned in the above link are not specific to POWER8 processor based servers, but also apply to AIX6.1, AIX 7.1 and AIX7.2 irrespective of POWER processor types.

In addition to the AIX APARs mentioned in above document, consider to apply the following APARs on the latest AIX 6.1, 7.1 and 7.2 versions.

IBM recommends to apply the latest technology level (TL) and Service Pack (SP) of AIX versions to fix many issues.

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

AIX 7.2 TL2 SP02 or later
AIX 7.1 TL05 SP02 or later
AIX 6.1 TL09 SP11 or later

For Oracle Database running **on AIX 7.2 TL0 SP01**, download the following APARs from IBM fix central and apply.

IV79639 - after live update ifix state may be left as Q; reboot required

IV79848 - mirrorvg/syncvg on minimal and migration install fails

V80412 - system crash application sets signal mask

IV81482 - MULTIBOS MAY FAIL TO MOUNT OR REMOVE A STANDBY INSTANCE

IV82224 - CORE DUMP IN MANY COMMANDS WHEN USING NIS

For Oracle Database running **on AIX 7.2 TL1 SP01**, download the following APARs from IBM fix central and apply.

IV94362: SOME APPLICATIONS MAY FAIL AFTER BEING LINKED ON AIX 7.2 TL1 APPLIES TO AIX 7200-01

If you are using **AIX 7.1 TL02 SP03 for Oracle Database**, obtain and apply the following APARs,

IV48895: A SPECIAL-PURPOSE LINKER FLAG WORKS INCORRECTLY. APPLIES TO AIX 7100-02

IV48898: ADD ABIILITY TO REORDER TOC SYMBOLS IN LIMITED CIRCUMSTANCES APPLIES TO AIX 7100-02

For Oracle Database **12.1 and 12.2,** obtain the following Interim fix (ifix) and apply on the AIX versions**,**

AIX 7.2 TL02: **IJ04933**: GETSOCKNAME RETURNS INVALID PATH NAME FOR AF_UNIX SOCKETS

AIX 71. TL05: **IJ04311**: GETSOCKNAME RETURNS INVALID PATH NAME FOR AF_UNIX SOCKETS

Database Corruption ORA-600 ORA-7445 errors after applying AIX SP patches - AIX 6.1TL09 SP08 or AIX 7.1 TL03 SP08 or AIX 7.1 TL04 SP03 or AIX 7.2.TL0 SP03 or AIX 7.2.TL01 SP01

- Refer My Oracle Support note 2237498.1

Affected AIX Levels	Fixed In	iFix / APAR (ftp://aix.software.ibm.com/aix/ifixes/)
o 6100-09-08	6100-09-09	IV93840
o 7100-03-08	7100-03-09	IV93884
o 7100-04-03	7100-04-04	IV93845

© Copyright 2018, IBM Corporation

- 7200-00-03 7200-00-04 IV93883
- 7200-01-01 7200-01-02 IV93885

AIX7.1 TL03: IV69422: CREATING AIOSERVER THREADS DELAYS WHILE HOLDING LOCK

- Fixed in AIX7.1 TL03 SP05
  - AIX7.1 TL02: IV69356 - Fixed in TL02 SP5
  - AIX6.1 TL08: IV70032 - Fixed in TL08 SP7
  - AIX6.1 TL09: IV11857 – Fixed in TL09 SP5

AIX7.1 TL04: IV76270: INCREASED UPPER LIMIT OF NUM_CMD_ELEMS FOR NPIV ADAPTER TO 2048

AIX6.1 TL09: IV76258

AIX7.1 TL03: IV51322 - NCARGS ALLOCATED FROM HEAP 0

AIX6.1 TL07: IV52681 – Fixed in TL07 SP10
AIX6.1 TL08: IV53167 – Fixed in TL08 SP05
AIX6.1 TL09: IV51157 – Fixed in TL09 SP03
AIX7.1 TL01: IV50813 – Fixed in TL01 SP10
AIX7.1 TL02: IV50347 – Fixed in TL02 SP05

AIX7.1 TL04: IV75433: IMPROVED DRIVER'S FAIRNESS LOGIC TO AVOID I/O STARVATION

AIX6.1 TL09: IV78454

AIX7.1 TL01

- IV23859: APPLICATIONS RUN SLOWLY WITH HIGH SYSTEM TIME
- IV21116: SYSTEM HANGS OR CRASHES WHEN APP USES SHARED CAPABILITY

AIX71 TL02

- IV21878: SYSTEM HANGS OR CRASHES WHEN APP USES SHARED CAPABILITY

Obtain and Install the APAR IZ78005 for AIX 6.1 TL05 SP1 for the issue - PERFORMANCE DEGRADATION IN IO WHEN AIO BUF SIZE IS > 128K or upgrade to AIX6.1 TL05 SP02 and UP

Obtain the following fixes for AIX 7.1 TL0 SP01 and install or upgrade to SP02 and above

- IZ87216: UNMAPPING ONE LUN ON DS5300 CAUSES ACCESS TO HANG ON OTHER LUNS
- IZ87564: 7.1. POW() SUBROUTINE IN AIX 7100-00-01-1037, BOS.ADT.LIBM

The following APAR needs on AIX 7.1 TL01 SP01 and AIX 6.1 TL07 SP01 while re-linking large Oracle library files using AIX “ld” command.

