



# Technical Notes IBM Oracle International Competency Center (ICC)

July 01, 2021

email address: [ibmoracle@us.ibm.com](mailto:ibmoracle@us.ibm.com)

## IBM Power Systems, AIX and Oracle Database 10g & 11.2.0.3 Performance Considerations

### Introduction

This document is intended to provide information, suggestions and website links to assist with IBM Power Systems running in an Oracle Database 10g and 11g environment. The primary focus is AIX with Oracle Database versions 10.2.0.4 and 11gR2 with POWER7 & POWER8 systems. Power Systems concerns can be minimized by following standard practices which applies to all POWER server generations. The document also includes some information related to Oracle Real Application Clusters (RAC).

Table 1, below, shows links to important documents to become familiar with.

Please Read First - Oracle and IBM References		
Oracle Real Application Clusters on IBM AIX – Best practices in memory tuning and configuring for system stability	<a href="http://www.oracle.com/tech/network/database/clusterware/overview/rac-aix-system-stability-131022.pdf">http://www.oracle.com/tech/network/database/clusterware/overview/rac-aix-system-stability-131022.pdf</a>	RAC tuning and configuration Guide
Oracle Real Application Clusters (RAC) and Oracle Clusterware Interconnect Virtual Local Area Networks (VLANs) Deployment Considerations	<a href="http://www.oracle.com/tech/network/database/clusterware/overview/interconnect-vlan-06072012-1657506.pdf">http://www.oracle.com/tech/network/database/clusterware/overview/interconnect-vlan-06072012-1657506.pdf</a>	RAC and VLAN deployment guide



Managing Raw Disks in AIX to use with Oracle Automatic Storage Management (ASM)	(MOS) Doc ID 1445870.1	My Oracle Support (MOS) <a href="https://support.oracle.com">https://support.oracle.com</a> (Registration required)
<b>Must Read for Oracle 11g</b>		
Oracle Database 11.2.0.3 with AIX 7.1 or AIX 6.1 using 1TB Segment Alias	<a href="https://www.ibm.com/support/pages/system/files/inline-files/Oracle_DB_and_Large_Segment_Aliasing_v1.0.pdf">https://www.ibm.com/support/pages/system/files/inline-files/Oracle_DB_and_Large_Segment_Aliasing_v1.0.pdf</a>	Oracle database 11.2.0.3 may run slow with high system time with 1 TB unshared segment aliasing enabled
Oracle DB & RAC 11gR2 on IBM AIX: Tips and Considerations	<a href="https://www.ibm.com/support/pages/node/6355069">https://www.ibm.com/support/pages/node/6355069</a>	11gR2 Planning and implementing

<b>Additional IBM Information</b>		
Review Recent ICC Flashes on Techdocs – the Technical Sales Library	<a href="https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library">https://www.ibm.com/support/pages/ibm-techdocs-technical-sales-library</a>	Select ‘Flashes’ on Technical sales support database option
Large page size	Review <a href="#">Oracle Large Page Usage</a> section for specifications	Larger page size has often provided performance improvements with Oracle
Latest POWER Firmware Release Levels	<a href="https://www14.software.ibm.com/webapp/set2/flrt/mtm#tab_p7">https://www14.software.ibm.com/webapp/set2/flrt/mtm#tab_p7</a>	It is recommended to be at the most current firmware level.
Oracle RAC and DB Support Tools Bundle	<a href="https://community.oracle.com/mosc/discussion/comment/12251628">https://community.oracle.com/mosc/discussion/comment/12251628</a>	Oracle RAC and single instance support tools
Improving ASM Disk Discovery Time Best Practices	MOS Doc ID <a href="#">1608549.1</a> <a href="#">Improving ASM Disk Discovery (ASM_DISKSTRING) Time Best Practices (Doc ID 1608549.1)</a>	Helpful ASM related information

Table 1: Suggested Power and Oracle considerations



## Oracle DB 11gR2 Standard Practices for IBM AIX

This section is a summary of standard practices for standalone Oracle DB 11gR2 instances on IBM AIX 6.1, AIX 7.1 and Power Systems. With the exception of references to Symmetric Multi-threading, quad-threaded mode (SMT4), all of the standard practices apply to POWER6 as well. While the primary focus is on standalone Oracle 11gR2 with filesystem-based storage, standalone Automatic Storage Management (ASM) is also addressed. These standard practices apply to Oracle DB 10.2.0.4 as well, and where they differ from those for Oracle DB 11gR2, patches specific to Oracle DB 10.2.0.4 are noted.

Detailed instructions for installing and configuring Oracle 11gR2 RAC with ASM are provided in [Release Notes: Oracle Real Application Cluster 11g Release 2 \(11gR2\) On IBM AIX](#), available from Oracle. The details of ASM installation and configuration in the document are mostly applicable to standalone ASM as well.

The following pages discuss memory, CPU, I/O, network, and miscellaneous settings. In addition, we list the AIX APARs required for Oracle 11gR2, the Oracle patches for 11gR2 on AIX 6.1 and 7.1, the Oracle patches for 10.2.0.4 and 11gR2, as well as recent suggestions and open issues.

### Memory

Specifications for kernel settings and Oracle large page usage are:

#### Kernel Settings

Kernel settings are listed in Table 2. These are commonly suggested values.

Parameter	Proposed Value	AIX 6.1 Default	AIX 6.1 Restricted	AIX 7.1 Default	AIX 7.1 Restricted
minperm%	3	3	No	3	No
maxperm%	90	90	Yes	90	Yes
maxclient%	90	90	Yes	90	Yes
strict_maxclient	1	1	Yes	1	Yes
strict_maxperm	0	0	Yes	0	Yes
lru_file_repage	0	0	Yes	N/A	N/A



lru_poll_interval	10	10	Yes	10	Yes
minfree	960	960	No	960	No
maxfree	1088	1088	No	1088	No
page_steal_method	1	1	Yes	1	Yes
memory_affinity	1	1	Yes	1	Yes
v_pinshm	0	0	No	0	No
lgpg_regions	0	0	No	0	No
lgpg_size	0	0	No	0	No
maxpin%	90	80	No	90	No
esid_allocator	1*	0	No	1	No

*Table 2 Kernel settings for Oracle*

\* The default value of 1 for esid\_allocator enables terabyte segment aliasing, reducing addressing lookasides. This value may be set to 1 in AIX 6.1 and is suggested for Oracle.

In general, AIX support suggests AIX 7.1 defaults for Oracle.

Three noticeable changes from AIX 6.1 to AIX 7.1 are:

- The elimination of the lru\_file\_repage tunable
- The default value of 90 for maxpin%, increased from 80% in AIX 6.1

### **Oracle Large Page Usage**

Spécifications for Oracle large page usage are:

- AIX 6.1 and 7.1 support three or four page sizes, depending on the hardware: 4 KB (default), 64 KB (medium), 16 MB (large), and 16 GB(huge). All four page sizes are supported by Power Systems from POWER5+ firmware level 240-202 onward.
  - Page sizes 64 KB and 16 MB have been shown to benefit Oracle performance by reducing kernel lookaside processing to resolve virtual to physical addresses. Oracle 11g uses 64 KB pages for dataspace by default.

- LOCK\_SGA = FALSE

- This is the default. This means that the SGA is not pinned in memory.
- AIX performance support generally suggests not to pin SGA.
- Automatic Memory Management (AMM) will use 64 KB pages for SGA if memory is available.
- **This is the suggested value**, since it has been found that 64 KB pages yield nearly the same performance benefit as 16 MB pages and require no special management.
- Oracle 10.2.0.4 with Oracle patch 7226548 will also use 64 KB pages for the SGA.

- LOCK\_SGA = TRUE

- AIX parameters to enable pinned memory and 16 MB large pages:

```
vmo -p -o v_pinshm=1 (allows pinned memory—requires reboot)
vmo -p -o lgpg_size=16777216 -o lgpg_regions=number_of_large_pages
where number_of_large_pages=INT[SGA size -1)/16MB))+1
```

- Capabilities required to allow Oracle to use 16 MB large pages (implement as root):

```
#chuser capabilities=CAP_BYPASS_RAC_VMM,CAP_PROPAGATE
Oracle MOS Reference: How to enable Large Page Feature on AIX-Based Systems \(Doc ID 372157.1\)
```

- AMM can be used with pinned 16 MB pages, provided the formula for calculating the number of large pages (above) is modified to number of large pages=memory\_max\_target+1.
- With Oracle 10.2.0.4, patch 7226548 is also required in order to use 16 MB pinned pages.

