



Consolidating Oracle to Linux on System z *Selecting a Database*

Oracle Database
Oracle Application Server
Oracle Applications

Denny Dutcavich
Technical Sales Specialist
dutch@us.ibm.com



Linux

© 2008 IBM Corporation



Trademarks

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DBE, e-business logo, ESCO, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/30, VM/ESA, VSE/ESA, Websphere, xSeries, z/OS, zSeries, zVM

The following are trademarks or registered trademarks of other companies

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries
LINUX is a registered trademark of Linux Torvalds
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.
Intel is a registered trademark of Intel Corporation
* All other products may be trademarks or registered trademarks of their respective companies.

NOTES:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput 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 improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

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

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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.

Chart 2

Linux



© 2008 IBM Corporation

Agenda

- Why Consolidate – the premise
- What to look for
- Best Fit
- Selecting an Application Database for Linux on System z
- What you need to know about testing
- Success Factors
- The Process




Why Oracle for Linux on System z

- Value Statement
 - The best TCO characteristics can be obtained from consolidating many servers with low CPU utilization and taking advantage of the virtualization capabilities of z/VM.
 - Lower hardware and software costs
 - Ease of operations
 - Simplified infrastructure
 - On Demand servers
 - However, Linux scales well in an LPAR or with z/VM and may resolve other issues or problems such as availability.
 - The new System z10 EC class machines compete with other technologies (We're fast)
 - Great scalability for consolidation or single large databases
 - Linux provides for a common skill base on all architectures it runs on
- System z differentiators
 - Inherited hardware quality of service
 - Proximity to z/OS
 - Green
 - Multicore Pricing now available for z10





Commercial Workloads under Virtualization



Linux on z/VM

- Very low latency hardware based virtualization
- Large tagged TLB
- Virtualized memory
- No minimum VM/LPAR size
- Small Linux code footprint
- I/O priority mechanism in z/VM
- Linux OS scalability

AIX on PowerVM

- Paravirtualization / hardware assists
- LPAR Size: 0.10 – 64 CPUs
- AIX scalability to 128 CPUs
- High clock speed / SMT
- Minimum entitlement required
- Non-goal mode WLM
- No IO priority capability in VIOS
- No virtualized memory support

z/OS on PR/SM

- Direct I/O virtualization
- Significant I/O offload
- Goal based WLM
- Intelligent Resource Director
- Significant z/OS RAS features
- Centralized vs. distributed model
- DB2 data sharing / Sysplex
- No virtualized memory support
- Relatively large z/OS footprint

Chart 5

© 2008 IBM Corporation



Oracle for Linux on System z Potential Candidates

Medium/High Peak to Average (HW / SW Savings)

- Apps with multiple environments that cumulatively have medium to high peak to average ratio
- Oracle/DB2 databases
- WebSphere
- Tivoli Monitoring
- Sandbox / Training servers

z/OS Affinity (Performance / Security / HW)

- CICS / IMS Gateways
- DB2 connect
- Applications with significant z/OS data affinity such as WAS
- SAP Application Servers with z/OS database
- DB2 on z/OS -zIIP offload candidates
- Communication Controller

Quality of Service (People / Business Impact)

- System z Hardware RAS
- z/VM - very mature Hypervisor
- Fewer critical Linux patches
- Service bureau virtual hosting
- Lotus Domino
- Infrastructure – LDAP, DNS, TSM, TIM/TAM
- Ease of DR

Complete TCO View (People/Other Savings)

- Leverage automation and shared infrastructure
- Speed to market
- Replicated middleware including WebSphere, Oracle, DB2, etc.

Linux Strategy (People / Flexibility)

- Significant industry growth in the Linux applications
- IFLs have strong IT industry adoption
- Common OS across platforms reduces administrative costs

Other (HW / SW Savings)

- MQSeries Queue Manager
- IBI WebFOCUS
- FTP

Chart 6

© 2008 IBM Corporation



Selecting an Application

- Where to start
 - Good planning is essential.
 - “Any man can make a long journey, it takes a smart man to know which direction”. (Confucius)
 - IBM can provide a no-cost application assessment
 - Sizing is critical – Techline (Through IBM rep) can provide the following
 - SURF - SCON Utilization Reduction Facility
 - zRace – Sizing with TCO
 - z/VM Planner for Linux Guest Sizing
 - Work directly with an applications team
 - Databases are driven by an application
 - Work as a team to demonstrate success
 - This can provide the start of a process to decide where new applications are implemented
 - Make sure you have a management sponsor
 - Success stories trump politics