- IV09541: FILE.ATION OVFLOW REPORTED IN ERROR WHILE LINKING LARGE 11/12/13 PTF PECHANGE
  - Fixed in AIX 7.1 TL01 SP02 and above
  - Fixed in AIX 7.1 TL01 SP04 and above
  - Fixed in AIX 6.1 TL07 SP02 and above

© Copyright 2018, IBM Corporation

AIX6.1 TL06:

- IV17291**:** a JFS2 filesystem may hang (dead lock), causing all commands accessing that filesystem to hang and to be un-killable. Requires AIX reboot to clean up

AIX6.1 TL07:

- IV23851: APPLICATIONS RUN SLOWLY WITH HIGH SYSTEM TIME
- IV20880: SYSTEM HANGS OR CRASHES WHEN APP USES SHARED CAPABILITY
  - This APAR requires use of shared symbol table and has been seen on AIX 6.1 TL07 when using Oracle Database 11gR2 with online patching. This includes Oracle Database versions 11.2.0.2 and 11.2.0.3
- IV26272: REDUCE EARLY WORKING STORAGE

AIX6.1 TL08:

- IV26735: REDUCE EARLY WORKING STORAGE
- IV31961: Possible problems in certain cases with UDP receive copying kernel data to the user space from an incorrect starting point. (Oracle RAC)

AIX7.1 TL01

- IV23859: APPLICATIONS RUN SLOWLY WITH HIGH SYSTEM TIME
- IV21116: SYSTEM HANGS OR CRASHES WHEN APP USES SHARED CAPABILITY

AIX71 TL02

- IV21878: SYSTEM HANGS OR CRASHES WHEN APP USES SHARED CAPABILITY

Oracle Database 11*g* Release 2 software Online Patching on AIX 6.1, AIX 7.1 and AIX 7.2

Oracle Database 11*g* Release 2 had introduced a new feature called “online patching” for some qualified interim patches. This “online patching” integrated with *OPatch* tool and provides the ability to patch the running processes of an Oracle instance without bringing down it. Each process associated with the Oracle instance checks for the patched code at a safe execution point, and then copies the code into its process space.

The “online patching” of Oracle Database 11*g* Release 2 is available for AIX 6.1 TL07 and AIX 7.1 TL01. The customers are advised to upgrade to the mentioned AIX level for using the “online patching” feature.

The “online patching” is not recommended to use with the lower TL level of AIX 6.1 TL07 and AIX 7.1 TL01.

For more detailed information on the “online patching” on AIX, refer the following flash document,

http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102085

© Copyright 2018, IBM Corporation

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, HMC, Server firmware, VIOS, GPFS™ and PowerHA. 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: http://www14.software.ibm.com/webapp/set2/sas/f/power5cm/home.html. Note, the POWER code matrix recommendation provides the maximum stability code combinations.

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 PowerPC®, POWER3™, POWER4™, POWER5™, POWER6™, POWER7™, POWER8™ and the latest POWER9™ processors from IBM. The compiler supports IBM Power Systems servers capable of running IBM AIX 5.2, AIX 5.3, AIX 6.1, AIX 7.1 and AIX 7.2.

To determine the ***current certifications*** for compilers on 11*g* Release 2, please review My Oracle Support note **43208.1**, "Certified Compilers". Currently, XL C/C++ 8.0 and 9.0 are certified for Oracle Database 11*g*. 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 XL C/C++ 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 11*g* Release 2

This section contains the Oracle Database 11*g* non-RAC technical information that needs to be considered in an AIX installation. Oracle 11*g* Release 2 introduces the Grid Infrastructure installation for both stand alone and Real Application Cluster (RAC) databases. Grid Infrastructure is an option which includes Automatic Storage Management (ASM), the listener, and Oracle Restart. Oracle Restart is a new feature that provides the ability to monitor, manage, and automatically restart the Oracle Database instance, ASM instance, and listeners.

Grid infrastructure is needed for the stand alone database server if the ASM is selected as a storage option for the database files.

© Copyright 2018, IBM Corporation

These are the basic documents to review for an Oracle Database 11*g* installation on AIX.

- "Oracle Grid Infrastructure Installation Guide 11*g* Release 2 (11.2) for IBM AIX Based Systems E48294-03, http://docs.oracle.com/cd/E11882_01/install.112/e48294/toc.htm
- "Oracle Database Installation Guide 11*g* Release 2 for AIX on POWER Systems " E48740-03, https://docs.oracle.com/cd/E11882_01/install.112/e48740/toc.htm
- Oracle Database Release Notes 11*g* Release 2 (11.2) for AIX on POWER Systems", E23560-06, https://docs.oracle.com/cd/E11882_01/relnotes.112/e23560/toc.htm
- My Oracle Support note **282036.1**, "Minimum Software Versions and Patches Required to Support Oracle Products on IBM pSeries" for the latest, up-to-date issues regarding Oracle Database 11*g* Release 2 and AIX.
- Oracle Database Editions and Options, https://docs.oracle.com/cd/E11882_01/license.112/e47877.pdf

Current certifications

To determine the ***current certifications*** for Oracle Database 11*g* Release 2 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 one of them from the list. It has the releases from 8.1.7.4 to 12.1.0.2., next field is "Platform", select "IBM AIX on POWER Systems (64-bit", it lists the AIX versions from 5.1 to 7.1.Choose one of them. Then click "Search" button. The search result will show the certification status in the link "See Certification Details for Notes and Support information" and the certification information of Enterprise Applications, Middleware/Application servers and support to the Oracle Management and Development Tools.

The following table shows the high level certification information.

OS	Product	Status
AIX 5.3 TL09 SP01 or later AIX 6.1 TL02 SP01 or later AIX 7.1 TL0 SP01 or later AIX 7.2 TL0 SP01 or later	Oracle Database 11*g* Release 2 (64-bit)	Certified

Table 1: Oracle Database Enterprise Edition 11g Release 2 certifications of AIX as of publication date.

The same certifications are in place for Standard Editions (SE) as Enterprise Edition (EE).

Here are some certification details to be aware of:

- These products are certified for AIX 5.3, AIX 6.1, AIX 7.1 and AIX 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 11*g* for AIX.
- Servers capable of more than 4 processors are certified only for EE.
- Logical Partitioning (LPARs), Dynamic Logical Partitioning (DLPARs) and Micro-Partitioning™ are supported.
    - Virtual IO Server (VIOS) is supported.

© Copyright 2018, IBM Corporation

Latest patch set

At the time of publication the base and the latest patch set is 11.2.0.4. To find the documentation for the latest Oracle Database 11*g* Release 2 patch set, 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". Then specify Document Number **161818.1**, "Oracle Server (RDBMS) Releases Support Status Summary". Then click on "11.2.0.X", this will bring up My Oracle Support note **880782.1**, "Support Status and Alerts for Oracle 11*g* Release 2(11.2.0.X)". Then go to the "IBM AIX Based Systems (64-bit)" section and locate the ReadMe (if any) for 11.2.0.4. The patch number is 13390677 and the files (7 files) are from p13390677_112040_AIX_1of7.zip and p13390677_112040_AIX_7of7.zip.

