



IBM eServer™ zSeries

## Consolidating Oracle Databases to Linux on System z and zSeries

Oracle9i and Oracle Database 10g

Denny Dutcavich

[dutch@us.ibm.com](mailto:dutch@us.ibm.com)

845.689.2226

© 2003 IBM Corporation

# Trademarks

**The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.**

|                           |   |                   |                |
|---------------------------|---|-------------------|----------------|
| AIX*                      | FICON Express                             | MQSeries*         | RS/6000*       |
| AIX 5L*                   | FlashCopy*                                | MVS               | pSeries        |
| CICS*                     | GDPS                                      | Netfinity*        | S/390*         |
| DB2*                      | Geographically Dispersed Parallel Sysplex | NetView*          | SP*            |
| DB2 Universal Database    | HiperSockets                              | OS/390*           | Tivoli*        |
| Domino                    | IBM*                                      | OS/400*           | WebSphere*     |
| e-business logo*          | IBM logo*                                 | Parallel Sysplex* | xSeries        |
| Enterprise Storage Server | IBM@server                                | PR/SM             | z/Architecture |
| FICON                     | IMS                                       | pSeries           | z/OS           |
|                           | Lotus*                                    | RACF              | z/VM           |
|                           |   |                   | zSeries        |

\* Registered trademarks of IBM Corporation

**The following are trademarks or registered trademarks of other companies.**

Linux is a registered trademark of Linus Torvalds

Penguin (Tux) compliments of Larry Ewing

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

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.

\* 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.



