

6th European GSE / IBM Technical University for z/VSE, z/VM and Linux on System z



22-24 October 2012 Hotel Hilton Mainz

Oracle & Linux on System z - The perfect fit

October 23, 2012

Dr. Holger Smolinski
smolinski@de.ibm.com
Global Client Center, IBM Germany R&D



Trademarks and disclaimers

■ Trademarks

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates.

IBM, IBM logo, bladecenter, eServer, iSeries, OS/400, PowerPC, POWER3, POWER4, POWER4+, pSeries, System p, POWER5, POWER5+, System x, System z, S390, xSeries, zSeries, zArchitecture, and z/VM are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Red Hat and the Red Hat logo are trademarks or registered trademarks of Red Hat, Inc. in the United States, other countries, or both.

Java and all Java-based products are trademarks of Sun Microsystems, Inc., in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

■ Disclaimers

Distributed server comparison is based on IBM cost modeling of Linux on zEnterprise vs. alternative distributed servers. Given there are multiple factors in this analysis such as utilization rates, application type, local pricing, etc., savings may vary by user.

Agenda

- **Server consolidation is continued towards core systems**
- **Oracle products are supported on Linux on System z and zBX**
- **TCO savings for database servers can be achieved with System z**
- **Lots of experience exists**
- **Common pitfalls can be easily avoided**

Linux on IBM System z

An ideal platform for enterprise-class IT optimization and cloud computing

Smarter computing based on an IT infrastructure that combines Linux® with the industry-leading IBM System z® and the out-standing IBM z/VM® virtualization technologies for server and workload consolidation, new Linux workloads and cloud at an attractive price.




Highlights

- Simplified IT infrastructure inside a single IBM System z
- Highly scalable, flexible and secure, sharing all system resources
- Tight integration of workloads
- Business continuance that help minimize revenue loss due to downtime
- Smarter computing at an attractive price — pay less as you run more

Saving Money and Reducing Complexity

- Run more applications/software at less expense
- Manage more virtual servers with fewer people
- Absorb workload spikes more easily
- Consolidate more servers per core
- Spend less on disaster recovery
- Occupy less floor space
- Save on energy

Helping you
“Do More with Less”
IT optimization and
Cloud computing




A refrigerator size box versus
versus a full room of servers.
The differences are quantum.

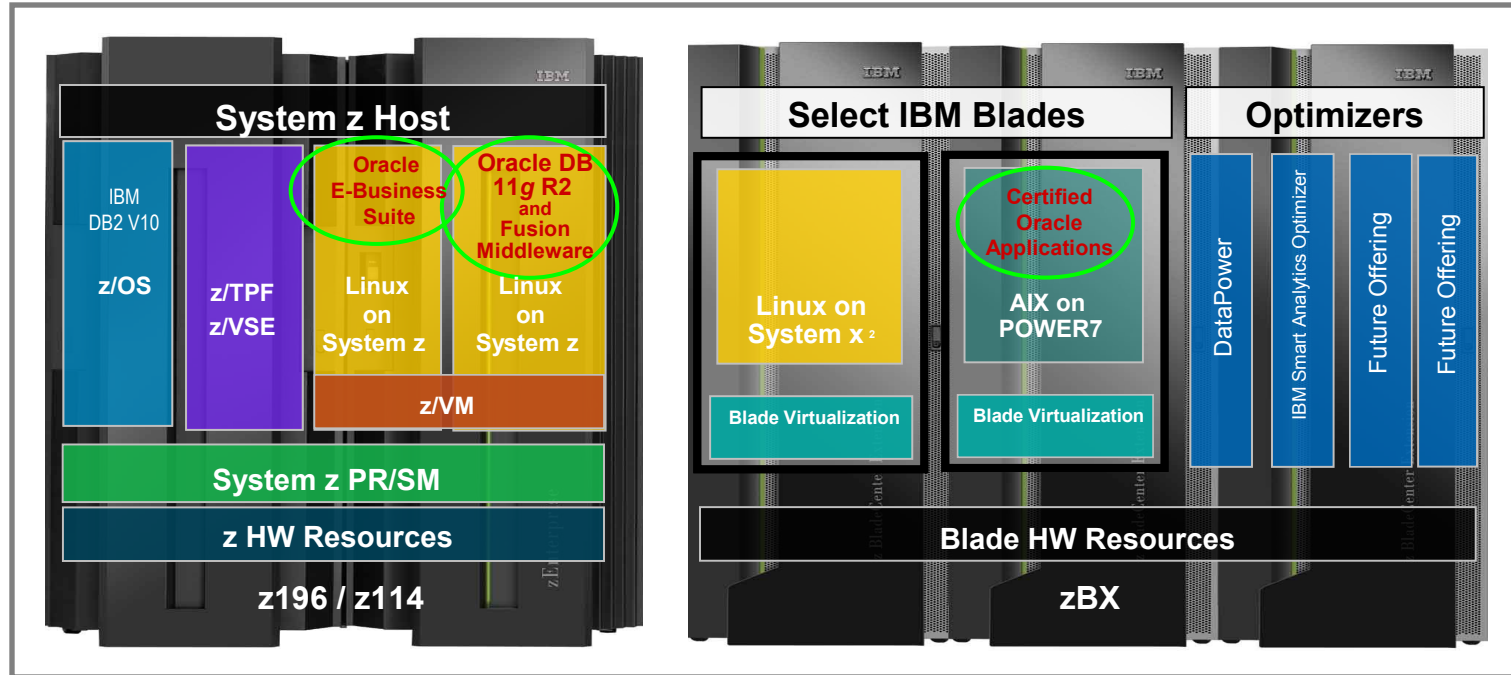
Why Oracle on Linux on IBM zEnterprise System

