

ORACLE®



ORACLE[®]

Support Update for Linux on System z

Damian Gallagher

Senior Technical Lead, Oracle on IBM System Z Support

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decision. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

AGENDA

- Oracle RDBMS
- Oracle E-Business Suite
- Virtual memory issues

Linux on Z Oracle RDBMS

- Became available towards the end of 1st Quarter 2011 (as promised 2 years ago 😊)
- Large number of downloads and installations
- Very few problems reported so far
- Installs very cleanly (more later)

Linux on Z Oracle RDBMS

- 10.2.0.5.0 patchset now available, and the January PSU (10.2.0.5.2) also.
- Requires a 10.2.0.2.0 base installation – you install 10.2.0.5.0 over the top (for a fresh installation)
- We strongly recommend that the 10.2.0.5.2 Patch Set Update is applied in conjunction with the 10.2.0.5.0 Patch Set, as there are a number of performance related patches in the PSU.

Patchset release

- There was an unfortunate release timing issue, which we are addressing – this affects both 10.2.0.5.2 and 11.2.0.2.2.
- Base Development build patchsets, Sustaining Engineering build PatchSet Updates, and they can't start until the first team has completed, and it's 6-8 weeks work to build and QA a PSU. Once the first one is complete, later PSU are on schedule.
- Principally the development environment for future releases has been overhauled to provide a more streamlined, automated process to reduce the effort required to produce releases and patchsets.

RPM Checker

- Designed to reduce the installation issues caused by having incomplete/incorrect packages installed – the documentation and installation notes are a little confusing, supplying as they do a list of packages.
- Vendors don't help by changing the way functions are packaged or provided – even in service packs 😊
- Doesn't install any code

RPM Checker

- Available for Database 10.2.0.4 on Red Hat EL5 and SuSE SLES10 SP2+ and SLES11 SP1
- Available for Application Server 10gR2 for Red Hat EL5 and SuSE SLES 10.
- Available for Database 11gR2 on Red Hae EL 5 Update 5, Suse SLES10 SP3+ and SLES11 SP1+.

Installation Notes

- Always do a software only install
- 10.2.0.2.0 + 10.2.0.5.0 + 10.2.0.5.2
- Only then DBCA etc – creating a DB at 10.2.0.2 and then upgrading is pointless.

Installation Notes (10g)

- The 10.2.0.2 OUI does not recognise SLES-11 as a valid OS version. Therefore, you must invoke OUI as:

```
%. /runInstaller -ignoreSysPrereqs
```

- Ignore any link errors
- Apply 10.2.0.5.0 over the top – this does recognise SLES 11

-

Installation notes

- A recent trend with install problems that originates from installing too many RPMs.
- For example: installing your own JDK version (prior to beginning the Oracle Software runInstaller) is not needed on Linux, and is not recommended on Linux. A pre-existing JDK often interferes with the correct JDK that the Linux Oracle Software runInstaller will place and use - usually because only the 64-bit JDK has been installed, and OUI needs the 31-bit.

Installation Notes

- No extra patch is required for the DIRECTIO support for zLinux.

Installation Notes

- You may receive the following message:
- "Checking monitor: must be configured to display at least 256 colors
>>> Could not execute auto check for display colors using command /usr/X11R6/bin/xdpyinfo. Check if the DISPLAY variable is set. Failed <<<<"
- This can be corrected by adding a symbolic link as follows: " ln -s /usr/bin/xdpyinfo /usr/X11R6/bin/xdpyinfo"

Installation Notes

- You may receive the following WARNING - "libxcb: WARNING! Program tries to lock an already locked connection..." this can be safely ignored. You may also receive this WARNING when running dbca, netca.

11Gr2 Installation Notes

- Operating system distributions:
- Red Hat RHEL5 (update 4 or later) - Note RH4 U8 is supported but RH5 recommended
- SUSE SLES10 (SP3)
- SUSE SLES11 (SP1)
-

11Gr2 Installation Notes

Storage requirements:

- Disk space
- 5.5 GB Oracle Grid Infrastructure for a Grid Home for Cluster and ASM Home 1GB /tmp directory.
- 4.6 GB database software.
- 1.7 GB of disk space for a preconfigured database that uses file system storage (optional)
-

11Gr2 Installation Notes

- Memory requirements
- OUI and manuals insist on 4GB
- Since Linux reserves 64Mb, even if you have 4GB, the OUI will still complain requiring an override of the OUI checks
- However, will still install in 2GB or less – we're doing tests to determine the actual limits and will update the MOS notes when complete

11Gr2 Installation Notes

- Library and tools from Oracle Technet
- Kernel driver availability differs:
- ASMLib available for Red Hat 5.6 – refer to My Oracle Support Note: Oracle ASMLIB - Support for RedHat Enterprise Linux 5 AS on IBM zSeries [ID 821652.1]
- Ships with SuSE 11 on DVD1

References

- Patch 8202632 – 10.2.0.5.0 Patch Set
- Patch 10248542 – 10.2.0.5.2 Patch Set Update
- Note 1086769.1 – RPM Checker
- Note 1082253.1 – Requirements for Installing Oracle 10gR2 RDBMS on SUSE SLES 10 on zLinux
- Note 741646.1 – Requirements for Installing Oracle 10gR2 RDBMS on RHEL 5 on zLinux
- Note 1290360.1 – Requirements for Installing Oracle 10gR2 RDBMS on SLES 11 on zLinux