Selecting an Application

- What kind of database to look for
 - Start with databases that are on lower utilized servers
 - Select several databases for a PoC
 - PoC should not be a performance test
 - It should be several databases and/or application servers
 - Look for databases with
 - 1's of Gbytes in size up to about 60 GBytes
 - SGAs up to approximately 1.5 GBytes
 - Work towards larger databases (100s of Gbytes)
 - TBytes not out of scope
 - Continue to assess servers for consolidation



Testing an Application

- Testing on Linux for System z
 - Test workloads selected for Linux on System z
 - This may be a difficult part of process
 - Benchmark testing (e.g. TPC-C) may not provide results needed to make a consolidation technology decision
 - AIM9
 - Benchmark Factory
 - TPM-C and others
 - If you can't test your application with either a production like database or test/dev/etc. consider
 - WAS/Trade
 - Swingbench – warehouse application
 - Loadrunner (getting expensive)

"I never blame myself if I am not hitting. I just blame the bat, and if that keeps up, I change bats."


Yogi Berra



Chart 9

© 2008 IBM Corporation

Design choices affect "benchmark" performance Mixed/WLM/Virtualization



- Working set(s) too large for cache
 - Requires more context switching
- Maximized Processor speed penalized
- "Fast" processor is under-utilized

Replicated

- Cache contains multiple working sets
- Processor speed optimized by cache
- RAS Space is "valued"
- All of "slow" processor is used

Consolidated

Chart 10

© 2008 IBM Corporation



Key Success Factors – z/VM and Linux

- Use **z/VM 5.3**
- Move to a current System z
- Memory is critical
 - Right-size the virtual guest; less is better even with z/VM 5.3
- Monitor resources – excellent tools available
 - IBM Performance Toolkit
 - ESAMON from Velocity Software
 - Omegamon from Tivoli
- Paging and swap space necessary
 - Use memory devices first
 - Use Best Practices for setting up paging space
- Avoid I/O bottlenecks
 - Distribute data in the ESS across arrays
 - Stripe the LVM2 (or use Oracle ASM)
 - Use either FICON or FCP (FCP preferred)



Key Success Factors - Oracle

- File Systems
 - Ext2/3 – EXT3 most common with LVM2
 - ASM – good choice
 - OCFS2 – not aware of z customers using
 - ReiserFS - slow
 - XFS can not do aio and dio at the same time (doc note 414673.1)
- Init.ora Parameters
 - FILESYSTEMS_IO = SETALL
 - Parallel query
 - Perform normal tuning in concert with Linux configuration
- RAC
 - Use for HA instead of instance failover
 - Consider protecting single instance with CRS
 - Integral part of grid if on a grid or consolidation strategy



From Plan to Production – A Process not a Project



Understand Benefits Select Applications Enlist Advocates**	Develop PoC Success Criteria** Define PoC Effort	Identify Servers Size Servers Execute PoC Status Meetings** Validate Results	Initial Production Plan Order of Move How to Move	Develop Architecture** New App Process
--	--	--	--	---

** Critical tasks

"Even if you're on the right track, you'll get run over if you just sit there"

Yogi Berra



Chart 13

© 2008 IBM Corporation

Linux

Information Sources

- <http://www.ibm.com/redbooks>
 - SG24-6552-00 Experiences with Oracle9i for Linux on zSeries
 - SG24-6482-00 Experience with Oracle Database 10g on Linux for zSeries
 - SG24-7191-00 Experiences with Oracle 10gR2 Solutions on Linux for System z
 - SG24-6669-00 Linux for IBM System z9 and zSeries
- <http://www.oracle.com/ibm>
 - IBM platform information
- <http://otn.oracle.com>
 - (Select "Downloads")
- <http://www.vm.ibm.com/perf/tips>
 - General z/VM Tuning Tips
- <http://www-124.ibm.com/developerworks/oss/linux390/index.shtml>
 - Lot's of information on Linux for zSeries
- <http://www-128.ibm.com/developerworks/linux/linux390/perf/index.html>
 - Hints and Tips for tuning Linux on System z
- <http://www.zseriesoraclesig.org>
 - Special Interest Group of Oracle users on the mainframe (z/OS and Linux)
- <http://www.mail-archive.com/linux-390%40vm.marist.edu/>
 - Marist List Server
- <http://www.oracleinsight.net/2008/02/06/the-mainframe-renaissance/>
 - The Mainframe Renaissance

Chart 14

© 2008 IBM Corporation

Linux