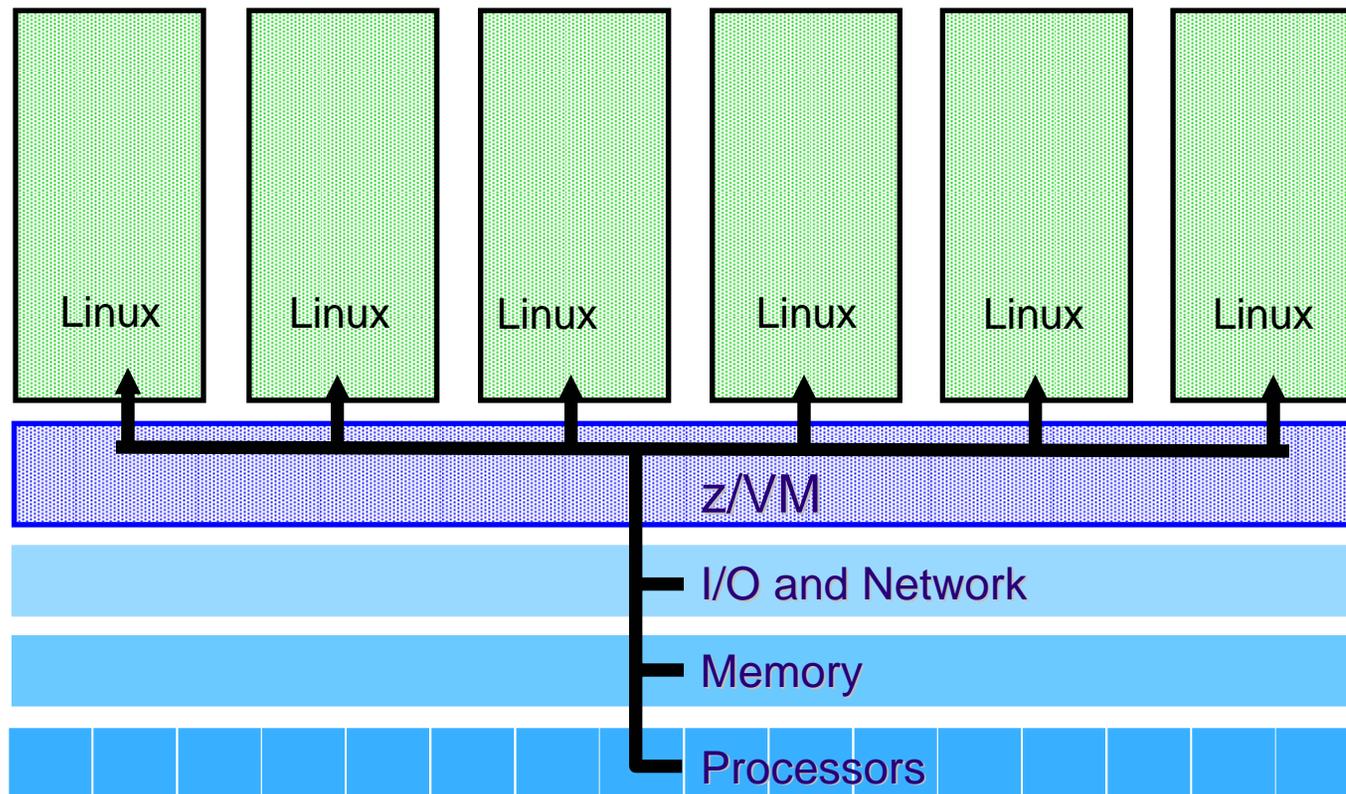
## Agenda

- Virtualization
  - A short course in virtualization on System z
- Linux on System z and Oracle
- Products Available on Linux for System z
  - Oracle9i Database Server
  - Oracle Database 10g Server
- Customer Status
- Selecting an Application Database for Linux on System z
- Additional Information