Note: Important Changes to Oracle Database Patch Sets Starting from 11.2.0.2, My Oracle Support note **1189783.1**

> "Starting with the first patch set for Oracle Database 11*g* Release 2 (11.2.0.2), Oracle Database patch sets are full installations of the Oracle Database software. In past releases, Oracle Database patch sets consisted of a set of files that replaced files in an existing Oracle home. Beginning with Oracle Database 11*g* Release 2, patch sets are full installations that replace existing installations."

Please read the full My Oracle Support note for the detailed information.

Also, check the Oracle Database 11*g* Release 2 release notes for AIX Based Systems, http://docs.oracle.com/cd/E11882_01/relnotes.112/e23560.pdf, for additional AIX-specific information.

In general, the ICC recommends that you keep your patch set current for your Oracle Database 11*g* Release 2 installation. Always perform thorough testing on a database update before deploying to production. Monitor the My Oracle Support note **282036.1** and the Release Notes for AIX Based Systems for the latest issues.

Critical Patch Update schedule

The Critical Patch Update (CPU) program is the method by which Oracle delivers security patch updates and security fixes for all their products. A CPU is a collection of patches for multiple security vulnerabilities. It also includes non-security fixes that are required (because of interdependencies) by those security patches. Oracle provides CPUs for all product offerings on a quarterly schedule.

For more information about CPUs and watch for Critical Patch Updates at this site: http://www.oracle.com/technology/deploy/security/alerts.htm.

Patch Set Update (PSU)

Beginning with October 2009 Critical Patch Update release, Oracle delivers Patch Set Updates for all platforms on the release date. PSUs are proactive cumulative patches containing recommended bug fixes that are released on a regular and predictable schedule. PSUs are on the same quarterly schedules as the Critical Patch Update (CPU). PSUs consist of CPU, generic patch bundles, RAC patch bundle, and Data Guard patch bundles. From CPUJan2016 onwards, the 5th digit of the Patch Set Update version is changed to reflect the release date in the format YYMMDD.

The current Patch Set Update at the time of publication of this document was **11.2.0.4.180717**. Patch number for the update of Database was **27734982** and Grid Infrastructure was **27967757**. See My

© Copyright 2018, IBM Corporation

Oracle Support note **756671.1** “Oracle Recommended Patches -- Oracle Database” and look for “Current Recommended patches.

Virtualization for Oracle Database 11*g* Release 2

As mentioned in the Current Certifications sub-section (see above) for Oracle Database 11*g* Release 2; LPARs, Micro-Partitioning, Dynamic LPAR and VIOS are supported. LPARs, Micro-Partitioning, Dynamic LPAR and VIOS are part of 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 11*g* Release 2 are:

- **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 has to 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.
- **Dynamic LPARs** (DLPAR) are a shared or dedicated LPAR to which changes can be made to the amount of processors, memory and virtual or physical adapters without requiring a reboot.
    - 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.

  VIOS 2.1 (2.1.2.10. FP 22) is a minimum certified version for Oracle 11*g* R2 (11.2.0.1) with AIX 6.1 TL03 SP01.

  VIOS 2.2 (2.2.0.10.FP24) with APAR IZ86878 is a minimum certified version supported for Oracle Database 11gR2 on AIX 7.1 along with fixes for APAR IZ87564 and IZ87216 or later.

  Note: For finding recommended level of VIOS for AIX version running on Power Systems servers with POWER7 or POWER8 processors, use IBM Fix Level Recommendation Tool (FLRT),

https://www-304.ibm.com/support/customercare/flrt .

© Copyright 2018, IBM Corporation

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 raw disk-based voting 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

- **Workload Partition** (WPAR) is a software-based virtualization feature. WPARs subdivide an AIX instance into multiple environments, each hosting applications and providing isolation from applications executing in other environments. **Live Application Mobility** allows you to relocate running WPARs from one LPAR to another. For more information on WPARs, see the *Introduction to Workload Partition Management in IBM AIX Version 6.1* Redbook, http://www.redbooks.ibm.com/abstracts/sg247431.html?Open

  Oracle DB 11*g*R2 (11.2.0.2) is certified to use with WPAR in AIX 6.1 TL02 SP02 and AIX 7.1.

  Apply AIX patch IZ52319 and IZ54871 and set “export AIXPERFSTAT_SCOPE=M”
  - Only IBM AIX JFS2 and supported NAS devices can be used with WPAR. Oracle’s ASM is not supported at this time. Live Application Mobility of WPAR is not certified with Oracle at this time.

- **Live Partition Mobility** allows you to migrate running AIX and Linux LPARs and their hosted applications from one physical server to another without disrupting infrastructure services. The migration transfers the entire partition state, including the processor context, memory, attached virtual devices, and connected users. For more information on LPM, see the *PowerVM Live Partition Mobility on IBM System p* Redbook ftp://ftp.software.ibm.com/systems/power/docs/hw/p8/p8hc3.pdf

Oracle DB 11*g* R2 (11.2.0.1) (Single Instance) is certified to use with IBM PowerVM LPM on the following code levels or higher are recommended
- AIX 6.1 TL4 SP01
- Oracle Database 11.2.0.1 and up
- VIOS 2.1.2.10.FP22
- HMC V7R7.2.0
- Firmware E*340_122

For AIX 7.1, the following or higher is certified:

- AIX 7.1 TL0 SP2
- Oracle Database 11.2.0.1 and up
- VIO Server 2.2.0.11-FP-24 SP-01 plus Interim fix IZ91190
- HMC V7R7.2.0.1
- Firmware 01EM350_103

Active Memory Expansion (AME)

AME is certified to use with Oracle Database 11*g* Release 2 with the following configuration,
- Oracle Database 11*g* Release 2 version 11.2.0.2

© Copyright 2018, IBM Corporation

- AIX 6.1 TL06 SP05
- POWER7 System firmware AM730_035

IBM AIX7.2 is certified to use with Oracle Database 11*g* Release 2 for LPAR, DLPAR, AME, and LPM with ASM, LPM with Spector Scale (formerly GPFS) VIOS and WPAR.