Fewer servers, fewer networks, fewer connections, and fewer software licenses may result in costs savings

- The best TCO characteristics can be obtained from consolidating many servers with moderate 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
- Linux scales well in an LPAR or with z/VM and may resolve other issues or problems such as availability
 - ▶ Great scalability for consolidation or single large databases
- **IBM System z differentiators**
 - ▶ Inherited Hardware Quality Of Service
 - ▶ Unmatched Virtualization Capabilities
 - ▶ It's Green
- Host complete Oracle solution suite on IBM zEnterprise™ by running Oracle Database on Linux on z/VM in conjunction with Oracle applications on zEnterprise blades
 - ▶ Simplified effort to manage a “fit for purpose” solution deployment on different architectures
 - ▶ Targeted for Oracle applications that interact with mainframe data and transactions




The zBX hybrid solution brings new options



Oracle solution landscape on a single zEnterprise System

- Run Oracle database and Oracle E-Business Suite on highly available and reliable zEnterprise 196 (z196) or zEnterprise z114 (z114)
- Run additional Oracle applications on AIX on zEnterprise BladeCenter (zBX)
- Fully benefit from zEnterprise, the first hybrid system, and manage all your Oracle solution servers with the zEnterprise Unified Resource Manager

The IBM Oracle Center (IOC)



OUR MISSION

Help IBM customers to deliver integrated solutions with Oracle Software Products on IBM Infrastructures



OUR STRENGTH

Cross platform team with strong knowledge on Oracle products and a wide network within IBM and Oracle ecosystem



OUR ACTIVITIES

- Convince : Briefings & Conferences
- Build : Architecture, Design, Sizing
- Demonstrate : Proof-of-Concept, Benchmarks
- Deliver : Publications & Workshops




COVERED PRODUCTS

- IBM Platforms (System z, Power, System x, Total Storage)
- Oracle Technologies (Oracle DB, RAC, ASM, Dataguard)
- Oracle Applications (EBS, Siebel & OBI & OWI)
- Entry point to other on Industry Solutions (BRM, iFlex, RETEK, Weblogic...)



Unified IBM / Oracle Architectures

IBM Oracle Center Customer Services



-  Architecture
-  Design Workshops
-  Technical Training
-  Live Demonstrations
-  Light Benchmark for Oracle
-  Collaterals
-  Cookbooks
-  Proof of Concept
-  Benchmarks



POWER7 technology delivers improved **energy efficiency** with unique features designed to enhance computation capacity

IBM Power system is **the most reliable** computer system to implement Siebel CRM solutions that are critical for business

IBM Power system is **the proven way to consolidate** swarming windows boxes within an enterprise

The unique and strong **IBM/Siebel Relationship** is a tremendous added value for clients

IBM Power Systems Family

A reliable, scalable, resilient server infrastructure is critical for businesses that demand powerful, flexible, reliable and secure computing solutions. IBM Power Systems technology is a smart choice for UNIX® and Linux® operating system based servers running Oracle Applications.