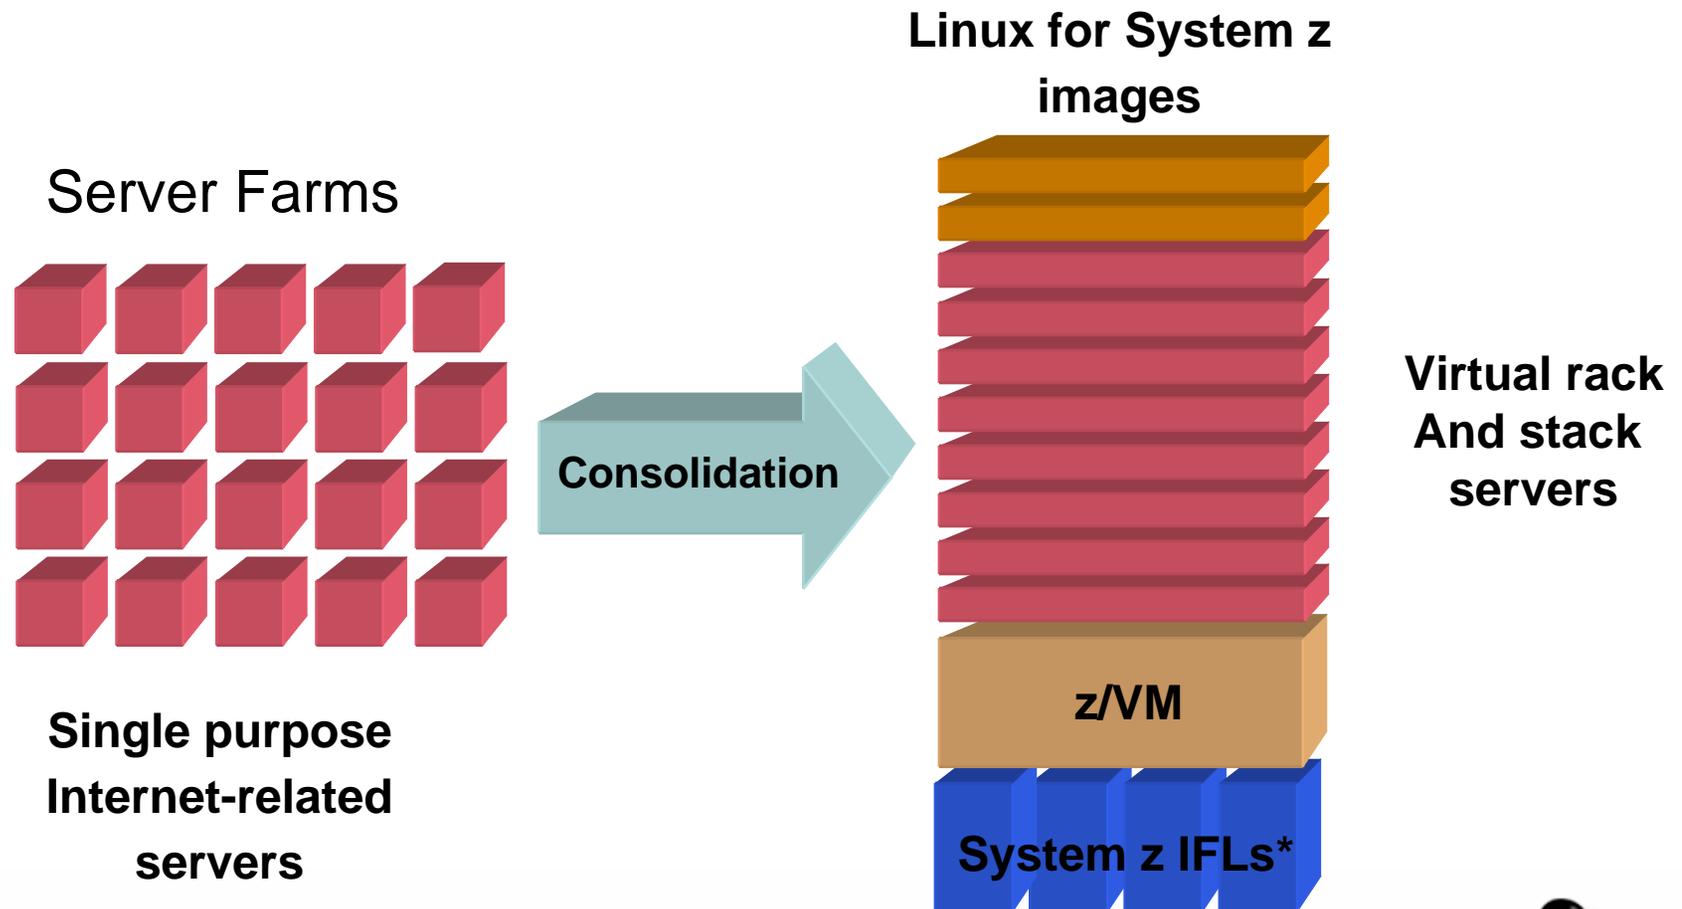
# Virtual Machine Partitioning

## Efficiently Exploiting the Entire Mainframe Complex

A *Virtual Machine* simulates the existence of a dedicated real machine, including processor functions, storage, and input/output resources.

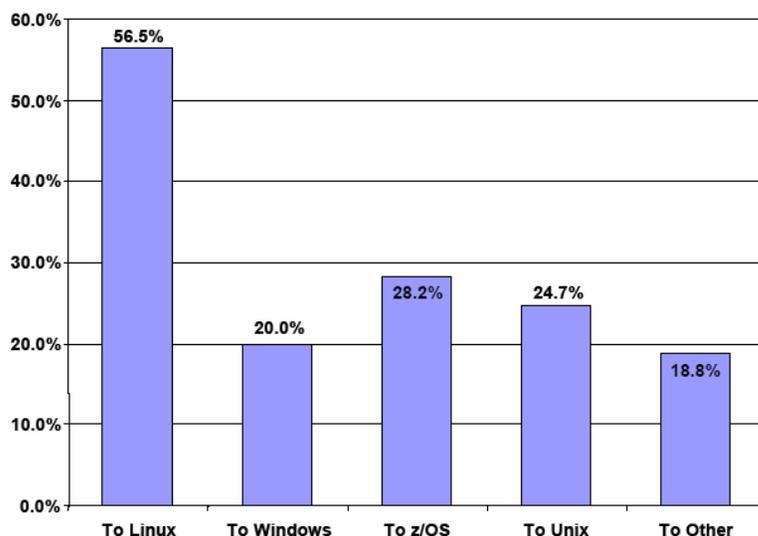


# Linux for System z Opportunity - 'The Server Farm in a Box'



## Customer Patterns

Figure 4. DBMS Platform Transition Plans



Note: Multiple responses permitted.  
Source: Gartner (March 2006)