- 64 KB Size

- Using 64 KB pages for data, text, and stack regions (applies to Oracle 10.2.0.4 as well)
- 64 KB page size for data, text, and stack regions is useful in environments with a large (for example, 64 KB+) SGA and many online transaction processing (OLTP) users. For smaller Oracle instances, 4 KB is sufficient for data, text, and stack.
- 64 KB page use for data, text, and stack is implemented separately from 64 KB pages for the SGA, and is done by means of an environment variable exported on behalf of the Oracle user. AME by default uses 4k page size.
- \$ export  
LDR\_CNTRL=DATAPSIZE=64K@TEXTPSIZE=64K@STACKPSIZE=64K  
oracle



## CPU

CPU specifications are:

- Symmetric multi-threading (SMT) mode: POWER7 supports SMT4, and this is the AIX default. AIX/Oracle performance support encourages starting with the default.
- Virtual processor folding: This is a feature of Power Systems in which unused virtual processors are taken offline until the demand requires that they be activated. The default is to allow virtual processor folding, and this should not be altered without consulting AIX support. **Processor folding = On**
- Specific to POWER7 SMT4: Certain Oracle 11G parameters, including DB\_WRITER\_PROCESSES and PARALLEL\_MAX\_SERVERS, are partly derived from CPU\_COUNT, and CPU\_COUNT is equal by default to the number of logical CPUs. CPU\_COUNT automatically adjusts to changes in virtual processor count and to SMT mode, up to three times the value on startup. Note that, when migrating from single-threaded platforms to Power Systems, or from POWER5 or POWER6 to POWER7 with SMT4, the value of CPU\_COUNT will also increase, affecting DB\_WRITER\_PROCESSES, PARALLEL\_MAX\_SERVERS, and other dependent parameters. Queries that are sensitive to a degree of parallelism might change behavior as a result of migration to POWER7. We suggest reviewing the PARALLEL\_MAX\_SERVERS parameter after migration, but to set DB\_WRITER\_PROCESSES to default.

## I/O

I/O specifications are:

- If ASM is not used, max interpolicy striping (also known as *pp spreading* or *poor man's striping*, is suggested when logical volumes are created. To get the most benefit from spreading physical partitions across the LUNs, use a small physical partition size, for example, 32 MB or 64 MB.
- Async I/O is used even with Concurrent I/O (CIO)
  - With AIX 6.1 and 7.1, start with the asynchronous I/O defaults. With AIX 6.1, there is a new implementation of AIO. AIO kernel extensions are loaded at system boot (always loaded), AIO servers stay active as long as there are service requests, and the number of AIO servers is dynamically increased or reduced based on demand of the workload. The aio\_server\_inactivity parameter defines after how many seconds idle time an AIO server will exit. AIO tunables are now based on logical CPU count, and hence it is usually not necessary to tune minservers, maxservers, and maxreqs as in the past.
  - Note that in AIX 6.1, there are two tunables for minservers and maxservers, aio\_minservers/aio\_maxservers for legacy threads, and posix\_aio\_minservers/posix\_aio\_maxservers for posix threads. Oracle uses legacy threads.
  - Only increase aio\_maxservers or posix\_aio\_maxservers with `ioo -p -o` (the default is 30 per logical CPU) if `pstat -a | fgrep aio` or `ps -k | fgrep aio` show that you are continually using maxservers.
  - Oracle parameter (init.ora)



- `disk_async_io = TRUE` (default value)
  
- Buffered file I/O on JFS2
  - The default `filesystemio_options=ASYNC`
  - In this case all data spaces, redo log file systems, and control file systems are using the kernel buffers rather than writing directly to disk.
  - In this case, it does not matter whether redo log file systems and control file systems are 512 b or 4 KB block size file systems.
  - Oracle on AIX/Power best performance is, however, usually achieved using CIO (though there are exceptions).
  
- Concurrent I/O (CIO) on JFS2
  - Set the Oracle parameter `filesystemio_options=SETALL`, or mount the filesystems (other than dump devices; may be required in older AIX /Oracle levels) with the CIO option. It is not necessary to both SETALL and mount filesystems with the CIO option, although no harm is done either way. Metalink note: 272520.1 indicate that mounting with CIO is needed, while IBM believes it is not needed. IBM is working with Oracle to fix the metalink note.
  - If using CIO with SETALL, CIO mount or both, you *must* create separate file systems for redo logs and control files (or a single filesystem for both), with an `agblksize` of 512 rather than the default 4 KB.
  - The `ioo` parameters `aio_fsfastpath` and `posix_aio_fsfastpath` accelerate CIO. It is enabled by default in AIX 6.1 and 7.1.
  - AIX 6.1, JFS2 with 11.2.0.2 and higher Bug notice:  
Bug 9310972 - ENHANCEMENT: INTRODUCING O\_CIOR FLAG WHEN OPENING DATAFILES IN AIX 6.1. With AIX 6.1, IBM introduced a new open flag `O_CIOR` which is same as `O_CIO`, but this allows subsequent open calls without CIO. The advantage of this enhancement is that other applications like `cp`, `dd`, `cpio`, `dbv` can access database files in read only mode without having to open them with CIO.

Starting with Oracle 11.2.0.2 when AIX 6.1 is detected, Oracle will use `O_CIOR` option to open a file on JFS2. Therefore you should no longer mount the filesystems with mount option `"-o cio"`.
  
- IBM mount advice for database files:
  - Data files: Use `CIO filesystemio_options=SETALL`, and default `agblksize (4k)`; mount with no options.
  - Redo logs: Create with `agblksize` of 512 and mount with no options. With SETALL, IBM is doing direct I/O for Redo logs.
  - Control files: Create with `agblksize` of 512 and mount with no options. With SETALL, IBM is doing direct I/O for control files.
  - Archive logs: Mount `-o rbrw` . Do not use CIO; use the `jfs2 rbrw` option

- Dumps: Mount -o rbrw
- The mount option noatime, suggested for Oracle 10g binaries, fixed in 11.2.0.2. Please reference “The Recent suggestions and open issues” section below for more information.
- IOO tunables j2\_nBufferPerPagerDevice and j2\_dynamicBufferPreallocation:
  - Do not change these unless there is a high delta in vmstat -v external pager filesystem I/Os blocked with no fsbuf. If this value is high, first increase j2\_dynamicBufferPreallocation from 16 (16k slabs) to 32; monitor. If increasing this does not help, then consider raising the value of j2nBufferPerPagerDevice which is the starting value for dynamic buffer allocation.
  - See help pages for information about these parameters. Do not change AIX 6.1 or 7/1 restricted tunables without the advice from IBM AIX support. In AIX 6.1, j2\_nBufferPerPagerDevice is a restricted tunable, while j2\_dynamicBufferPreallocation is not.
  - Here are some default values for three ioo parameters:
    - j2\_dynamicBufferPreallocation=128
    - numfsbufs=1024 (legacy jfs)
    - maxpgahead=16 (legacy jfs)
- ASM considerations for standalone Oracle 11gR2:
  - For identifying, renaming, and securing ASM raw devices, see [Managing Raw Disks in AIX to use with Oracle Automatic Storage Management \(ASM\)](#).
  - ASM will use asynchronous I/O by default, so filesystemio\_options=ASYNCR (default) is appropriate.
  - For clustered ASM (e.g. RAC) configurations, SCSI reservation must be disabled on all ASM hdisk and hdiskpower devices (e.g. reserve\_policy=no\_reserve). The standalone use of ASM, hdisks and hdiskpower devices does not need to have SCSI) reservation disabled.
  - The following initialization parameters need to be adjusted for ASM:
    - Add 16 to the value of processes
    - Add an additional 600 KB to the value of large pool size
    - Add to shared pool size the aggregate of the values returned by these queries:
      - SELECT SUM(bytes)/(1024\*1024\*1024) FROM V\$DATAFILE;
      - SELECT SUM(bytes)/(1024\*1024\*1024) FROM V\$LOGFILE a, V\$LOG b WHERE a.group#=b.group#;
      - SELECT SUM(bytes)/(1024\*1024\*1024) FROM V\$TEMPFILE WHERE status='ONLINE';
    - For disk groups using external redundancy, every 100 GB of space needs 1 MB of extra shared pool, plus 2 MB
    - For disk groups using normal redundancy, every 50 GB of space needs 1 MB of extra shared pool, plus 4 MB



- For disk groups using high redundancy, every 33 GB of space needs 1 MB of extra shared pool, plus 6 MB
- Source:  
[http://docs.oracle.com/cd/E18283\\_01/server.112/e16102/asminst.htm#CHDBBIBF](http://docs.oracle.com/cd/E18283_01/server.112/e16102/asminst.htm#CHDBBIBF)

## Network

This section outlines the minimum values applicable to network configurations.