The Power Systems family with 1- to 64-core scalability is competitively priced and designed with 64-bit IBM POWER6 processor technology. This enables businesses to lower software, energy and space costs through performance leadership and system options that can have a dramatic increase in overall utilization. **A choice of more than 8,000 IBM AIX® and Linux applications support a broad array of business requirements.** Models can be selected for specific workloads, from front-end application servers to back-end database servers.

Oracle and IBM Alliance

In 1996, IBM and Siebel (created in 1993) began working jointly together to create reliable solutions that deliver business value and provide a competitive edge for their clients. IBM and Oracle continue to work jointly together.

The two companies have over **19,000 common installations** ranging from small to very large enterprises. The combination of Oracle's application expertise with IBM's comprehensive portfolio of services, optimized hardware and storage has helped clients reduce their total cost of ownership and mitigate risk.

IBM's service organization, the IBM Global Business Services, is an Oracle Certified Advantage Partner and has a proven track record with over **5,000 experienced professionals that have completed over 7,500 Oracle projects.**



For further information, please contact:
christophe.menichetti@fr.ibm.com
 IBM PSSC Customer Center - Montpellier

Center systems to boost performance...
 without necessarily adding servers.
 value and investment protection for...
 tion processing. They are designed to...
 - Simplify Deployment



ent. That means greater efficiency or...
 to one server, use savings in power and...
 nds.

nd BladeCenter systems designed for...
 m x3850 X5 and System x3690 X5 rack...
 Hot swap components and robust...
 he eX5 systems.

or 8-socket availability than monolithic

ou can run larger Oracle workloads,...
 s and IT complexity.

Cost reduction can be achieved...

The economics of Linux on z114 for consolidation and cost reduction

- Consolidate an average of **30 distributed servers** or more on a single core, or **hundreds** in a single footprint.
- Deliver a virtual Linux server for approximately **\$500 per year** or as little as a **\$1.45 per day per virtual server** (TCA)¹

TCA Analysis:


Consolidate 40 Oracle server cores onto 3 Linux cores on z114

Lower acquisition costs of hardware and software vs. distributed servers

– **up to 51% less than Nehalem²** in new footprint (Enterprise Linux Server (ELS)) –

– or **up to 62% less** when adding to existing footprint (Solution Edition for Enterprise Linux (SEEL)) –