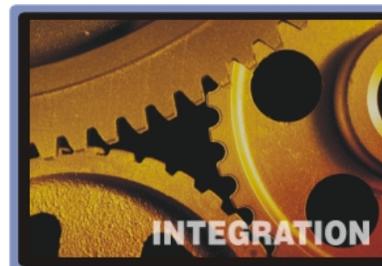
“We continue to see the mainframe chosen as older applications are rewritten or purchased. Finally, a percentage of those moving to Linux are moving to a zLinux LPAR on the z9 Series hardware for the same reasons of resource reduction. Many of these shops are moving Oracle from Unix to zLinux, often consolidating hundreds of Unix servers to a single mainframe.”

Source: Gartner Report G00138213, 29 March, 2006

## What Oracle Products are Available for Linux on System z?

### ■ Products

- Oracle9i Release 2
  - 31-bit database only
  - SuSE SLES8 and SLES9 only
- Oracle Database 10g Release 1
  - 64-bit database only
  - SuSE SLES8 and SLES9
- Oracle Database 10g Release 2
  - SuSE SLES9 and RHEL4
- AS 10g – projected 2006
- Applications
  - Oracle Collaboration Suite - now
  - Siebel7 CRM - now
  - EBusiness Suite – projected in 3Q06
- Oracle Clustered File System V2 (OCFS2)
  - In Linux Kernel today for Oracle RAC
  - Use as a general purpose FS in SLES10 and RHEL5



## Is Anyone Really Doing Oracle on Linux on System z?

- Yes!
  - Many customers in production with both Oracle9*i* and Oracle 10<sub>g</sub>
  - Many customers performing PoCs with Oracle 10<sub>g</sub> R1 and R2
- References available
  - State/Local/Federal Government
  - Utilities
  - Telco
  - etc
- What does an installation look like?
  - A large customer using z/VM is Telemar in Brazil with a32 IFL only z9-109 System z,
    - aprx 40 virtual machines in 10 IFL LPAR for Oracle
    - Using Oracle for billing and customer support
    - Other IFLs used for mail, Websphere Portal and dev/test
  - Govt customer moving over 100 databases to z9-EC on 5 IFLs
  - Many single IFL installations

*“I had zero questions about IBM hardware reliability. It’s the most reliable hardware that we have in our environment.”*

– Marq Youngblood, Chief Information Officer, Oklahoma Department of Human Services



# Oklahoma Department of Human Services



OKLAHOMA DEPARTMENT OF HUMAN SERVICES

## Improving responsiveness and reducing costs

The zSeries running Linux offers a reliable, stable and open-standards platform that employees can now rely on to provide 24x7 availability for critical data. Furthermore, the new solution has improved the department's responsiveness to urgent matters, such as those involving child abuse cases, and it has:

- Reduced the risks associated with unscheduled downtime – lost productivity could have cost an estimated US\$95,000
- Increased reliability of the new infrastructure to help avoid ongoing maintenance charges of US\$87,000 per year
- Improved services by providing more timely responses to its citizens
- Simplified the management of its IT infrastructure
- Decreased overall IT costs
- Enhanced performance by 20 percent at a reduced cost in some areas
- Added more capacity on the IBM mainframe and reduced database backup time
- Reduced the need to rely on multiple vendors and services

“Going forward, it's a better economic solution because we're running the Linux operating system on the IBM platform,” says Youngblood. “With Linux and IBM, our costs over time will be much better managed and significantly reduced.”