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:
http://www.oracle.com/technetwork/database/virtualizationmatrix-172995.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 11*g* Release 2 RAC

This section contains the Oracle Database 11*g* Release 2 RAC technical information that needs to be considered in an AIX installation. Starting with Oracle Database 11*g* Release 2, 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 11*g* Release 2 RAC installation on AIX.

- “Oracle Grid Infrastructure Installation Guide 11*g* Release 2 (11.2) for IBM AIX Based Systems E48294-03. http://docs.oracle.com/cd/E11882_01/install.112/e48294/toc.htm
- “Oracle Real Application Clusters Installation Guide 11*g* Release 2 (11.2) for Linux and UNIX”, E41962-05. http://docs.oracle.com/cd/E11882_01/install.112/e41962/toc.htm
- My Oracle Support note **282036.1**, “Minimum Software Versions and Patches Required to Support Oracle Products on IBM pSeries” for the latest, up-to-date issues regarding Oracle Database 11*g* and AIX.

Note, these documents apply equally as well to AIX 5L version 5.3, AIX 6.1, AIX 7.1 and AIX 7.2.

Current certifications

To determine the ***current certifications*** for Oracle Database 11*g* Release 2 RAC 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 “More..” and click “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. It has the releases from 8.1.7.4 to 12.1.0.2.0. next field is “Platform”, select ”IBM AIX on POWER Systems (64bit)”, it lists the AIX versions from 5.3 to 7.1.Choose one of them. Then click “Search” button. The search result will show the certification status in

© Copyright 2018, IBM Corporation

the link “See Certification Details for Notes and Support information” in the search result page. More information are available about the certification status of General Parallel File System (GPFS), PowerHA, Oracle Clusterware and Veritas Storage Foundation for Oracle RAC in AIX platform.

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

To find the latest certification information of IBM Spectrum Scale, PowerHA System Mirror with Oracle Database, refer certification section of My Oracle Support.

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

Oracle Redundant Interconnect usage feature

With Oracle Database 11*g* Release 2, Patch Set One (11.2.0.2), Oracle introduced an integrated Redundant Interconnect Usage feature, which provides a Highly Available (HA) IP network functionality for the Oracle interconnect. Previous to this version, Oracle RAC and Oracle Clusterware depended on AIX and respective OS features to provide a highly available network interface for the Oracle interconnect. With Oracle Database 11*g* Release 2 Patch Set One, customers have the choice to either continue to use the AIX provided HA network interface, or to use Oracle’s integrated Redundant Interconnect Usage feature, which will provide full high availability for an Oracle RAC Database and Oracle ASM of version 11.2.0.2 or higher.

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 as with pre-11.2.0.2 releases during upgrades. This will allow the Redundant Interconnect Usage to allocate an (HA)IP 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 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 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 RAC database. Under certain circumstances, virtualization solutions and other techniques provided by the hardware or OS vendor may mitigate these

© Copyright 2018, IBM Corporation

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 RAC 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 11*g* Release 2 RAC. For more information on Oracle Clusterware, visit:
http://www.oracle.com/technetwork/products/clusterware/ocluster11gover-1988195.html

For AIX 6.1, to resolve Oracle Bug 7006789 “VIP does not failover when public network cables are removed on one node”, follow this configuration recommendation.

> Oracle RAC clusterware has strict timeout requirements for VIP address failover in case of a public network failure. When DNS servers are unreachable due to a public network failure, DNS name resolution calls such as getaddrinfo may hang for the default AIX query timeout duration of 5 minutes. Name resolution calls made by Oracle processes can thus delay the VIP failover. To reduce such delays, the DNS query timeout can be reduced to 1 minute, by adding the following options line in /etc/resolv.conf for all RAC cluster nodes:
>
> "options timeout:1"
>
> No reboot is necessary to activate this change.
>
> If you need even faster VIP failover the timeout can be further reduced to a value of 0; provided your network infrastructure (network and DNS servers) has the speed to serve name queries within a few (5-6) seconds. If you use a value of 0 for timeout and your DNS or network is slow to respond, DNS name lookups will start to fail prematurely.

IBM PowerHA (formerly High Availability Cluster Multi-Processing)

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

For Oracle database 11*g* Release 2 RAC, raw logical volumes are only supported when performing an upgrade installation from Oracle database 10*g* R2 RAC or 11*g*R1. Raw logical volumes are not available when performing a fresh install of Oracle Database 11*g* Release 2 RAC.

Where customers desire raw logical volumes for their Clusterware files, database files or recovery files, PowerHA will be required. The certification details for the different versions of PowerHA on Oracle Database 11*g*R2 RAC are:

© Copyright 2018, IBM Corporation

PowerHA SystemMirror 7.1

The software requirement for PowerHA 7.1 and Oracle Database 11*g* R2 RAC is as follows,

- AIX 6.1 TL06 SP03 or newer, must include APAR IZ88711, IV08558 and IV07567
- AIX 7.1 TL01 SP02 or newer, much include APAR IZ89165,IV01423 and IV08187
- AIX7.2 TL0 SP01 or later
- VIOS 2.2.0.13 FP24 SP03 or newer , must include APAR IV08703
- PowerHA 7.1 SP3 or newer, must include APAR IV00109
- RSCT 3.1.0.3 or newer for AIX 6.1 and 7.1
- Oracle Database 11*g* R2 RAC version 11.1.0.1 or newer, must install Oracle rootpre patch mentioned in the document ID 1384060.1

PowerHA SystemMirror 6.1

Oracle Database 11*g* Release 2 RAC version 11.2.0.1 is certified with PowerHA SystemMirror 6.1 with following software stack

- AIX 5.3 TL10 SP05 or later, must include fix for APAR IZ89268
- AIX 6.1 TL04 SP07 or later, must include fix for APAR IZ87768
- AIX 7.1 TL01 SP02 or newer, much include APAR IZ89165,IV01423 and IV08187
- AIX 7.2 TL0 SP01 or later
- VIOS 2.2.0.10 FP24 SP01
- PowerHA 6.1 SP03 or later
- RSCT 2.4.12.0 or later for AIX 5.3
- RSCT 2.5.4.0 or later for AIX 6.1

RSCT 3.1.0.0 version for AIX 7.1 HACMP 5.4.1

- **Certified on AIX 5.3 and AIX 6.1.**

The certification notes for HACMP 5.4.1 and **AIX 5.3** list the following requirements.