- Note 1306465.1 - Getting Started - 11gR2 Grid Infrastructure, ASM and DB (IBM: Linux on System z)
- Note 1290644.1 - Requirements for Installing Oracle 11gR2 on SLES11 on IBM: Linux on System z (s390x)
- Note 1306889.1 - Requirements for Installing Oracle 11gR2 on RHEL 5 on IBM: Linux on System z (s390x)
- Note 1308859.1 - Requirements for Installing Oracle 11gR2 RDBMS on SLES 10 SP3 on IBM: Linux on System z (s390x)
- Note 169706.1 - Oracle Database on Unix AIX, HP-UX, Linux, Mac OS X, Solaris, Tru64 Unix Operating Systems Installation and Configuration Requirements Quick Reference

Oracle Enterprise Manager

- If you use Grid Control (OMS) 10.2.0.5 and Grid Agent 10.2.0.5 you can monitor a 11.2.0.2 database . See My Oracle Support note 953846.1 Grid Control 10.2.0.5.0 now Supports DB 11.2.0.

Product notes

- 11gR2 on Linux on Z is a full product – all the standard features are included - Real Application Testing, for example.
- Oracle Messaging Gateway fully supported, which will cheer up some people in the room 😊

Oracle Applications News

- E-Business Suite
- 12.1.3 (RUP) now available
- Fusion Middleware certifications: Web Cache, BPEL, Discoverer 10.1.2.3
- Oracle E-Business Intelligence is now available
- Oracle Embedded Data Warehouse now available
- Certified with Database 11.2.0.2
- Complete product stack supported on Linux on Z (12.1.2+)
- Split tier (12.0.4+), i.e the Application Server on distributed platform

References

- Oracle E-Business Suite Installation and Upgrade Notes Release 12 (12.1.2) for IBM: Linux on System z (Doc ID 1116895.1)

Virtual Memory Management

- This is probably the single highest generator of service requests to the Linux on Z support team
- Symptoms range from high CPU utilization through to guest 'hangs' and CRS node reboots
- Linux doesn't handle memory shortage issues very well, or rather, not very politely 😊

Indicators: Load

- **Load Profile**
- From 3,069 Executes per second
- To 2,345


- i.e reduced throughput

Indicators: CPU Wait

- **Top 5 Timed Events (AWR)**

Event	%Total Call Time
db file sequential read	17.5
latch: cache buffers chains	16.5
log file sync	16.0
latch: library cache	11.1
read by other session	9.3
Sum:	70.4
CPU Wait:	29.6

Indicators: Load spike



Indicators: Anomalies

- *** 2009-02-31 14:02:26.487
- Warning: log write time 4500ms, size 4KB
- *** 2009-02-31 14:03:00.576
- Warning: log write time 610ms, size 16KB
- *** 2009-02-31 14:03:44.949
- Warning: log write time 690ms, size 4KB
- *** 2009-02-31 14:04:24.765
- Warning: log write time 780ms, size 4KB
- *** 2009-02-31 14:04:54.096
- Warning: log write time 4050ms, size 4KB

Indicators: memory use

- MemTotal: 5138052 kB
- MemFree: 119928 kB
- Buffers: 186252 kB
- Cached: 2593492 kB
- SwapCached: 94376 kB
- Active: 3662612 kB
- Inactive: 768628 kB
- HighTotal: 0 kB
- HighFree: 0 kB
- LowTotal: 5138052 kB
- LowFree: 119928 kB
- SwapTotal: 4988856 kB
- SwapFree: 4609940 kB
- Dirty: 604 kB
- Writeback: 12 kB
- AnonPages: 1647708 kB
- Mapped: 1786944 kB
- Slab: 167372 kB
- CommitLimit: 7557880 kB
- Committed_AS: 4820648 kB
- PageTables: 372424 kB

Remedies

- Major SGA components not very swap friendly – Oracle will use the LRU pages at the same time as the memory manager wants to swap them.
- With high user process volumes comes higher page table use - so more memory 'fixed'
- Workload spikes can force a machine into memory scavenge mode - kswapd looking through a large allocated memory area for contiguous pages to swap for a new process allocation
- Oom-killer, obviously, is really bad news, as it targets the biggest memory user
- Avoid CMM/CMMA until SLES 11 (which needs 10.2.0.5 min)

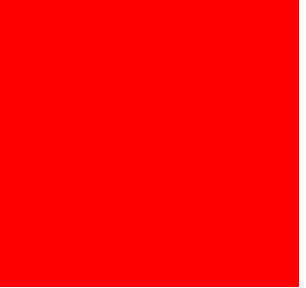
- Hugepages available from 11GR2
- On a 1MB boundary, previously 4KB page
- This reduces page table space, and hence overhead
- And thus reduction in system cycles, especially when under memory pressure.

A bit about Support

- Support specialist organisation for Linux on Z
- ZLUG
- Zseriesoraclesig.org
- Communities



ORA



The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions.

The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

ORACLE®