### Kernel configurations

These values are generally suggested for Oracle, and can be considered as starting points (pls note all udp settings are specific for RAC):

- `sb_max`  $\geq$  1MB (1048576) and must be greater than maximum tpc or udp send or recvspace (if you are using RAC and very large `udp_recvspace`, you might need to increase `sb_max`)
- `tcp_sendspace` = 262144
- `tcp_recvspace` = 262144
- `udp_sendspace` = `db_block_size` \* `db_file_multiblock_read_count`
- `udp_recvspace` = 10 \* (`udp_sendspace`)
- `rfc1323` = 1 (see [Recent suggestions and open issues](#))
- Ephemerals (non-defaults suggested for a large number of connecting hosts or a high degree of parallel query; also to avoid install-time warnings)
  - `tcp_ephemeral_low`=9000
  - `tcp_ephemeral_high`=65500
  - `udp_ephemeral_low`=9000
  - `udp_ephemeral_high`=65500

Jumbo frames are Ethernet frames larger than the standard maximum transmission unit (MTU) size of 1500 bytes. They can be up to 9000 bytes. They are used to reduce the number of frames to transmit a given volume of network traffic, but they only work if enabled on every *hop* in the network infrastructure. Jumbo frames help to reduce network and CPU overheads.

## Miscellaneous Specifications

**AIXTHREAD\_SCOPE=S (set in Oracle profile)**

**ulimits (smit chuser or edit /etc/security/limits to create a stanza for Oracle)**

- -1 (unlimited) for everything except core

**Maximum number of PROCESSES allowed per user (smit chgsys)**

- maxuproc >= 2048; 16 KB is a commonly suggested value for Oracle Environment variables
- LDR\_CNTRL=DATAPSIZE=64K@TEXTFSIZE=64K@STACKSIZE=64K  
\$ORACLE\_HOME/bin/oracle. (preferred approach)
- There is an option to ledit the Oracle binaries so they use 64 KB pages directly. Note that whenever a patch is applied or an Oracle relink is performed, this ledit will have to be performed again.
- # ledit -btextpsize=64k -bdatapsize=64k -bstacksize=64k  
\$ORACLE\_HOME/bin/oracle.
- AME by default will use 4k page size instead of 64k page size. The use of 4KB page size may impact performance. AME certification on 11gR2 versions was done with 4kpage size.

### Disk and Adapter Resources

- Hdisk – lsattr –El hdisk<>
  - Queue depth might vary among default of 8, 16, 20, and 24, depending on the storage vendor. A queue depth of 2 on SAN devices usually indicates a driver mismatch, but is the default for some Hitachi HDS on AIX and should be increased to 8 as a starting point. Queue wait and queue overflow detected through iostat –DI might indicate a need to increase queue depth.
  - max\_transfer might need to be adjusted upward depending on the largest I/O requested by Oracle
    - A typical starting point for Oracle on AIX is 0x100000
  - As of AIX 5.3, the optimal setting for LTG size is dynamically calculated during the varyonvg process and does not need to be manually set. The varyonvg '-M' parameter should not be used as it will over-ride the dynamically calculated LTG size. It is recommended that all hdisks within a given VG have the same 'max\_transfer' (and other attribute) values. In order to change hdisk attribute values, any associated filesystems should be unmounted and the VG varied off.
- FC Adapter – lsattr –El fcs<> and fcstat -e
  - max\_xfer\_size should be increased from default 1MB to 2MB. The default adapter DMA memory size is 16 MB which increases to 128 MB when a non default max\_xfer\_size is used. Larger DMA size can be important for performance with many concurrent large block I/Os.
  - num\_cmd\_elems might need to be increased if fcstat -e reports a persistent non-zero value for *No Command Resource Count*. Please verify with your storage provider possible limits and recommended storage best practices before changing num\_cmd\_elems.
  - If fcstat –e reports a persistent, non-zero value for *No DMA Resource Count* contact support.

### Live Partition Mobility (LPM)

- Oracle Miscount timer - default of 30 seconds can be extended to 60 seconds if required

- Use of Dedicated Adapter recommended instead of SEA
- Network bandwidth if minimal can cause restrictions
- LPM with RAC Reference on IBM Techdocs - WP102094
- IBM PowerVM Virtualization Introduction and Configuration SG 24-7940

## Performance

### Potential Oracle Database Performance Issue – Oracle Table ‘x\$ksmsp’

**Issue:**

“If customers are using any query that accesses the table: 'x\$ksmsp' they should disable such queries..

Oracle provides the use of x\$ksmsp to allow customers to give a listing of the RAM heap to see how; free space is allocated within the shared pool, the sizes of available chunks on the freelist for the shared pool and RAM.

Unfortunately, this leads to a myriad of issues including; system hangs, heap issues (locking) etc in production systems and selecting from x\$ksmsp on a production system is to be avoided.

This basically does the following for each subheap in the pool, it will:

- 1) Grab the shared pool latch for that subheap
- 2) Walk through ALL chunks of memory in that subheap calling back to ksmsp to extract information about the chunk of memory
- 3) Release the shared pool latch for that subheap.

Even on a minimal sized shared pool this means you are holding the shared pool latch for a significant amount of time which then blocks anyone needing to perform any shared pool memory operation that requires that latch.”

**Precautions:**

- 1) It is NOT recommended to run queries on X\$KSMSP when the database instance is under load.
- 2) Performance of the database will be impacted, especially currently with very large SGAs.
- 3) Bug 14020215 was filed for ORA-600 errors and unplanned outages running queries directly on X\$KSMSP.
- 4) There is a view, X\$KSMSP\_NWEX, in later versions of 11g that is safer to use for investigation of memory usage.
- 5) Oracle STRONGLY recommends you not run these queries unless specifically requested by Oracle Support to do so.



**Recommend:**

Suggest checking via a crontab entry running a query against a variety of X\$ tables - eliminating any query against the x\$ksmsp resolves such latch contention issues.

## **Standalone Database and RAC Miscellaneous Tips and Reference**

### Troubleshooting Oracle RAC Tips

- Collect and review 'Last Gasp' file  
<http://www.oracle.com/technetwork/database/rac-aix-system-stability-131022.pdf> (pg 7)
- Review CRS Logs for events with timestamps
- OSWatcher - network
- AWR Report
- RDA - collects Oracle logs

### Performance

- Oracle Redo Logs and Control Files – placement on separate file systems
- Apply N-Apply Bundle Patch or 11.2.0.4 (includes Bundle Patch)
- OCSSD Bin – Process & Thread Priority

### Oracle Timers

- Network Heartbeat 1 per/sec with 30 sec default timeout
- CSS Voting Disk Heartbeat 1per/sec 27 sec default timeout
- AIX Driver read/write 30 Second timeout
- Oracle Miscount value – default 30 sec, used 60 sec for LPM large lpar's successfully then changed back to 30 sec
- Oracle CRS OCSSD – default 200 seconds is timeout value
- Note: ORAchK replaces the popular RACcheck tool - (Doc ID 1268927.2 in MOS)

## **AIX fixes for Oracle 10gR2 and 11gR2**

Some of the common AIX fixes for Oracle 10g and 11g follow. The APAR number will be unique to the specific AIX TL level with the most current TL level providing a rollup of the earlier TL & SP levels. The APAR's listed may apply to either Oracle dbase version. If an APAR is not listed for a TL, then the fix is already included in that TL, or that TL is not vulnerable to the problem. The minimum recommended AIX levels currently are: AIX 6.1 TL09 (terminal release) and AIX 7.1 TL03. Please see the appendix for older TL APAR information.