- AIX 5.3 TL 6 or later, specifically bos.rte.lvm 5.3.0.60. 11*g* Release 2 requires TL09 SP01 or later.
- HACMP 5.4.1, available in media or APAR IZ02620.
- RSCT 2.4.5 for AIX 5.3
- Reliable Scalable Cluster Technology (RSCT) 2.4.7.3, rsct.basic.rte and ifix for APAR IZ01838. This APAR is incorporated into 2.4.8.1.
- Patch 6718715 for rootpre.sh is required with HACMP during a fresh install of Oracle RAC Clusterware or when upgrading to Oracle Database 11*g* Release 2.

The certification notes for HACMP 5.4.1 and **AIX 6.1** list the following requirements.

- AIX 6.1 TL02 SP1 or later.
- HACMP 5.4.1 SP01, available in media or APAR IZ02620.
- RSCT 2.5.2 for AIX 6.1
- Patch 6718715 for rootpre.sh is required with PowerHA during a fresh install of Oracle RAC Clusterware or when upgrading to Oracle Database 11*g* Release 2.
- To resolve Oracle Bug 7006789 “VIP does not failover when public when public network cables are removed on one node”, follow this configuration recommendation.

© Copyright 2018, IBM Corporation

Oracle RAC clusterware has strict timeout requirements for VIP address failover in case of a public network failure. When DNS servers are unreachable due to a public network failure, DNS name resolution calls such as getaddrinfo may hang for the default AIX query timeout duration of 5 minutes. Name resolution calls made by Oracle processes can thus delay the VIP failover. To reduce such delays, the DNS query timeout can be reduced to 1 minute, by adding the following options line in /etc/resolv.conf for all RAC cluster nodes:

"options timeout:1"

No reboot is necessary to activate this change.

If you need even faster VIP failover the timeout can be further reduced to a value of 0; provided your network infrastructure (network and DNS servers) has the speed to serve name queries within a few (5-6) seconds. If you use a value of 0 for timeout and your DNS or network is slow to respond, DNS name lookups will start to fail prematurely.

For more information, refer the My Oracle Support note **404474.1** "Status of Certification of Oracle Clusterware with PowerHA" and read the content in the topic "Using PowerHA and Oracle Database 11*g* Release 2 RAC".

PowerHA version compatibility matrix with AIX versions can be found in the link http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD101347

Multi-node Disk Heartbeat (MNDHB)

In order to reduce the likelihood of an unnecessary cluster partition failover, multiple IP networks and at least one non-IP network are recommended for heart beating. The most convenient way of configuring non-IP networks is to use disk heart beating, as it removes the problems of distance associated with RS-232 networks. Refer to these documents for additional information on MNDHB:

- PowerHA Best Practices, https://www.redbooks.ibm.com/redpapers/pdfs/redp5117.pdf
- My Oracle Support note 404474.1, "Status and Certification of Oracle Clusterware with HACMP 5.3 & 5.4" for detailed MNDHB configuration instructions.

The PowerHA web page is at https://www.ibm.com/us-en/marketplace/powerha.

IBM Spectrum Scale (formerly General Parallel File System)

If using a file system for your Oracle Database 11*g* Release 2 RAC data files (as opposed to raw logical volumes or ASM), you'll need to use a cluster file system (CFS). A CFS 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 11*g* RAC data files. The IBM General Purpose File System (GPFS) is an Oracle RAC 11*g* certified CFS.

© Copyright 2018, IBM Corporation

GPFS 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/GPFS 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).

IBM Spectrum Scale 4.2 is certified to Oracle 12c R1 with AIX 7.1 and 7.2

The following stacks are recommended

- IBM Spectrum Scale 4.2.0.4
- Oracle Database 12c R1 RAC version 12.1.0.2.0
- AIX7.2 TL2 SP01
- AIX7.1 TL05 SP01
- VIOS 2.2.4

IBM Spectrum Scale 4.1 is certified with Oracle 11*g*R2 with AIX 6.1 and AIX 7.1.

- IBM Spectrum Scale 4.1
- Oracle Database 11*g*R2 RAC version 11.2.0.4
- AIX 7.2 TL0 SP01, AIX7.1 TL02 SP04
- AIX6.1 TL08 SP04
- VIOS 2.2.1.7 or later

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

The GPFS web site is https://www.ibm.com/support/knowledgecenter/en/SSFKCN/gpfs_welcome.html

Make sure to review the ***current GPFS advisories*** in the GPFS FAQ available from the GPFS web site.

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

Virtualization for Oracle Database 11*g* Release 2 RAC

As mentioned in the Current Certifications sub-section (see above) for Oracle Database 11*g* Release 2 RAC; LPARs, Micro-Partitioning and Dynamic LPAR are supported in Oracle Database 11*g* Release 2 RAC environments. LPARs, Micro-Partitioning, Dynamic LPAR and VIOS are part of 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 11*g* Release 2 RAC are:

- **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.

© Copyright 2018, IBM Corporation

- **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 has to 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.

- **Dynamic LPARs** (DLPAR) are a shared or dedicated LPAR to which changes can be made to the amount of processors, memory and virtual or physical adapters without requiring a reboot.
  - 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 11*g* Release 2 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 11*g* Release 2 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 11*g* Release 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 associated raw disk based Voting and OCR. This may be done for both ASM and GPFS 3.2. Customer must confirm IBM support of the configuration, and install any required AIX and Oracle updates before using.

VIOS 2.1 and VIOS 2.2 are certified on Oracle Database 11*g* Release 2 RAC with AIX 6.1, and 7.1 respectively.