## Idaho Power Company



- Idaho Power Company – Boise Idaho
  - Provides power to 883,000 customers
- System
  - Z900 2064 – 103 with 2 IFLs
  - Early users of FCP/SCSI
- Why did the move from HP and Intel?
  - The mainframe “ities”
  - Centralized management of the distributed systems
  - Cost
    - Breakeven point was 2 Oracle databases (6 Oracle servers) moved from Intel
    - Also implemented DB2 connect and LDAP servers
- What they currently run
  - 30 Linux virtual machines
  - Mostly Oracle
    - Off the shelf applications
  - Data Center Power Monitor
  - Source Code Mgmt system In test
  - To test Oracle 10g for real time meter readings



## 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.
  - However, Linux scales well in an LPAR and may resolve other issues or problems such as availability.
- Linux runs on multiple IBM architectures and provides
  - A common set of tools such as GNU compiler, etc.
  - A common runtime
  - A common skill set for SA and application development
- Which can provide
  - Reduced costs
  - Simplified infrastructure
    - Simplified systems management through centralization
    - Enhanced business continuity and recoverability
  - Speed to market through technical exploitation and innovation

*“We chose Linux for the same reasons that we chose zSeries: the superior availability and reliability. We will be running mission-critical services in this environment, so the importance of these factors cannot be over-estimated.”* Tim Simpson, IT Support Manager, Dundee City Council



## Why Oracle for Linux on System z

- Customer Statement from the Marist Listserver

I would like to relay what is real concerning zLinux costs and what is theory or conjecture. I run a z/900-101 with one IFL along with my z/900- 101 and z/900-102 in a Parallel Sysplex. **First it makes no sense to bring up just 1-2 zLinux instances on an Enterprise Server. To make it cost effective, I figure it is around 20-30 and cheaper if there is more.** So that lets out just using LPARs without z/VM. With z/VM V5, it is cheap. Here is the breakdown: ...

Ok, I run a z/VM LPAR for Production with many Virtual Linux Servers. I am up to about 32 Virtual Servers and to do web serving, DB2 with IFL running about 10%. So it leaves me room to bring up more. I am looking to transfer a Windows \$40K Oracle license to the IFL. Oracle charges \$40K for an engine. **So you can do porting for free on any platform and only need the license when it goes production. So if we port the application to zLinux, the production license comes too. Now once I have the license, the next port means I can terminate a \$40K license on Windows. More likely I will just bring up more Oracle's on the zLinux using the one (1) license** I will have. What'a deal. The same goes for DB2, I can run many at the same price. **Then there is nothing too shabby about running Virtual Routers and Virtual Switches under z/VM (free).** The network group casting a dim eye wanting to see something physical. I can talk amongst them using Hypersockets or a memory to memory transfer. I run an LPAR Test & Development LPAR with all the zLinux virtual machines. My SYSPROGs wanted a SYSPROG LPAR for new z/VM's although the learned about z/VM 2nd level. WOW. **Is it dirt-cheap, no. Is it cheaper, yes.**

- Customer Statement from the Marist Listserver

In our case, **just the savings in Oracle by moving from Sun** plus some savings from zOS (moved some work over saving \$300k in zOS licenses) **more than compensated for the new costs.**



## Selecting an Application

- Where to start
  - Performance on zSeries CPUs comparable to CPUs on other platforms of similar speed.
    - CPU speed is not the entire story – it's in the architecture!
    - System z has definite advantage with applications that have mixed CPU and I/O
  - System z and z/VM provide excellent virtualization capabilities
    - Look for applications that are on lower utilized servers
    - Development and Test are good choices to start
  - Good planning is essential.
  - IBM can
    - Perform sizing estimates
    - Assist with planning and initial installation needs



*Cost benefit analysis was proven to be true. A corporation can successfully use LINUX technology in a z/Series environment to provide increased dBase and application performance, and realize real server consolidation and reduced licensing and support costs.*