<b>AIX 6.1 TL08 &amp; TL09</b>	<b>TL08</b>	<b>APAR</b>	<b>TL09</b>	<b>APAR</b>		
SMT4 PERFORMANCE ENHANCEMENTS	Base	IV10656				
WAITPROC IDLE LOOPING CONSUMES CPU	Base	IV10657				
SRAD LOAD BALANCING ISSUES ON SHARED LPARS	Base	IV10652				
MISCELLANEOUS DISPATCHER/SCHEDULING PERFORMANCE FIXES	Base	IV10651				
64K PAGING TAKING PLACE WHEN AVAILABLE SYSTEM RAM EXISTS	Base	IV23604				
SYSTEM CRASH IN AS_FORK_ALIAS IF ESID_ALLOCATOR IS ENABLED	Base	IV26735				
SHLAP64 UNABLE TO PROCESS ORACLE REQUEST LEADING TO KERNEL HANG	Base	IV35888	Base	IV35888		
UDP MULTICAST BUG MAKES ORACLE RAC UNSTABLE, INCLUDING NODE EVICTIONS	Base	IV30297	Base	IV30219		
THREAD_CPUTIME() RETURNS INCORRECT VALUES	SP03	IV38778	Base	IV36225		
DISABLE MULTICAST LOOPBACK FOR MPING symptom: Dropping packet due to direction mismatch. Was expecting r but got s.	Base	IV34046				
<b>AIX 7.1 TL02</b>	<b>TL02</b>	<b>APAR</b>				
WAITPROC IDLE LOOPING CONSUMES CPU	Base	IV11460				
SRAD LOAD BALANCING ISSUES ON SHARED LPARS	Base	IV11991				
MISCELLANEOUS DISPATCHER/SCHEDULING PERFORMANCE FIXES	Base	IV11988				
ADDRESS SPACE LOCK CONTENTION ISSUE	Base	IV11455				
SYSTEM CRASH IN AS_FORK_ALIAS IF ESID_ALLOCATOR IS ENABLED	Base	IV23735				
REDUCE EARLY WORKING STORAGE PAGING	Base	IV26731				
UDP MULTICAST BUG MAKES ORACLE RAC UNSTABLE, INCLUDING NODE EVICTIONS	Base. SP02	IV35893				
THREAD_CPUTIME() RETURNS INCORRECT VALUES	Base	IV30318				
DISABLE MULTICAST LOOPBACK FOR MPING symptom: Dropping packet due to direction mismatch. Was expecting r but got s.	SP03	IV38960				
ADD ABILITY TO REORDER TOC SYMBOLS IN LIMITED CIRCUMSTANCES	Base	IV48898				
VMM RELALIAS LOCKING TUNABLE TO IMPROVE MMAP/UNMAP PERFORMANCE	Base	IV34380				
<b>AIX 7.1 TL03 (con't)</b>	<b>TL03</b>	<b>APAR</b>				
NETWORK PERFORMANCE DEGRADATION ON FC5899 (AUSTIN) ADAPTER APPLIES TO AIX 7100-03	SP01	IV58687				



XMGC NOT TRAVERSING ALL KERNEL HEAPS. APPLIES TO AIX 7100-03 14/04/17 PTF PECHANGE	SP01	IV53587				
UDP MULTICAST: SHORT PACKET FOR SOME LISTENERS. APPLIES TO AIX 7100-03	SP01	IV33047				
DATA PTR INCORRECTLY INCREMENTED IN UDP RECEIVE. APPLIES TO AIX 7100-03	SP01	IV34454				
TATX/LOOKUPS/FILE OPENS APPLIES TO AIX 7100-03	SP01	IV44289				
POSSIBLE STARVATION OF LARGE I/OS UNDER HEAVY WORKLOAD APPLIES TO AIX 7100-03	SP01	IV44347				
PORT/DEVND FC5899 DRIVER HOG CPU WHEN ENTSTAT ON CLOSED	SP01	IV60218 /60052				
UDP SEND PERFORMANCE ENHANCEMENTS APPLIES TO AIX 7100-03	SP01	IV54257				
DISABLE MULTICAST LOOPBACK FOR MPING symptom: Dropping packet due to direction mismatch. Was expecting r but got s.	Base	IV36204				

### Non Critical APAR's

<b>AIX 6.1</b> Non Critical	<b>TL08</b>	<b>APAR</b>	<b>TL09</b>	<b>APAR</b>		
TCP RETRANSMIT PROCESSING IS VERY SLOW	Base	IV14524				
SYSTEM CRASH DUE TO FREED SOCKET WHEN SOCKETPAIR() CALL USED	Base	IV19341				
DISABLE MULTICAST LOOPBACK FOR MPING symptom: Dropping packet due to direction mismatch. Was expecting r but got s.			Base	IV36225		
LPARSTAT -H AND -H WILL NOT SHOW HYPERVISOR STATISTICS APPLIES TO AIX 6100-09			Base	IV53394		
CAT /PROC/SYS/FS/JFS2/MEMORY_USAGE MAY RETURN INVALID ARGUMENT	SP05	IV57785	SP03	IV54359		
UMOUNT FAILS WITH DEVICE BUSY ERROR EVEN WITHOUT ACTIVE PROCESS APPLIES TO AIX 6100-08	Base	IV46203	Base	IV39905		
JAVA WON'T INSTANTIATE IF PROT_NONE USED FOR SHARED MMAP REGION	Base	IZ55237				
SYSTEM HANGS OR CRASHES WHEN APP USES SHARED SYMTAB CAPABILITY. APPLIES TO AIX 6.1-08	Base	IV21846				
LINK FAILS WITH UNDOCUMENTED COMPILER FLAG AND THREAD-LOCAL STG	Base	IV39893				
RUNTIME LINKING FAILED TO BIND THE BSS SYMBOL EXPORTED FROM MAIN	Base	IV39856				
A SPECIAL-PURPOSE LINKER FLAG WORKS INCORRECTLY.			Base	IV42840		
HANG UNDER SOME CIRCUMSTANCES WHEN A C++ DTOR UNLOADS LIBRARIES	Base	IV64454	Base	IV63322		
ATTEMPT FAILED.OPENED STATE DOES NOT RETRY NPIV LOGIN, IF FIRST APPLIES TO AIX 6100-09	Base	IV38695	Base	IV37549		



<b>AIX 7.1 Non Critical</b>	<b>TL02</b>	<b>APAR</b>				
TCP RETRANSMIT PROCESSING IS VERY SLOW	Base	IV15184				
MULTICAST UDP PACKETS NOT DELIVERED TO ALL LISTENERS IN WPAR ENV	Base	IV16417				
SYSTEM CRASH DUE TO FREED SOCKET WHEN SOCKETPAIR() CALL USED	Base	IV19357				
CAT /PROC/SYS/FS/JFS2/MEMORY_USAGE MAY RETURN INVALID ARGUMENT	SP05	IV56930				
UMOUNT FAILS WITH DEVICE BUSY ERROR EVEN WITHOUT ACTIVE PROCESS APPLIES TO AIX 7100-01	Base	IV40102				
SYSTEM HANGS OR CRASHES WHEN APP USES SHARED SYMTAB CAPABILITY. APPLIES TO AIX 7100-01	Base	IV21878				
LOADING 5.3 TLS ENABLED LIBS BY 5.2 APPS CAUSED CORE DUMP IN 32B APPLIES TO AIX 7100-01	Base	IV30320				
LINK FAILS WITH UNDOCUMENTED COMPILER FLAG AND THREAD-LOCAL STG APPLIES TO AIX 7100-01	Base	IV40005				
RUNTIME LINKING FAILED TO BIND THE BSS SYMBOL EXPORTED FROM MAIN APPLIES TO AIX 7100-01	Base	IV39987				
HANG UNDER SOME CIRCUMSTANCES WHEN A C++ DTOR UNLOADS LIBRARIES	Base	IV64248				
ATTEMPT FAILED.OPENED STATE DOES NOT RETRY NPIV LOGIN	Base	IV38879				
<b>AIX 7.1 Non Critical (con't)</b>	<b>TL03</b>	<b>APAR</b>				
A SPECIAL-PURPOSE LINKER FLAG WORKS INCORRECTLY	Base	IV42782				
CAT /PROC/SYS/FS/JFS2/MEMORY_USAGE MAY RETURN INVALID ARGUMENT	SP3	IV55030				
UMOUNT FAILS WITH DEVICE BUSY ERROR EVEN WITHOUT ACTIVE PROCESS	Base	IV40079				
ADD ABILITY TO REORDER TOC SYMBOLS IN LIMITED CIRCUMSTANCES APPLIES TO AIX 7100-03	Base	IV44690				
HANG UNDER SOME CIRCUMSTANCES WHEN A C++ DTOR UNLOADS LIBRARIES.APPLIES TO AIX 7100-03	Base	IV63130				
ATTEMPT FAILED.OPENED STATE DOES NOT RETRY NPIV LOGIN, IF FIRST APPLIES	Base	IV37484				