Plus, additional savings in DR, floor space, power, cooling and labor costs



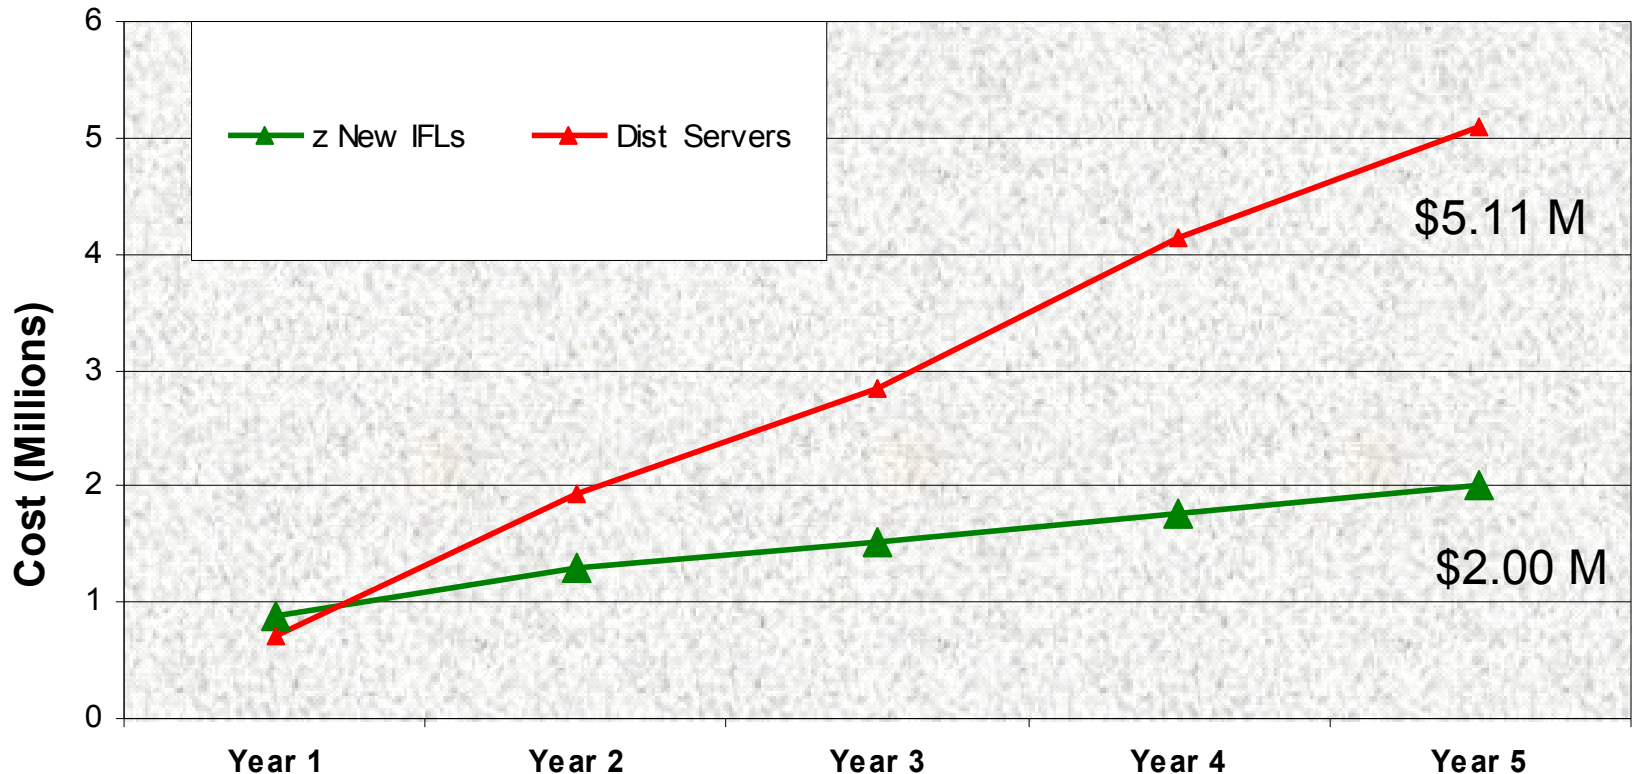
¹ Based on US Enterprise Linux Server pricing. Pricing may vary by country. Model configuration included 10 IFL cores running a mixed workload averaging 31 virtual machines per core with varying degrees of activity. Includes zEnterprise hardware and z/VM virtualization software. Does not include Linux OS or middleware software.

² Distributed server comparison is based on IBM cost modeling of Linux on zEnterprise vs. alternative distributed servers. Given there are multiple factors in this analysis such as utilization rates, application type, local pricing, etc., savings may vary by user.

5 Year TCO: System z IFLs vs Distributed Source Servers

Accumulated TCO Cost Comparison

IBM cost model based on real customer data



- System z saves +\$3.1 M (About 60% less than source servers)
- Start migration 2012 and complete within one year
- Yearly run rate lower on System z
- ROI less than 1 year

Deploy Oracle Software to the “Best Fit” Technology Tier to meet Business Requirements

IT Optimization and server consolidation with Linux on System z provide an excellent price performance for Oracle software.

**Oracle Database
11g Release 2
Version 11.2.0.2**

**Oracle E-Business Suite 12.1.x
native on
Oracle Database 11g Release 2**



- Excellent virtualization capabilities provides highest resources efficiency and utilization
- Exploitation of existing disaster recovery plans
- Low security breach risks/costs via most secure design, and high availability features
- Low power consumption, cooling, and data center floor space
- Low total cost of ownership, total cost per user, and total cost per transaction

Dundee City Council delivers value through new technologies

The new infrastructure runs a range of Linux applications and Oracle databases – supporting key systems such as social services 24x7.