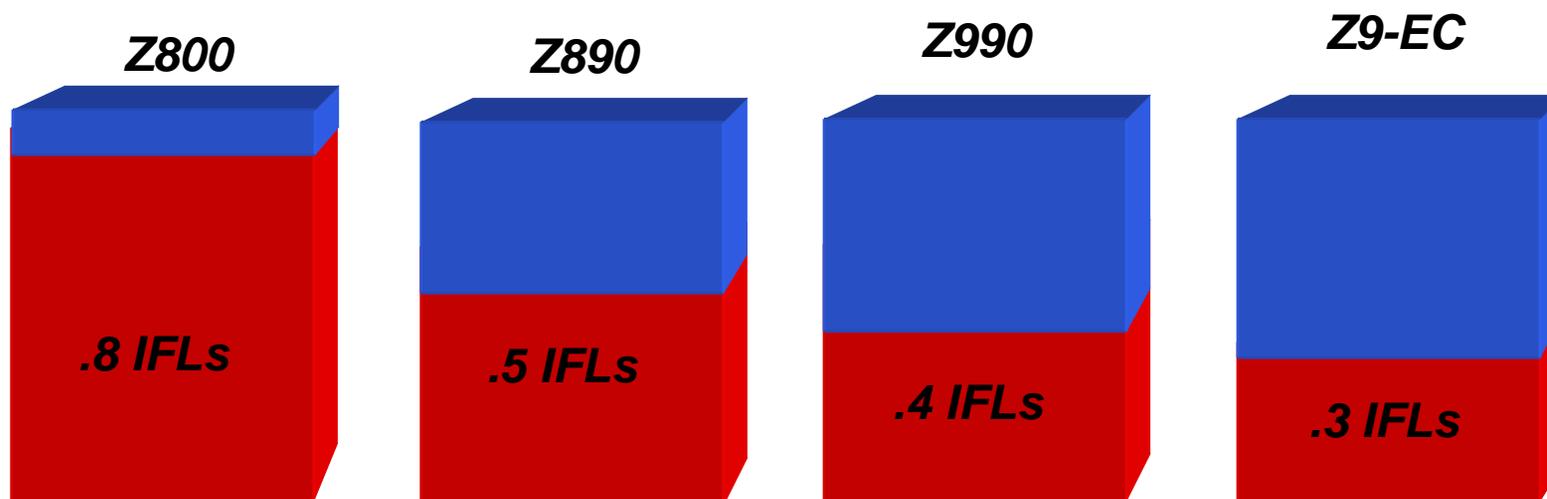
Ron Keogh – Project Manager, Idaho Power Company



## Selecting a System z

Consolidating 4 Sun Boxes running Oracle. Each with

- 2 CPs each @ 750 Mhz (total of 8 CPs)
- 15 % utilization



*In all cases - consolidated 8 Sun CPs into less than 1 IFL  
Over 3times that on the new EC class System z*