### Oracle PSU Information on MOS

IBM Oracle ICC Technical Notes  
<https://www.ibm.com/support/pages/node/6355341>  
 © Copyright 2021, IBM Corporation



Quick Reference to Patch Numbers for Database PSU, SPU(CPU), Bundle Patches and Patchsets (Doc ID 1454618.1)

**Recommended N-Apply Bundle #3 for AIX 11.2.0.3. with critical fixes (Patch 20165285)**

Last Updated Jan 21, 2015  
 Release Oracle 11.2.0.3.9  
 Platform IBM AIX on POWER Systems (64-bit)  
 Recommended for Oracle Database 11.2.0.3.0

The following prerequisite patches should be installed in the order shown before installing this patch.

- 1 17540582 DATABASE PATCH SET UPDATE 11.2.0.3.9 (INCLUDES CPUJAN2014)
- 1 19121548 DATABASE PATCH SET UPDATE 11.2.0.3.12 (INCLUDES CPUOCT2014)

**Bugs Resolved by This Patch**

- 11689561 HASH JOIN SLOWER IN 11.2 COMPARED TO 10.2 ON AIX
- 12596494 GENERALLY HIGHER CPU USAGE IN 11.2.0.2 THAN 10.2.0.4
- 12865682 PERFORMANCE ISSUES DUE TO BYTE-SWAPPING
- 13400729 SCATTERED READS SLOWER IN 11.2 THAN 11.1
- 13443029 ONLINE PATCHING / HOT PATCHING: USE SHRSYMTAB TO REDUCE USLA HEAP USAGE
- 13494030 SLOW PROCESS STARTUP WITH ONLINE PATCHES ON AIX
- 13940331 VALUE FOR SETTING THREAD SCHEDULING IS INCORRECT IN SLTSTSPAWN
- 14764540 SGA ATTACH ADDRESSES SHOULD BE AT 1TB BOUNDARY
- 14799449 ORA-7445 [SSKGDS\_SNM()+592] AND INSTANCE TERMINATED BY PMON
- 16036950 N-APPLY PATCH 11.2.0.3 FOR CRITICAL AIX BUGS BUNDLE #1

**Oracle 10gR2 and 11.2.0.3 Bug Fixes AIX Related**

Bug #	Generic/AIX	Description	Version affected	Fixed in	Patch Available	Ref. MOS Note id.
7226548	AIX	enables 64 KB and 16 MB large page support	10.2.0.4	11.1	10.2.0.4 10.2.0.5	282036.1
7568734	AIX	AIX: Sporadic spikes of 'log file sync' on AIX with heavy commit concurrency	10.2.0.4	10.2.0.5	10.2.0.4	34592.1
6784747	AIX	Wasted memory in koh-kghu allocations. Free extents of memory might not get used (Power7 customers should apply this patch, included in 10.2.0.5 patchset)	10.2.0.4	10.2.0.5	10.2.0.4	





10411618	Generic	Add different wait schemes for mutex waits. <b>Note</b> - 11.2.0.2.2 PSU breaks this patch and additional patch (12431716) is required	11.2.0.1 11.2.0.2	12.1 11.2.0.3	11.2.0.1 11.2.0.2	727400.1 10411618.8
13443029 10190759	AIX	Oracle 11.2.0.2 & 11.2.0.3 requires USLA heap patch Bug 13443029 (requires AIX 6.1 TL07 SP2 or AIX 7.1 TL01 SP2 OR Bug 10190759 (disables hot patching)	11.2.0.2 11.2.0.3		11.2.0.2 11.2.0.3	
12740358	Generic	DBMS_UTILITY.FORMAT_CALL_STACK is slower than 10g	11.2.0.2 11.2.0.3	12.1	11.2.0.2 11.2.0.3	1450540.1
9842771	Generic	Wrong SREADTIM & MREADTIM stats in AUX_STAT\$	11.2.0.1 11.2.0.2	11.2.0.3	11.2.0.1 11.2.0.2	
12596494	AIX	Generally higher CPU usage in 11.2.0.2	11.2.0.2	11.2.0.4	11.2.0.2 11.2.0.3	1392633.1
12412983	AIX	AIX 'asynch descriptor resize' wait not necessary on AIX	11.2.0.2 11.2.0.3	12.1	11.2.0.2 11.2.0.3	1318709.1
11689561	AIX	Hash join consumes more Processor in IBM AIX	11.2.0.2 11.2.0.3	11.2.0.4	11.2.0.2 11.2.0.3	
13877328	AIX	Database performance slows down steadily over time	11.2.0.3	See MOS note for work-around		1467807.1
13840529	Generic	Database hang on cache buffer chains and row cache objects	11.2.0.3	See MOS note for work-around		1467807.1
9495594	Generic	Performance degradation running anonymous PL/SQL blocks	11.2.0.1 11.2.0.2	11.2.0.3		1378094.1
13354348	AIX	Unaccounted gap between elapsed to CPU time in 11.2 on AIX	11.2.0.2	11.2.0.3		
13846587	AIX	AWR report shows enlarged "DB Time" values & TKPROF report show enlarged execute CPU times	11.2.0.3	12.1	11.2.0.3	
12980183	Generic	SPM cannot reproduce execution plans	11.2.0.2 11.2.0.3		11.2.0.2 11.2.0.3	
13004894	Generic	Wrong results with SQL TRACE/10046	11.2.0.3		11.2.0.3	1320966.1
12363485	Generic	Wrong result using union all	10.2.0.5 11.2.0.1 11.2.0.2	12.1	10.2.0.5 11.2.0.1 11.2.0.2	1320966.1
10269193	Generic	Wrong results with outer join and CASE expression optimization	11.2.0.1 11.2.0.2	12.1	11.2.0.1 11.2.0.2	