“We can run 60 virtual machines on just four System z processors – whereas an equivalent x86-based architecture might require several processors for each server! So the savings can be considerable,” - Tim Simpson, IT Support

Manager, Dundee City Council

Dundee City Council

Creates a cost-effective IT architecture with IBM System z and IBM XIV Storage System technologies

The need:

Like all UK local authorities, Dundee City Council needs to handle increasing demand for IT and eGovernment services, while also reducing costs in line with central government targets. When the lease on its server and storage hardware needed to be renewed, the Council saw an opportunity to enhance its capabilities and increase value for money.

The solution:

Dundee worked with IBM to upgrade its mainframe environment with two powerful IBM® System z10® servers, and introduced the IBM XIV® Storage System to replace a mixed storage environment. The new infrastructure runs a range of Linux applications and Oracle databases – supporting key systems such as social services 24x7.

The benefits:

- Improves performance by more than 50 percent, providing capacity for growth without increasing IT costs
- Provides excellent availability and disaster recovery capabilities: in the event of a disaster at the main site, all systems can be restored at another location within 20 minutes
- Reduces Oracle licensing costs, as numerous virtual Linux servers can run on each IFL processor

“The combination of the z10 processors and the XIV grid architecture gives us 50 percent better performance... which means we can run 50 percent more workload for the same price.”

— Tim Simpson, IT Support Manager,
Dundee City Council

Solution components:

- IBM® System z10® Business Class
- IBM XIV® Storage System



Computacenter helps clients cut costs with the Enterprise Linux Server offering from IBM

Business challenge:

With many of its clients keen to move their business-critical systems onto Linux, Computacenter wanted to find a way to deliver world-class availability and security for Linux environments at a competitive price-point.

Solution:

IBM helped Computacenter deploy an IBM Enterprise Linux Server (ELS) at its Hatfield Solution Centre, enabling the company to run proofs-of-concept demonstrating the superiority of Linux on the mainframe over traditional distributed server architectures.

Benefits:

- The ELS can host thousands of virtual Linux environments within a single physical footprint – improving energy efficiency by up to 90 percent and supporting Green IT objectives.
- Powerful IBM Integrated Facility for Linux (IFL) processors reduce the cost of per-processor licensed software (such as Oracle) by up to 97 percent.
- IBM z/VM enables new Linux environments to be provisioned in minutes with no need to procure new hardware, increasing enterprise agility.

“The IBM ELS has the potential to provide the best of both worlds for our clients – combining the traditional reliability, availability and serviceability of System z with the openness and flexibility of Linux.”

— Paul Casey, Datacentre Platforms Practice Leader, Computacenter

Solution components:

- IBM Enterprise Linux Server (ELS)
- IBM Integrated Facility for Linux (IFL)
- IBM z/VM



Transzap boosts uptime with IBM System z

Transzap, Inc.

Business challenge:

Transzap offers its customers a comprehensive suite of financial software tools. As a small business with tens of billions of dollars in client transactions flowing through their systems each year, Transzap needed an economical, reliable platform to provide clients with high availability while enabling the capacity to accommodate growth within their software as a service business model.

Solution:

Transzap decided to consolidate on an IBM System z® platform to provide the stability and scalability needed to accommodate triple digit volume growth, enabling them to focus on the business of software innovation. Transzap migrated to System z and virtualized its critical applications on Linux® on System z, a platform that supports Transzap's dynamic Java™ and Oracle® environments.

Benefits:

- Helps Transzap to serve more than 69,000 users across 6,800 companies
- Provides higher levels of uptime for their customers
- Offers peace of mind through 24x7 world-class hardware support

“We intend to deliver a 99.9% application uptime guarantee to our customer base, thanks to the availability characteristics of System z.”

*— Peter Flanagan, CEO of
Transzap, Inc.*

Solution components:

- IBM System z®
- Linux® on System z
- IBM z/VM®



Bank of New Zealand ...