VIOS2.2:

- VIOS 2.20.10.FP24 with APAR IZ86878
- Oracle Database 11*g* Release 2 version 11.2.0.1 or later
- AIX 7.1 SP01 with fixes for APAR IZ87564 and IZ87216 or later

VIOS 2.1:

© Copyright 2018, IBM Corporation

- VIOS 2.1.2.10.FP22
- Oracle Database 11*g* Release 2 RAC 11.2.0.1 or later
- AIX 6.1 TL03 SP1 or later

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

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

- **Live Partition Mobility (LPM)** allows you to migrate running AIX and Linux LPARs and their hosted applications from one physical server to another without disrupting infrastructure services. The migration transfers the entire partition state, including the processor context, memory, attached virtual devices, and connected users.

LPM is certified to work with Oracle Database 11*g* Release 2 RAC in the following environment,

- Oracle Database 11*g* Release 2 RAC version 11.2.0.1 with ASM and up
- Oracle Database 11*g* Release 2 RAC version 11.2.0.3 and up with GPFS
- AIX 6.1 TL06 SP09
- VIOS 2.2.0.12 FP24 SP02
- HMC V7R7.3 SP01
- Firmware AM730-049

For AIX 7.1, the following or higher is certified

- Oracle Database 11*g* Release 2 RAC version 11.2.0.1 with ASM and up
- Oracle Database 11*g* Release 2 RAC version 11.2.0.3 and up with GPFS
- AIX 7.1 TL0 SP2
- VIO Server 2.2.0.11-FP-24 SP-01 plus Interim fix IZ91190
- HMC V7R7.2.0.1
- Firmware 01EM350_103

AIX 7.2 is certified to use with Oracle Database 11g Release 2 (11.2.0.4) for LPM.

http://www.redbooks.ibm.com/abstracts/sg247460.html?Open

- **Active Memory Expansion (AME)**

AME is a new technology included in IBM POWER7 and up servers with AIX 6.1, AIX 7.1 and AIX 7.2 for expanding a system's effective memory capacity. Active Memory Expansion employs memory compression technology to transparently compress in-memory data, allowing more data to be placed into memory and thus expanding the memory capacity of POWER7 systems. Utilizing Active Memory Expansion can improve system utilization and increase a system's throughput.

AME is certified to use with Oracle Database 11*g* Release 2 RAC with the following configuration,

- Oracle Database 11*g* Release 2 RAC version 11.2.0.2

© Copyright 2018, IBM Corporation

- AIX 6.1 TL06 SP05
- Firmware AM730_035

When AME is enabled prior to AIX 7.1 TL02 and AIX 6.1 TL08, the value for the AIX parameter 'vmm_klock_mode' would be 2. However, beginning with AIX 7.1 TL02, 'vmm_klock_mode' will be set to 3 by default when AME is enabled. This is not suitable for the stability of Oracle RAC node or Grid deployed in the AIX logical partitions. Therefore, when AME is enabled on AIX 7.1 TL02, Oracle RAC or Grid users should set vmm_klock_mode to 2 as in the following example: 'vmo -r -o vmm_klock_mode=2'.

Use the following AIX command to display the value: 'vmo -L vmm_klock_mode'.

No changes are required for Oracle Single Instance partitions or if using AIX 6.1 TL08. For more details on the meanings of each value, see the output from the command 'vmo -h vmm_klock_mode'

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

Virtualization features are powerful, yet their implementation can get quite complex in Oracle Database 11*g* Release 2 RAC. If you have any questions about implementing virtualization with Oracle database 11*g* Release 2 RAC, contact the IBM Oracle International Competency Center at ibmoracl@us.ibm.com. For additional, detailed PowerVM information, see these documents and websites:

- 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 PowerVM Best Practices, http://www.redbooks.ibm.com/redbooks/pdfs/sg248062.pdf

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 11*g* Release 2 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 5.3, AIX 6.1, 7.1 and 7.2.

1. 100 Mbps and 1 Gigabit
2. 10 Gigabit Ethernet
3. 40 Gigabit Ethernet
4. IP over InfiniBand (IPoIB)

© Copyright 2018, IBM Corporation

5. RDSv3

IPoIB and RDSv3 are supported for Oracle Database 11g Release 2 version 11.2.0.4 with AIX 6.1 TL9 SP1 and above or AIX 7.1 TL3 SP1 and above with Oracle patch 12909465

For information about implementing IPoIB in AIX, refer the following documents:

- Implementing InfiniBand on IBM System p (Redbook), www.redbooks.ibm.com/abstracts/sg247351.html?Open
- Configuring InfiniBand for AIX (article), http://www.ibm.com/developerworks/aix/library/au-infiniband/?S_TACT=105AGY06&

The Oracle Redundant Interconnect (HAIP) feature introduced in 11.2.0.2 is not compatible with InfiniBand on AIX at this time. Customers can continue to use 11*g*R1 with IB on AIX

Integrated Virtual Ethernet

The Integrated Virtual Ethernet (IVE) is a collection of Power Systems hardware, software and hypervisor features that provides integrated high-speed Ethernet adapter ports with virtualization capabilities. The IVE appears in AIX system commands as the Host Ethernet Adapter (HEA) or Local HEA (LHEA). The IVE offers:

- IVE Adapter Ethernet port options:
  - Two 1 Gbps ports or
  - Four 1 Gbps ports or
  - Two 10 Gbps ports
- External network connectivity for LPARs using dedicated ports without the need of a VIOS.
- Industry standard hardware acceleration, loaded with flexible configuration possibilities.
- The speed and performance of the GX+ bus, faster than PCI Express x16.

For additional information about IVE/HEA:

- Integrated Virtual Ethernet Adapter, Technical Overview and Introduction (Redbook), http://www.redbooks.ibm.com/redpapers/pdfs/redp4340.pdf

Ensure your IVE/HEA is deployed consistent with the instructions in My Oracle Support note **282036.1**, “Minimum Software Versions and Patches Required to Support Oracle Products on IBM pSeries”.

For configuring IVE/HEA on Power Systems, refer to the IBM technical document “Setting up IBM POWER6 10 Gigabit Ethernet ports and AIX 6.1 etherchannel for Oracle RAC private interconnectivity” , Document ID : WP101734

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101734

Some additional consideration as part of the 10 GigE setup as follows,

- LACP timeout: Use the “long timeout” switch setting for the amount of time to wait before sending LACPDUs.
- 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.