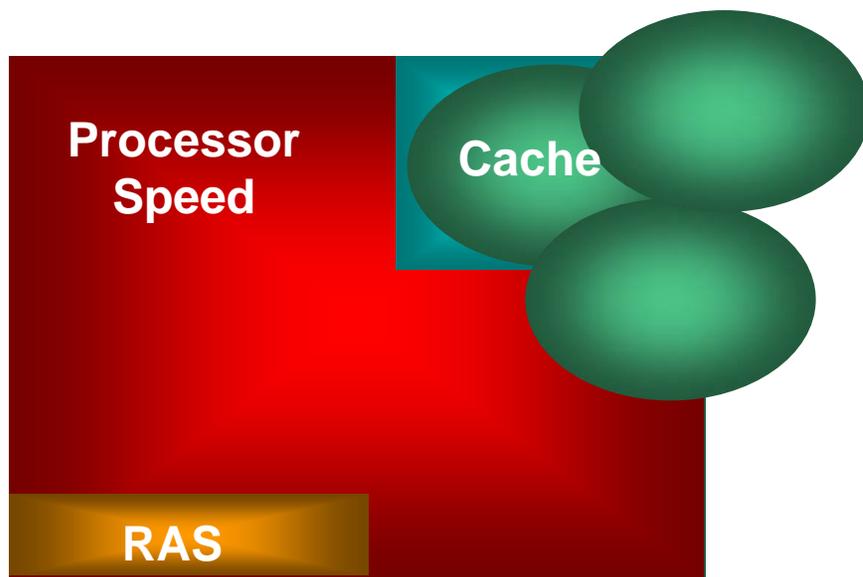


## Selecting an Application

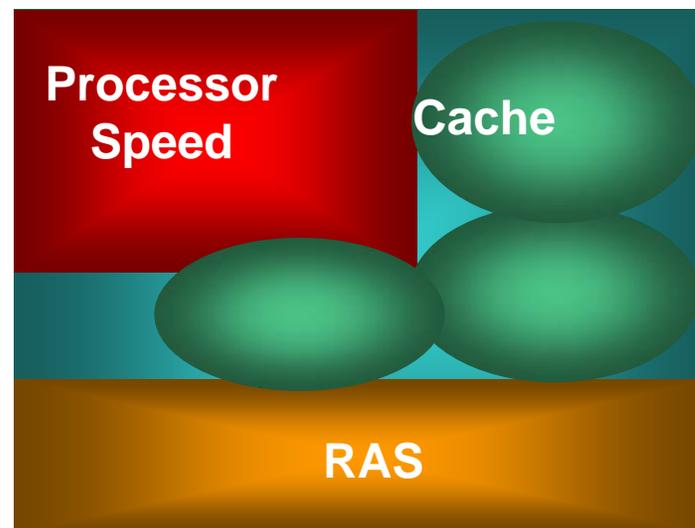
- One last thought
  - Test workloads selected for Linux on System z
    - Most difficult part of process
  - Benchmark testing (e.g. TPC-C) may not provide results needed to make a consolidation technology decision



## 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



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

Replicated

Consolidated



## Key Success Factors

- Use z/VM 5.2
- Memory is critical
  - Small virtual guest size – less is better even with z/VM 5.2
  - Let z/VM manage memory not Linux
    - Use VDISK for Linux swap
- Monitor resource
  - Make changes (i.e.tuning) that may be necessary
- Paging and swap space necessary
  - Both should use memory devices
  - Use Best Practices for setting up paging space
- Avoid I/O bottlenecks
  - Distribute data in the ESS across arrays
  - Consider striping with LVM
  - Use either ECKD or SCSI/FCP

## Upcoming Events

- San Francisco Linux World – August 13<sup>th</sup> to 17<sup>th</sup>, 2006
  - Both Chat on Oracle on System z
- z/OS Expo – October 9<sup>th</sup> - 13<sup>th</sup>
  - Caribe Royal Resort in Oralando
- Share – August 18<sup>th</sup>, 2006
  - Roadmap by Oracle
- System z Oracle Special Interest Group – April 2007
  - Clearwater FL (not booked yet)
  - Three to four day event

## 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
  - SG24-6669-00 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://awlinux1.alphaworks.ibm.com/developerworks/linux390/perf/tuning\\_rec\\_dasd.shtml](http://awlinux1.alphaworks.ibm.com/developerworks/linux390/perf/tuning_rec_dasd.shtml)
  - Hints and Tips for Selecting and Tuning I/O options
- <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

## Telemar Stays Competitive with zSeries and Linux

### Telemar

Telemar is the largest provider of fixed-line telecommunications services in South America. Telemar's 16-state region, which includes Rio de Janeiro, covers approximately 64 percent of the country and is home to more than 93 million citizens.

### Challenge

Telemar was formed by the privatization and merger of 16 state-owned phone companies. It had undergone radical transformation and centralization to be competitive in the deregulated telecommunications industry. To control its overhead costs and remain competitive, Telemar needed to centralize its IT systems and lower its total cost of ownership (TCO).

### Solution

- Telemar consolidated 16 geographically dispersed servers onto a centralized System z 109-EC server running SUSE Linux Enterprise Server V8 and z/VM 5.2.
- They migrated their Eureka billing application consisting of many Oracle Databases, along with email, Websphere Portal and other office applications.
- Telemar installed new web-based user interfaces to applications.
- The solution included Tivoli Storage Manager for backup and recovery

### *Benefits:*

Telemar sought an open-standards-based solution that would maximize the manageability, scalability, security and availability of its key business systems.

- Reduced need for server capacity by one-third
- Virtual servers are easier to manage than the physically distributed environment
- Reduced Eureka application run time for monthly billing cycles by five hours.
- Lowered operating and administration for maintaining email server applications.