Reduces Carbon Footprint with Red Hat on the Mainframe

Business need:

Address environmental and space issues in the datacentre and achieve the corporate goal of becoming carbon neutral by 2010. Migration Path: From distributed Intel and SUN SPARC servers to Red Hat Enterprise Linux 5 running under z/VM on IBM z9 and z10 mainframes

Solution:

Software: Red Hat Enterprise Linux 5, Red Hat Network (RHN) Satellite, Oracle database, WebSphere Application Server, ESB, Process Server, TX and MQ Hardware: 1x z9 and 1x z10 mainframe (with 3 x IFL engines in each)

Benefits:

- Recovered 30 percent of datacenter floor space
- Reduced power consumption by 38 percent
- 20 percent return on investment (ROI) over the life of the platform
- Simplified, more efficient deployment

“The choice to invest in Red Hat was largely based on its commitment to the ongoing development of the platform and its strong support capabilities, particularly in reference to supporting Red Hat Enterprise Linux on the mainframe,”

- Lyle Johnston, infrastructure architect at BNZ

Common pitfalls can be avoided...

Keep in mind that resources are virtualized... - that implies multiplexing rather than multiplying

■ Oracle, like any other product, which uses in-memory operations assumes, that memory is real memory

→ Do not overcommit memory, or expect additional latency

- In Linux – any page fault requires I/O to be waited for:
 - Virtual memory of guest should be System + SGA + PGA + Application memory
Think about additional applications which run occasionally and might force out parts of the database
 - Each instance of a connection (pool) requires additional memory
Think about the amount of concurrent connections and do the sizing correctly
see:
<http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101619>
 - When using Oracle 11g, enable large page support with HUGETLB
 - Avoids duplication of page tables per forked process (Linux emulation under z/VM)
- In z/VM – any pseudo page fault will require IO to be waited for
 - Physical memory should be at least as large as 'hot' memory footprint
 - QUEUE 0 memory should be at least as large as 'hot' memory footprint

Rightsize the Memory for your Oracle guests - wrong

Example values taken from a real customer case – need review before being transferred to different situation

Linux System	512 MB
Oracle SGA	7168 MB
Oracle PGA	2458 MB
Oracle ASM	512 MB
Additional Tooling	512 MB
Linux Management Overhead (5%)	758 MB
Connections: $560 * 5.5 \text{ MB}$	3080 MB
Total:	~15 GB

Rightsize the Memory for your Oracle guests - right

Example values taken from a real customer case – need review before being transferred to different situation

Linux System	512 MB
Oracle SGA	7168 MB
Oracle PGA	2458 MB
Oracle ASM	512 MB
Additional Tooling	512 MB
Linux Management Overhead (5%)	758 MB
Connections: 560 * 5.5 MB	3080 MB
Linux Page Tables:	4253 MB
Total:	~19 GB

How large page support helps

- **One page table entry (8 Byte) is required per process and accessed page – $8/4096 \sim 0.2\%$ of process footprint.**
 - Page table memory is reserved (non swappable) memory
- **In Linux Threads can share page tables, processes don't.**
 - But threads cannot be used for any purpose...
- **Each Oracle connection is handled in its own process (not: thread!)**
 - processes share the memory, but not the page tables
- **With 512 connections sharing the memory the same amount is required for page tables**
 - some savings from memory which is not accessed
- **When handling many connections, page table overhead becomes a significant contribution to memory consumptions.**
 - Check `/proc/meminfo` for the actual use – in example mentioned before:
PageTables: 4355724



How large page support helps

- **With a page size of 1MB, one PTE (8Byte) per 1MB of memory is used**
 - PTE space down by a factor of 256, crossover at 128k processes
- **z/VM does not support huge pages (yet)**
- **Linux guests can define huge pages regardless of z/VM support**
 - z/VM will manage these pages as normal 4kB pages for paging etc...
 - Linux will use one Segment Table Entry per process and 1MB
 - All STEs referring to the same shared area share a set of PTEs
 - Memory consumption for N processes: $8 \times N + 4096$ Byte per MB
- **Define huge pages in Linux:**
 - Use kernel parameter `hugepages=8192` in `zipl.conf` for 8GB of huge pages
 -

...and low latency communications is required

■ **Virtualization and resource sharing may cause delays**

- Page faults and pseudo page faults (see previous page)
- Sufficient CPU capacity is required
 - Set LPAR weights and z/VM guest shares appropriately
 - Get alerted when there is 'steal time' in Oracle guests
- In z/VM disable page reordering for Oracle guests
 - Caution: Impacts z/VM memory management functions
- Use HiperSockets or z/VM Guest-LAN where possible

■ **Keep in mind the z/VM System limits**

- Consider LPAR for very large Oracle guests
- Consider LPAR for only few Oracle guests
- Consider LPAR for RAC members, OR design systems layout very carefully.