© Copyright 2018, IBM Corporation

- 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”.
- 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 or 10GE HEA 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. Customer 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 release

Refer the IBM white paper “*Oracle 12.1.x and 11.2.0.4 Database Performance Considerations with AIX on POWER8*” for Oracle Database bugs at https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102608.

Consider applying fixes for other Oracle Database bugs on lower version of Oracle Database 11g Release 2 as indicated below,

Bug 6782569: “11*g* install AIX VIPS won’t start”. This fix is required when using the IBM Logical Host Ethernet Adapter (LHEA) for the Oracle Public or VIP interfaces. This fix is required when configuring IBM Logical Host Ethernet Adapter on the IVE/HEA.

Bug 6319685: LGWR posts do not scale on some platforms. LGWR posts do not scale well giving longer than needed latency on “log file sync” waits (which involve posting LGWR).

In order to use AIX VIPA, users need to apply patch 10203061, this is also contains in 11.2.0.2.1 patch set.

Apply the Oracle fix for the bug 9457492: OCR DISKGROUP WAS DISMOUNT WHILE I/O ERROR OCCURED ON ONE NODE

Mutex Wait:Bug 10411618: ADD DIFFERENT WAIT SCHEMES FOR MUTEX WAITS - Note: 11.2.0.2.2 PSU breaks this patch and additional patch (12431716) is required. See Master Note: WAITEVENT: "library cache: mutex X" [ID 727400.1]

Bug 1062676.1 - ORAAGENT or ORAROOTAGENT High Resource (CPU, Memory etc) Usage.

In Oracle Database version 11.2.0.3 RAC shows high CPU usage in “ologgerd” and heavy write to $GRID_HOME/crf/db - 9+MB/s. Applying the patch 11.2.0.3.1 GI+PSU resolves this issue.

Grant DBA to user can cause latch row cache objects: Granting “DBA” role to all application users who make large number of concurrent transactions resulting a huge wait with an event “latch: row cache objects” which affects performance of the database. To resolve this performance 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).

© Copyright 2018, IBM Corporation

Contact Oracle support for the interim/one off fix.

Tuning tips

Tuning Oracle Database 11*g* Release 2 and Oracle Database 11*g* Release 2 RAC for AIX is documented in detail in many documents. 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 11*g* Release 2 that provides the database administrator with a simple storage management interface that is consistent across all server and storage platforms.

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

In Oracle Database 11*g* Release 2 , 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 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.

For more information on ASM new features , refer to the Oracle document "*Oracle Database New Features Guide 11g Release 2 (11.2)"*

https://docs.oracle.com/cd/E11882_01/server.112/e41360/toc.htm

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 5.3 TL07 or later, AIX 6.1, 7.1 and 7.2.

The two commands "lkdev" and "rendev" are introduced to effectivly 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 at a later time. 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" 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 output from AIX disk commands indicating that a disk is being used by Oracle

© Copyright 2018, IBM Corporation

For more details of using these commands, refer the “My Oracle Support” note **1445870.1**.

For information about ASM, see Oracle’s ASM web page:

http://docs.oracle.com/cd/E11882_01/server.112/e18951/asmcon.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 11*g* often requires multiple server and user processes running at the same time. Therefore Oracle Database 11*g* 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 “Tuning Resources” section of this document for many white papers. The IBM white paper “*Tuning IBM AIX 5L for an Oracle Database* whitepaper” 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.

In AIX 5.3 AIO is disabled by default. However, in AIX 6.1 and above AIO is enabled by default. When upgrading to AIX 6.1, the AIO setting will not be changed.

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.

Concurrent I/O and direct I/O

Refer to “Tuning Resources” section of this document has many useful links for whitepapers such as file system I/O tuning information in the whitepapers *Tuning IBM AIX 5L for an Oracle Database* , the *Oracle Architecture and Tuning on AIX,* and the *Direct I/O or Concurrent I/O on AIX 5L* My Oracle Support note [272520.1]. For Oracle Database 11*g*, the database defaults to asynchronous I/O (AIO) enabled and concurrent I/O (CIO) disabled. In general, a good starting point is to set the filesystemio_options=setall, in your init*.ora configuration file. This setting will enable AIO (which is the default) and CIO operation. CIO operation is built upon direct I/O (DIO) with the additional function of inode locking. Note, there may be workloads (eg. sequential reads) where cached I/O performs better than CIO.

When using CIO/DIO, the Oracle setting of DB_FILE_MULTIBLOCK_READ_COUNT (the maximum number of blocks read in one I/O operation during a sequential scan) needs to be considered. Also, the alignment of the database blocksize and the file system block size (agblksize) has to be considered. These two topics are addressed by the documents in the previous paragraph. They are also addressed in the *Oracle 9i & 10g on AIX 5L: Tips and Considerations* whitepaper.

When not using CIO/DIO, look at the suggested settings in the “AIX sequential read ahead” section of the *Oracle Architecture and Tuning on AIX* whitepaper.

From Oracle Database 11g Release 2 version 11.2.0.2 and later, Oracle opens the files using "O_CIOR" which is similar to "O_CIO", but allows subsequent open calls without CIO, so that you no longer need to mount the JFS2 filesystems with mount option "-o cio. " This allows other OS tools and third party tools to access the database files without any issues.

© Copyright 2018, IBM Corporation

File system cache size

In AIX 5.3 the default Virtual Memory Manager (VMM) settings are not optimal for a database environment. In AIX 6.1 the VMM defaults have been changed to be much more suitable for a database workload. When upgrading from AIX 5.3 to AIX 6.1, the VMM settings will not be changed. So, VMM settings need to be changed in the upgraded AIX 6.1 as needed for Oracle database. Refer the following section “Tuning resources” and the document “Tuning IBM AIX 5.3 and AIX 6.1 for Oracle Database (whitepaper) link.

For an Oracle Database workload, we need to ensure the computational pages used for Oracle code, SGA and PGA remain resident in memory. The Oracle Database buffer cache already provides caching of database files. Therefore, the file system cache size should be tuned (using the VMM settings) to favor computational pages over file pages. Check the “Memory and Paging” chapter in the *Oracle Architecture and Tuning on AIX* whitepaper for recommended VMM settings. Note, these settings are a **suggested starting point**. If you have already tuned your system, do not revert to these VMM settings.