9910484	Generic	SPM uses excessive space in SYSAUX	11.1.0.7 11.2.0.1 11.2.0.2	12.1	11.1.0.7 11.2.0.1 11.2.0.2	
10242967	Generic	ORA-6502 from DBMS_XPLAN.DISPLAY_SQL_PLAN_BASELINE	11.1.0.7 11.2.0.2		11.1.0.7 11.2.0.2	
13377816	Generic	Excessive Memory use by DIAG / DIA0	11.1.0.7 11.2.0.1 11.2.0.2 11.2.0.3		11.1.0.7 11.2.0.1 11.2.0.2 11.2.0.3	
13551402	Generic	High "log file parallel write" and "log file sync" after upgrading to 11.2 with Veritas/Symantec ODM	11.2.0.2 11.2.0.3	12.1 11.2.0.4		
13940331	AIX	AIX: OCSSD threads are not set to the correct priority	11.2.0.2 10.2.0.4 10.2.0.5	12.1	10.2.0.5, 11.2.0.2, 11.2.0.3	1427855.1
16910419	AIX	HAIP FAILED TO START: HAIP failed to start after applied patch 16842566 on top of PSU6	11.2.0.3	11.2.0.3.6		
10627020	AIX	GPFS w/RAC using Asynch I/O - ORA-1115, ORA-1110, ORA-27091 ORA-27072 IBM AIX RISC SYSTEM/6000 ERROR: 4: INTER	11.2.0.2 10.2.0.4 10.2.0.5	12.1	10.2.0.5, 11.2.0.2, 11.2.0.3	
14764540	AIX	SGA ATTACH ADDRESSES SHOULD BE AT 1TB BOUNDARY	11.2.0.3		11.2.0.3	14764540.8 1528081.1
14799449	AIX	INSTANCE TERMINATED BY PMON AFTER SIGSEGV ON [SSKGDS_SNM()+592]	11.2.0.3	12.2	11.2.0.3	
16078109	AIX	PL/SQL APPLICATION RUN SLOWER AFTER INCREASING CPU CORES FROM 26 TO 42	11.2.0.2.0 11.2.0.2.2 11.2.0.3		11.2.0.2.0 11.2.0.2.2 11.2.0.3	
16697958	AIX	OSYSMOND.BIN FAILS ON A PERFSTAT SYSTEM CALL WHEN LPAR HAS A VIRTUAL INTERFACE - CHM process osysmond.bin fails -ERROR: errno=14				Ref: APAR IV24576
18072720	AIX	OSYSMOND.BIN OCCURS EXCESSIVE ERECV/POLL SYSCALL ON 11.2.0.3.6 - osysmond.bin high CPU usage- node was rebooted due to lack of system resource.	11.2.0.3			
14251087		ORA-27072: FILE I/O ERROR, ERROR: 4: INTERRUPTED SYSTEM CALL – fixed in 12.1		10g & 11g		
11902008		SMON CRASH WITH ORA-600 [KCBGCCR_3] DURING TX RECOVERY 10.2.0.4. is fixed in 11.2.0.4.				



18153848		GI_HOME SHOWING 11.2.0.3.8 AFTER SUCCESSFUL INSTALLATION OF PSU 9 - doc error fixed in PSU10				
16005924		ORA-38754: FLASHBACK DATABASE NOT STARTED; REQUIRED REDO LOG IS NOT AVAILABLE	11.2.0.3			
Doc ID 1452790.1		Warning: Heavy Swapping Observed on System after Upgrade to 11.2.0.3				
		<b>Merged 11.2.0.3 Bug Fixes</b>				
12412983	AIX	AIX 'asynch descriptor resize' wait not necessary on AIX				
13400729	AIX	Scattered reads slower in 11.2 than 11.1				
		<b>OR</b>				
14478927		Merge request on top of database PSU 11.2.0.3.3 for bugs above two bus (12412983 & 13400729)				
11689561		Hash join consumes more CPU in IBM AIX				
12596494		Generally Higher CPU Usage in 11.2.0.2 than 10.2.0.4. Reference Merge Patch 13947840				
		<b>OR</b>				
13947840		Merge Request on top of 11.2.0.3.0 for bugs 11689561 12596494				
13877328		Database performance slows over a period of time. This may not apply to you now, but you need to be aware of this. Please review the bug text				
13840529		Database hang on cache buffer chains and row cache objects. Review bug text for workaround				
13004894		10046 trace changes results of gl query with patch 9776940.				
13743357		PSRC Query returns wrong results on first execution after starting instance				
		<b>OR</b>				
17063116		Merge request on top of database PSU 11.2.0.3.7 for bugs 13004894 13743357				
		<b>Merged 11.2.0.4 Bug Fixes</b>				



16901346		BUG 16901346 - OSYMOND PROCESS TAKING ALMOST 5% OF CPU BECAUSE OF HIGH OPEN FDS COUNT			
17733927		BUG 17733927 - CSS CLIENTS TIMEOUT UNDER HEAVY CONNECTIVITY LOADS ON AIX			
		OR			
19202720		Merge patch is: BUG 19202720 - MERGE REQUEST ON TOP OF 11.2.0.4.0 FOR BUGS 16901346 17733927			
		<b>11.2.0.4 &amp; 12c Bug Fixes</b>			
17294810		BUG 17294810 - PUTTING THREAD_CPU TIMER IMPLEMENTATION BACK TO 11.2.0.4			
17501296		After upgrade to 11.2.0.4 unable to delete any rows from table with context index due to error: "PLS-00306	11.2.0.4		
18066615		Bug 18066615 : WRONG RESULTS AFTER UPGRADE TO 11.2.0.4	11.2.0.4		
19181582		DEADLOCK BETWEEN LGON ON 'LGWR WORKER GROUP ORDERING'			
17551261		ORA-904 WHILE RUNNING SELECT QUERY which is in 11.2.0.4. and fixed in terminal release of 12.1.01	11.2.0.4		
14602788		HIGH CPU CONSUMPTION BY QMON PROCESS AFTER UPGRADE TO 11.2.0.4 fixed in 12.1 AQ:11GQ: Q00* PROCESS SPIN IN BOTH INSTANCES			
14275161		ORA-7445 [EVAOPN3()+135] INSERTING FROM PL/SQL AFTER UPGRADE TO 11.2.0.4 fixed in 12.1 - ORA-600 [RWOIRW: CHECK RET VAL] ON CTAS	11.2		
18828868	Patch	EXPLOSION IN AQ BACKGROUND SLAVES RESULTING IN ORA-18/ORA-20	11.2.0.3		
18194267		AIX-11.2.0.4-SIHA:ORAAGENT.BIN MEMORY LEAK	11.2.0.4		
18828868	Patch	EXPLOSION IN AQ BACKGROUND SLAVES RESULTING IN ORA-18/ORA-20	11.2.0.4		
18261183	Patch	GRID INFRASTRUCTURE ORAAGENT.BIN GENERATES <DEFUNCT> THREADS	11.2.0.4		
17802407		ORA-7445 [QERIXFETCHFASTFULLSCAN+8212] AFTER UPGRADE TO 11.2.0.4	11.2.0.4		
13807031	Patch	RMAN DUPLICATE DOES NOT FOLLOW SYMBOLIC LINKS FOR REDOLOG FILES IN 11.2 - fixed in 12.1	11.2.0.4		



## Recent suggestions and open issues

- Hardware prefetch: The POWER7 chip has the ability to recognize streaming memory access patterns with a unit stride or stride N, and initiate the `dcbt`, or `dcbstst` prefetch instructions automatically. It controls how aggressive the hardware will prefetch (that is, how many cache lines will be prefetched for a given reference), and is controlled by the Data Streams Control Register (DSCR). For Oracle and JAVA workloads, due to irregular memory accesses, it has sometimes been found beneficial to disable hardware streaming memory prefetch. Aggressive prefetching can inhibit performance because the prefetched lines might not be referenced. The workaround is to turn off Hardware Prefetch using `dscrctl -b -n -s 1`. The AIX Active System Optimizer (ASO) has the capability to automatically/dynamically tune Hardware Prefetch. (as of AIX 7.1 TL2 SP1 and AIX 6.1 TL8 SP1).
- Large segment aliasing allows each memory segment lookaside buffer to address up to 1 TB of memory, reducing segment lookaside buffer faults and improving memory access. This is enabled by default on AIX 7.1, and is enabled using `vmo -p -o esid_allocator=1` in AIX 6.1. A recently discovered problem with Oracle 11gR2 and large segment aliasing concerning a related tunable, `shm_1tb_unsh_enable` is currently being addressed. This problem is specific to Oracle 11.2.0.3 on AIX 6.1 and AIX 7.1 and is addressed in [Bug 13877328](#).
- `noatime` file system mount option: The Oracle/AIX bug (9548634: IBM/AIX: EXPENSIVE GETCWD() CALLS FROM SNLFCNCDIR() ) which necessitated the mount option has been fixed in 11.2.0.2. The current suggestion is to use the `noatime` mount option for the file system hosting the Oracle binaries only, and only on Oracle database versions 10.2.0.x and 11.1.0.x.
- `rfc1323=1` is a long-standing network tuning suggestion for Oracle on AIX, although the default remains 0 (in global "no" parameter list) in AIX 6.1 and 7.1. A network retransmission latency issue has recently been discovered when `rfc1323=1` on AIX 6.1 TL6 and APAR IV00755 is also present. Our recommendation is to use `rfc1323=1` (Oracle recommendation), but to ensure that "IV13121: TCP RETRANSMIT PROCESSING IS VERY SLOW 12/05/30 PTF PECHANGE" (or equivalent) is applied. `rfc1323` is default for the ISNO values set by the interface device driver, for VETH and 10Gbit it is enabled (and for NFS enabled by default).