## State of Minnesota gains rapid Return on Investment with zSeries, Linux Solution

### State of Minnesota InterTechnologies Group

The State of Minnesota's InterTechnologies Group (InterTech) has been providing telecommunications resources and service to the state's public sector for more than 30 years.

### Challenge

The infrastructure for the application for management and administration services which supported all 84 counties of Minnesota. It was running on Compaq servers scattered across six different locations and was contributing to:

- Poor response time
- Scalability issues

Due to budget constraints, the solution needed to solve these problems and demonstrate a rapid return on investment.

### Solution

InterTech migrated its social services applications to IBM and is running:

- eServer zSeries 990 server running an Integrated Facility for Linux (IFL) and the SUSE Linux Enterprise V8 operating system.
- z/VM 5.1 operating system
- Virtualization
- Consolidated six locations down to one

### Benefits:

- IBM was able to deliver a full ROI in less than 18 months
- Reducing the number of physical servers and corresponding Oracle instances, InterTech was able to reduce costs associated with software licensing and hardware maintenance
- Application response time has improved
- The open-source Linux environment provides high levels of stability



# Swisscom IT Services Leverages zSeries with Linux to Cut Costs and Improve Client Services

## Swisscom IT Services

Swisscom IT Services AG is a leading information technology services provider in Switzerland, with competencies in telecommunications and financial services.

## Challenge

Swisscom IT Services had been hosting its client's application environment on 21 geographically dispersed servers, making maintenance and management difficult. Furthermore, the client's chosen platform, HP True64, was nearing the end of life.

Swisscom was concerned about: cost, scalability, and ease of migration.

## Solution

Swisscom already had two zSeries 990 servers in-house. They were able to migrate all 21 HP servers onto the existing zSeries 990 servers without any impact to their installed applications.

- The solution included running two logical partitions with z/VM v5.1 to create a total of five integrated facilities for Linux (IFLs) running SUSE LINUX Enterprise Solution v8 operating system.
- The IFLs are used to create virtual Linux servers for testing and development.
- IBM TotalStorage San Switches are used to connect the Linux environment to their Hitachi Data Systems storage.

## Benefits:

Swisscom IT Services has significantly reduced its overall operating costs and improved client delivery performance.

- They reduced their Oracle licensing fees by reducing the number of processors from 21 to five.
- By relying on the Linux operating system, Swisscom has improved the price/performance ratio of its application environment.
- The new environment enables Swisscom to provide their clients with a fast, flexible, stable server environment.
- Swisscom now has the ability to reappportion resources to new projects in a matter of minutes in accordance to customer demand.



# Portal built on zSeries/Linux Solution Exceeds Expectations

## EDS Brazil

With more than 15 years of experience, over 6,200 employees, and offices throughout the country, EDS Brazil is one of the leaders in the information technology services industry in Brazil.

## Challenge

EDS Brazil had more than ten intranets distributed across various department, most of the running under small local servers. This caused delays in information sharing and business process.

- From a technology perspective, problems included a lack of integrated and central control; difficult and time consuming maintenance and upgrade processes; and delays in integrating new users and implementation of new functions.

## Solution

EDS Brazil implemented a Portal project to integrate existing departmental intranets and all internal systems.

Phase 1: A team built the integrated intranet based on the IBM eServer zSeries platform, z/VM, SUSE LINUX, Apache, PHP, LDAP, and an Oracle Database.

Today: The project runs on an IBM eServer zSeries 900 with Novell SUES LINUX v8 and five integrated facilities for Linux – and supports more than 15 applications from HR to customer help, marketing, and administration.

## *Benefits:*

This IBM portal solution has met and exceeded EDS Brazil's business needs.

- The portal has helped EDS Brazil centralize information, making it easier to find, manager, update, and control.
- Time to process vacation requests reduced by 50%, time to upgrade information from the help desk has decreased by a factor of three, the average time to find general information is less than five minutes.
- Systems management and maintenance have been greatly simplified.
- Developers need fewer resources to producer more and better output.
- More than ten servers have been freed up for other use.
- EDS benefits from high availability, greater processing performance, and stability.