RAC IPC

Oracle RAC 11*g* uses the User Datagram Protocol (UDP) for inter process communication (IPC) between nodes. The *Oracle Architecture and Tuning on AIX* whitepaper indicates how to tune the UDP kernel settings in the “Network Tuning” section.

IBM Spectrum Scale (formerly General Parallel File System)

The *IBM Spectrum Scale Concepts, Planning and Installation Guide* https://www-01.ibm.com/support/knowledgecenter/STXKQY_4.1.1/ibmspectrumscale411_content.html contains a “GPFS use with Oracle” chapter that contains Oracle configuration and tuning considerations. 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 for 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 11*g* and Oracle RAC when using system-wide contention scope (AIXTHREAD_SCOPE=S). See the “Tuning Memory” chapter of the *Tuning IBM AIX 5L for an Oracle Database* whitepaper for a detailed description of the AIXTHREAD_SCOPE parameter.

Tuning resources

These are the documents referenced throughout this section.

1. Tuning IBM AIX 5.3 and AIX 6.1 for Oracle Database (whitepaper), http://public.dhe.ibm.com/partnerworld/pub/whitepaper/162b6.pdf
2. Oracle 12.1.x and 11.2.0.4 Database Performance Considerations with AIX on POWER8 https://www-03.ibm.com/support/techdocs/atsmastr.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/atsmastr.nsf/WebIndex/WP101513

© Copyright 2018, IBM Corporation

4. Oracle's USLAHEAP patches available on AIX
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102066
5. Disabling unshared segment aliasing on AIX for Oracle Database 11.2.0.3
http://www-01.ibm.com/support/docview.wss?uid=tss1wp102172
6. Oracle Architecture and Tuning on AIX (whitepaper),
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100883
7. IBM POWER7 AIX and Oracle Database performance considerations
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102171
8. IBM Knowledge Center , IBM Spectrum Scale
9. https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale_welcome.html
10. 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
11. IBM GPFS and Oracle RAC on AIX 5L and IBM pSeries (My Oracle Support note 302806.1),
https://support.oracle.com/CSP/ui/flash.html.
12. IBM Spectrum Scale Frequently Asked Questions
https://www.ibm.com/support/knowledgecenter/en/STXKQY/gpfsclustersfaq.html
13. Managing AIX Devices used by Oracle Automatic Storage Management (ASM)
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102158
14. Direct I/O or Concurrent I/O on AIX 5L (My Oracle Support note 272520.1),
https://support.oracle.com/CSP/ui/flash.html.
15. Oracle 9*i* & 10*g* on IBM AIX5L: Tips and Considerations (whitepaper),
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100556
16. Oracle DB 11g and 12c on IBM Power Systems built with POWER8 technology and IBM Flash System 840
https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102440

These are supplemental tuning resources.

AIX 7.2 Performance Management Guide
ftp://public.dhe.ibm.com/systems/power/docs/aix/72/prftungd_pdf.pdf

- AIX 7.1 differences guide
http://www.redbooks.ibm.com/redbooks/pdfs/sg247910.pdf
- AIX 6.1 Performance Management (AIX documentation),
http://publib.boulder.ibm.com/infocenter/systems/scope/aix/topic/com.ibm.aix.doc/doc/base/performance.htm
- AIX 5L Performance Management Guide (AIX documentation),
http://publib.boulder.ibm.com/infocenter/pseries/v5r3/topic/com.ibm.aix.doc/doc/base/performance.htm
- AIX 5L Practical Performance Tools and Tuning Guide (Redbook),
http://www.redbooks.ibm.com/abstracts/sg246478.html?Open
- Configuring IBM TotalStorage for Oracle OLTP Applications (whitepaper),
https://www.ibm.com/developerworks/community/files/basic/anonymous/api/library/d7e42915-b8b6-4e55-a820-a6d9eea3cde0/document/72ae6bf5-1c52-43d4-83e8-ce9d6c3fbed0/media

© Copyright 2018, IBM Corporation

- Improving Database Performance with AIX Concurrent I/O (whitepaper), https://www.ibm.com/developerworks/community/forums/ajax/download/77777777-0000-0000-0000-000014642230/7102e046-b2f0-4594-87fb-14dc41ff05ce/attachment_14642230_systems_p_os_aix_whitepapers_db_perf_aix.pdf
- Diagnosing Oracle Database Performance on AIX Using IBM NMON and Oracle Statspack Reports (whitepaper), https://www-01.ibm.com/support/docview.wss?uid=tss1wp101720
- RAC Starter Kit and Best Practices, “My Oracle Support” note 811293.1
- Additional information about Oracle10g and 11*g* database compatibility with AIX, refer “My Oracle Support” note 889464.1

© Copyright 2018, IBM Corporation

Summary

This document gathers together the key technical topics that need to be considered in planning or implementing Oracle Database 11*g* R2 or Oracle Database 11*g* R2 RAC with 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.

© Copyright 2018, IBM Corporation

Resources

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

- Oracle Database 11*g* Documentation Library,
  http://www.oracle.com/pls/db112/homepage

- Oracle Clusterware and RAC documentation,
  http://www.oracle.com/pls/db112/portal.portal_db?selected=16&frame=

- Oracle RAC SIG,
  http://www.oracleracsig.org

- Raw Devices and Cluster Filesystems with RAC, My Oracle Support note 183408.1,
  https://support.oracle.com/CSP/ui/flash.html.

- Oracle My Oracle Support note 341507.1: Oracle Products on Linux on IBM POWER,
  https://support.oracle.com/CSP/ui/flash.html.

- Oracle and IBM System Storage
  http://www-03.ibm.com/systems/storage/solutions/compatibility/oracle/index.html

© Copyright 2018, IBM Corporation

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 sales collateral and sizing information for Oracle products on AIX. He also provides support for AIX and Power Systems projects at the IBM Oracle ICC and onsite at Oracle Headquarters in Redwood Shores and Pleasanton, CA.

© Copyright 2018, IBM Corporation

Trademarks and special notices

© Copyright. IBM Corporation 1994-2018. 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, and System p5 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.

© Copyright 2018, IBM Corporation