### Additional Information

This Technical Note was authored by Wolfgang Tertel, IBM. For more information on this Technical Note, please send your questions to the IBM Oracle International Competency Center at [ibmoracle@us.ibm.com](mailto:ibmoracle@us.ibm.com).

### External IBM Support website website for this document:

IBM Power Systems, AIX and Oracle Database 10g & 11.2.0.3 Performance Considerations  
<https://www.ibm.com/support/pages/node/6355341>

## Appendix

### OCSSD Bin – Process & Thread Priority

Oracle RAC ocssd.bin, cssdagent, cssdmonitor, and osysmond.bin all (and should) have priority (PRI) of '0' and a scheduling policy (SCH) of '-' as well as all threads for ocssd.bin, cssdagent, and cssdmonitor also have a PRI '0'.

The threads are often overlooked and forgotten. The main process maybe correct but the threads associated with those processes also need to run (inherit) at those elevated priorities.

Known 11.2.0.3 OCSSD Bug 13940331

Priority should be 0 and not 60 (see PRI column in data).

Please ensure the ocssd threads are also running at "0" and "—"?

```
ps -p `ps -ef | grep ocssd.bin | grep -v grep |awk '{print $2}'` -mo THREAD
```

Bug 13940331 - AIX: OCSSD threads are not set to the correct priority (Doc ID 13940331.8)

On AIX has found that the threads created for OCSSD were not inheriting the correct priority this may cause evictions when the machine is running under high workload.

The processes show the functions "ntevque->netevpque->poll" in their call stack and may hang.

EXAMPLE:

```
svp0090bdc:root# ps -p `ps -ef | grep ocssd.bin | grep -v grep | awk '{ print $2 }'` -mo THREAD
  USER      PID      PPID      TID ST  CP  PRI  SC      WCHAN      F      TT
BND COMMAND
  oracle  8454212 11141280      - A   0   0  35      *
10240103      -      - /opt/GRID/app/grid/11.2.0.3/bin
  -      -      - 53412031 S   0   0   1
f1000f0a10032f40 8410400      -      -
  -      -      - 67174423 S   0  60   1
f1000f0a10040140 8410400      -      -
  -      -      - 69402727 S   0  60   1
f1000f0a10042340 8410400      -      -
  -      -      - 69730365
S    0  60  1      - 418400      -      -
  -      -      - 70189183 S   0  60   1
f1000f0a10042f40 8410400      -      -
  -      -      - 71958629
Z    0  60  1      - c00001      -      -
  -      -      - 74317843 S   0  60   1
f1000f0a10046e40 8410400      -      -
  -      -      - 76742815 S   0  60   1
f1000f0a10049340 8410400      -      -
```



```

- - - 82182299 S 0 0 1
f1000f0a1004e640 8410400 - - -
- - - 82313403
Z 0 60 1 - - c00001 - - -
- - - 85000239
Z 0 60 1 - - c00001 - - -
- - - 92143809
Z 0 60 1 - - c00001 - - -
- - - 104005671
Z 0 0 1 - - c00001 - - -
- - - 40894877 S 0 60 1
f1000f0a100a7040 8410400 - - -
- - - 41156961 S 0 60 1
f1000f0a100a7440 8410400 - - -
- - - 42992033
S 0 60 1 - - 418400 - - -
- - - 43188529
Z 0 60 1 - - c00001 - - -
- - - 43254061 S 0 60 1
f1000f0a100a9440 8410400 - - -
- - - 43319609
Z 0 60 1 - - c00001 - - -

```

### 11g to 12c Upgrade

Heavy paging observed after upgrading DB 11.2.0.3.4 and AIX 6100-09-01-1341 to DB 12.1.0.1 and AIX 7100-03-03-1415. Both Oracle and IBM Support have arrived at the same conclusion that additional memory is recommended or the workload rebalanced to reduce paging which impacted system performance.

### Out of Standard Support TL APAR's for AIX 6.1 & 7.1

AIX 6.1	TL04	APAR	TL05	APAR	TL06	APAR
BIND64 CORES WITH -BLAZY OPTION ON AIX61	SP9	IZ89302	SP5	IZ89300	SP5	IZ89514
					SP4	IZ88711
LOCKL PERFORMANCE ISSUE					SP6	IZ91983
THERE IS A TIMING ISSUE BETWEEN THE SYNC DAEMON AND A MAPPE			SP5	IZ90483	SP4	IZ94396
SMT4 PERFORMANCE ENHANCEMENTS					SP5	IZ97088
WAITPROC IDLE LOOPING CONSUMES CPU			SP8	IV01111	SP7	IV06197
SRAD LOAD BALANCING ISSUES ON SHARED LPARS			SP8	IV06194	SP7	IV06196
MISCELLANEOUS DISPATCHER/SCHEDULING PERFORMANCE FIXES			SP8	IV11068	SP7	IV10259

IBM Oracle ICC Technical Notes  
<https://www.ibm.com/support/pages/node/6355341>  
 © Copyright 2021, IBM Corporation



ADDRESS SPACE LOCK CONTENTION ISSUE			SP8	IV1086 1	SP7	IV03903
64K PAGING TAKING PLACE WHEN AVAILABLE SYSTEM RAM EXISTS	SP4	IZ71191	Base	IZ7160 3	Base	IZ72650
SYSTEM CRASH IN AS_FORK_ALIAS IF ESID_ALLOCATOR IS ENABLED					SP9	IV23852
REDUCE EARLY WORKING STORAGE PAGING					SP10	IV27655
64K PAGING TAKING PLACE WHEN AVAILABLE SYSTEM RAM EXISTS	Base	IZ71191	Base	IZ7160 3	Base	IZ72650
LSUSER MAY FAIL IF THERE IS A NON-EXISTANT MEMBERNETGROUP	SP10	IZ97619	SP06	IZ9752 5	SP05	IZ96256
SHLAP64 UNABLE TO PROCESS ORACLE REQUEST LEADING TO KERNEL HANG	SP11	IZ86799	SP07	IZ8825 8	SP06	IV04047
THREAD_CPUTIME() RETURNS INCORRECT VALUES					Base	IV35038
<b>AIX 6.1</b>	<b>TL07</b>	<b>APAR</b>				
BIND64 CORES WITH -BLAZY OPTION ON AIX61	Base	IZ88880				
LOCKL PERFORMANCE ISSUE	Base	IZ94177				
THERE IS A TIMING ISSUE BETWEEN THE SYNC DAEMON AND A MAPPE	Base	IZ91991				
SMT4 PERFORMANCE ENHANCEMENTS	Base	IZ96715				
WAITPROC IDLE LOOPING CONSUMES CPU	SP2	IV10172				
SRAD LOAD BALANCING ISSUES ON SHARED LPARS	SP2	IV10173				
MISCELLANEOUS DISPATCHER/SCHEDULING PERFORMANCE FIXES	SP2	IV10292				
ADDRESS SPACE LOCK CONTENTION ISSUE	SP2	IV10606				
SYSTEM CRASH IN AS_FORK_ALIAS IF ESID_ALLOCATOR IS ENABLED	SP5	IV23850				
REDUCE EARLY WORKING STORAGE PAGING	SP6	IV26272				
LSUSER MAY FAIL IF THERE IS A NON-EXISTANT MEMBERNETGROUP	Base	IZ96441				
SHLAP64 UNABLE TO PROCESS ORACLE REQUEST LEADING TO KERNEL HANG	Base	IZ97712				
THREAD_CPUTIME() RETURNS INCORRECT VALUES	Base	IV30712				
VMM RELALIAS LOCKING TUNABLE TO IMPROVE MMAP/UNMAP PERFORMANCE	Base	IV31584				
<b>AIX 7.1</b>	<b>TLO</b>	<b>APAR</b>	<b>TL01</b>	<b>APAR</b>		
BIND64 CORES WITH -BLAZY OPTION ON AIX61	SP3	IZ89165				
LOCKL PERFORMANCE ISSUE			Base	IZ9674 1		
THERE IS A TIMING ISSUE BETWEEN THE SYNC DAEMON AND A MAPPED FILE	SP3	IZ94598	Base	IZ9496 3		
SMT4 PERFORMANCE ENHANCEMENTS	SP3	IZ96658	Base	IZ9721 8		