Breaking news – lessons learnt recently

- **Since Oracle 11g the Oracle DB has improved its behavior**
 - In Previous versions each thread executed a system call `gettimeofday()` when it needed accurate timer information – which is somehow CPU intensive
 - Recent versions of Linux have accelerated this by including `gettimeofday` in `vdso` calls (virtual dynamic shared object) – system calls without crossing the kernel boundary
 - Since Oracle 11g there is a single designated process (VKTM) acting as a time keeper and providing accurate time information to all threads and processes.
 - VKTM periodically calls `gettimeofday()` for course updates of time
 - when started as a high priority process it uses `nanosleep()`-calls in between to establish a high resolution timer
 - **Result: ~0.5% of CPU use by VKTM process**
- **Example customer: consolidating 190 Oracle DB instances would use $190 \times 0.5\% = 0.95$ processors for high resolution timer process**
 - Remember the discussion about tickless systems few years ago?
- **Solution option: remove VKTM from high_priority_processes by setting hidden Oracle system parameter**
 - Caution: May impact your DB function – check with Oracle support for safety!



Resources

- **Generic Mail box IBM Oracle Center (PSSC, Montpellier, France)**

- ioc@fr.ibm.com

- **Get started – White and Performance papers, Redbooks, Education, Special Interest Group**

- ibm.com/systems/z/os/linux/solutions/zlo_gs.html

- **FAQ from IBM and Oracle customers about Oracle running on IBM System z**

- [ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/3f3b9f07d6ffdf8c862577110073e1cd/\\$FILE/IBM%20zFAQ%20joint%20document%20June%2006%202010.pdf](http://ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/3f3b9f07d6ffdf8c862577110073e1cd/$FILE/IBM%20zFAQ%20joint%20document%20June%2006%202010.pdf)

- **WebLogic on IBM System z Linux Under z/VM-Quick Start**

- [ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/3f59d30df69853608625772800627b45/\\$FILE/WebLogic%20on%20IBM%20System%20z%20Linux%20Under%20zVM-Quick%20Start.ppt](http://ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/3f59d30df69853608625772800627b45/$FILE/WebLogic%20on%20IBM%20System%20z%20Linux%20Under%20zVM-Quick%20Start.ppt)

- **IBM System z running Linux with Oracle Technology Solutions**

- [ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/013136a5bc74513a8625761f005046ae/\\$FILE/ORA%20System%20z%20082409.pdf](http://ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/013136a5bc74513a8625761f005046ae/$FILE/ORA%20System%20z%20082409.pdf)

z/VM and Linux: Performance Best Practices When Running Oracle

- [ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/3371404eb07310cb8625764e0016f642/\\$FILE/BestPracticesOraclea.pdf](http://ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/3371404eb07310cb8625764e0016f642/$FILE/BestPracticesOraclea.pdf)

- **BM Light Benchmark Offering for Oracle Solutions on System z**

- [ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/393248ad2b4fc36b86257538001a8942/\\$FILE/ORA%20Light%20Benchmark%20System%20z%20040809.pdf](http://ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/393248ad2b4fc36b86257538001a8942/$FILE/ORA%20Light%20Benchmark%20System%20z%20040809.pdf)


- **The HA Stack on Linux for System z and Oracle**

- [ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/d02d497932b32b128625755100267f93/\\$FILE/The%20HA%20Stack%20with%20Oracle%20on%20Linux%20for%20System%20z.ppt](http://ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/d02d497932b32b128625755100267f93/$FILE/The%20HA%20Stack%20with%20Oracle%20on%20Linux%20for%20System%20z.ppt)

- **Selecting an Oracle Database for Server Consolidation on Linux for System z**

- [ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/6ddb491c9380e45886257493007d61f9/\\$FILE/ServerCons_SelectDatabase.ppt](http://ibm.com/support/techdocs/atmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/6ddb491c9380e45886257493007d61f9/$FILE/ServerCons_SelectDatabase.ppt)

Questions?



Holger Smolinski

Certified IT Specialist

Linux on System z

*Schönaicher Strasse 220
71032 Böblingen, Germany*

*Phone +49 (0)7031-16-4652
Holger.Smolinski@de.ibm.com*