WAITPROC IDLE LOOPING CONSUMES CPU	SP5	IV09133	SP2	IV1048 4		
SRAD LOAD BALANCING ISSUES ON SHARED LPARS	SP5	IV09114	SP2	IV1080 2		
MISCELLANEOUS DISPATCHER/SCHEDULING PERFORMANCE FIXES	SP5	IV11176	SP2	IV1080 3		
ADDRESS SPACE LOCK CONTENTION ISSUE	SP5	IV11177	SP2	IV1079 1		
64K PAGING TAKING PLACE WHEN AVAILABLE SYSTEM RAM EXISTS	SP3	IZ84452				
SYSTEM CRASH IN AS_FORK_ALIAS IF ESID_ALLOCATOR IS ENABLED	SP7	IV23960	SP5	IV1126 1		
REDUCE EARLY WORKING STORAGE PAGING	SP8	IV26581	SP6	IV2701 4		
64K PAGING TAKING PLACE WHEN AVAILABLE SYSTEM RAM EXISTS	Base	IZ84452				
LSUSER MAY FAIL IF THERE IS A NON-EXISTANT MEMBERNETGROUP	SP03	IZ96373	Base	IZ9674 5		
SHLAP64 UNABLE TO PROCESS ORACLE REQUEST LEADING TO KERNEL HANG	SP06	IZ92569	Base	IZ9792 3		
THREAD_CPUTIME() RETURNS INCORRECT VALUES	Base	IV35196	Base	IV3486 9		

### Non Critical APAR's

<b>AIX 6.1 Non Critical</b>	<b>TL04</b>	<b>APAR</b>	<b>TL05</b>	<b>APAR</b>	<b>TL06</b>	<b>APAR</b>
TCP RETRANSMIT PROCESSING IS VERY SLOW					SP8	IV18483
FILE.ATION OVERFLOW REPORTED IN ERROR WHILE LINKING			SP8	IV1057 6	SP7	IV10539
SYSTEM CRASH IN NETINFO_UNIXDOMNLIST	SP11	IV00634	SP7	IZ9735 3	SP6	IZ97166
MULTICAST UDP PACKETS NOT DELIVERED TO ALL LISTENERS IN WPAR			SP9	IV1956 9	SP8	IZ89950
SYSTEM CRASH DUE TO FREED SOCKET WHEN SOCKETPAIR()					SP9	IV16603
IOCP GETMULTIPLECOMPLETIONSTATUS() NEVER RETURNS	SP6	IZ74508	SP2	IZ7622 7	Base	IZ74932
Unmapped DS3/4/5K LUN may block I/O to other LUNs in Array.					SP5	IZ86957
CRASH IN NETINFO_UNIXDOMNLIST WHILE RUNNING NETSTAT	SP11	IV00634	SP7	IZ9735 3	SP6	IZ97166
<b>AIX 6.1 Non Critical (con't)</b>	<b>TL07</b>	<b>APAR</b>				
TCP RETRANSMIT PROCESSING IS VERY SLOW	SP4	IV14297				
SYSTEM CRASH @GET_NET_INFO+000278	SP4	IV08486				
FILE.ATION OVERFLOW REPORTED IN ERROR WHILE LINKING	SP4	IV09580				
SYSTEM CRASH IN NETINFO_UNIXDOMNLIST	Base	IZ99445				



MULTICAST UDP PACKETS NOT DELIVERED TO ALL LISTENERS IN WPAR ENV	SP4	IV08682				
SYSTEM CRASH DUE TO FREED SOCKET WHEN SOCKETPAIR() CALL USED	SP5	IV21128				
CRASH IN NETINFO_UNIXDOMNLIST WHILE RUNNING NETSTAT	Base	IZ99445				
TOPAS/NMON FAILS WITH ASSERT FAILURE FOR VIRTUAL INTERFACES APPLIES TO AIX 6100-07 (Ref Bug 16697958)	SP06	IV24576				
UMOUNT FAILS WITH DEVICE BUSY ERROR EVEN WITHOUT ACTIVE PROCESS APPLIES TO AIX 6100-08	Base	IV39754				
SHLAP PROCESS FAILS WHEN SHARED SYMBOL TABLE FEATURE IS USED. APPLIES TO AIX 6100-07	Base	IV28319				
LINK FAILS WITH UNDOCUMENTED COMPILER FLAG AND THREAD-LOCAL STG	Base	IV39104				
RUNTIME LINKING FAILED TO BIND THE BSS SYMBOL EXPORTED FROM MAIN	Base	IV33433				
<b>AIX 7.1 Non Critical</b>	<b>TLO</b>	<b>APAR</b>	<b>TL01</b>	<b>APAR</b>		
TCP RETRANSMIT PROCESSING IS VERY SLOW	SP8	IV20595	SP4	IV13121		
SYSTEM CRASH @GET_NET_INFO+000278			SP4	IV08797		
FILE.ATION OVERFLOW REPORTED IN ERROR WHILE LINKING			SP4	IV09541		
SYSTEM CRASH IN NETINFO_UNIXDOMNLIST	SP4	IV06032	Base	IZ99636		
MULTICAST UDP PACKETS NOT DELIVERED TO ALL LISTENERS IN WPAR ENV	SP8	IV26418	SP4	IV16250		
SYSTEM CRASH DUE TO FREED SOCKET WHEN SOCKETPAIR() CALL USED	SP7	IV21131	SP5	IV21235		
Unmapped DS3/4/5K LUN may block I/O to other LUNs in Array.	SP3	IZ86419				
CRASH IN NETINFO_UNIXDOMNLIST WHILE RUNNING NETSTAT	SP4	IV06032	Base	IZ99636		
UMOUNT FAILS WITH DEVICE BUSY ERROR EVEN WITHOUT ACTIVE PROCESS APPLIES TO AIX 7100-01			Base	IV37940		
ORACLE ASM SLOW TO START UP APPLIES TO AIX 7100-01		IV01369	Base			
JAVA WON'T INSTANTIATE IF PROT_NONE USED FOR SHARED MMAP REGION APPLIES TO AIX 7100-01		IV16737	Base	IV38857		
SYSTEM HANGS OR CRASHES WHEN APP USES SHARED SYMTAB CAPABILITY. APPLIES TO AIX 7100-01			Base	IV21116		
SHLAP PROCESS FAILS WHEN SHARED SYMBOL TABLE FEATURE IS USED. APPLIES TO AIX 7100-01			Base	IV28925		
LOADING 5.3 TLS ENABLED LIBS BY 5.2 APPS CAUSED CORE DUMP IN 32B APPLIES TO AIX 7100-01	Base	IV35283	Base	IV35057		



LINK FAILS WITH UNDOCUMENTED COMPILER FLAG AND THREAD-LOCAL STG APPLIES TO AIX 7100-01	Base	IV42032	Base	IV3913 6		
RUNTIME LINKING FAILED TO BIND THE BSS SYMBOL EXPORTED FROM MAIN APPLIES TO AIX 7100-01			Base	IV4141 5		
A SPECIAL-PURPOSE LINKER FLAG WORKS INCORRECTLY. APPLIES TO AIX 7100-01			Base	IV4507 2		
ADD ABILITY TO REORDER TOC SYMBOLS IN LIMITED CIRCUMSTANCES APPLIES TO AIX 7100-01			Base	IV45073		
ATTEMPT FAILED.OPENED STATE DOES NOT RETRY NPIV LOGIN			Base	IV41362		