



8470: Overview of Oracle Solutions available on Linux on IBM System z and experiences with Oracle Database Proof of Concept Implementation Projects

Wednesday, March 2, 2011, 8:00 AM-9:00 AM

Gaylan Braselton
IBM Sales Consultant, NA
Oracle on System z

gbrasel@us.ibm.com



Trademarks

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

| | | | |
|--|---|------------------------------------|----------------|
| AIX* | Geographically Dispersed Parallel Sysplex | OS/390* | VSE/ESA |
| APPN* | HiperSockets | Parallel Sysplex* | VTAM* |
| CICS* | HyperSwap | PR/SM | WebSphere* |
| DB2* | IBM* | Processor Resource/Systems Manager | z/Architecture |
| DB2 Connect | IBM eServer | RACF* | z/OS* |
| DirMaint | IBM e(logo)server* | Resource Link | z/VM* |
| DRDA* | IBM logo* | RMF | z/VSE |
| Distributed Relational Database Architecture | IMS | S/390* | zSeries* |
| e-business logo* | InfoPrint* | Sysplex Timer* | |
| ECKD | Language Environment* | System z | |
| Enterprise Storage Server* | MQSeries* | System z9 | |
| ESCON* | Multiprise* | TotalStorage* | |
| FICON* | NetView* | Virtualization Engine | |
| GDPS* | On demand business logo | VSE/ESA | |

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of Intel Corporation in the United States, other countries, or both.

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 Linus Torvalds in the United States, other countries, or both.

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.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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.

Topics -- Oracle Solutions On Linux on z

- IBM Oracle Partnership and Commitment
- Oracle System z Solutions and Support
- Other Customers doing Oracle on Linux on z
- Proof of Concept Process for Oracle on Loz
- Sizing and Implementation Process

Complete Solutions, Collaborative Partners

Enduring Relationship

- Oracle 24 Years, PeopleSoft 24 Years, Siebel 14 Years
- Mission: Provide the strongest combination of solutions to maximize value of client investments

Mutual Executive Commitment

- Regular Senior Executive Reviews and Development & Sales interlocks
- Dedicated, executive-led IBM Alliance Team

Over 30,000 Joint Applications Customers Worldwide

- Hardware and Software support via Applications Unlimited

Market Leading Services Practice

- ~ 5,500 Successful Joint Services Projects
- > 10,000 IBM Consultants Worldwide dedicated to Oracle Solutions

Vibrant Technology Collaboration

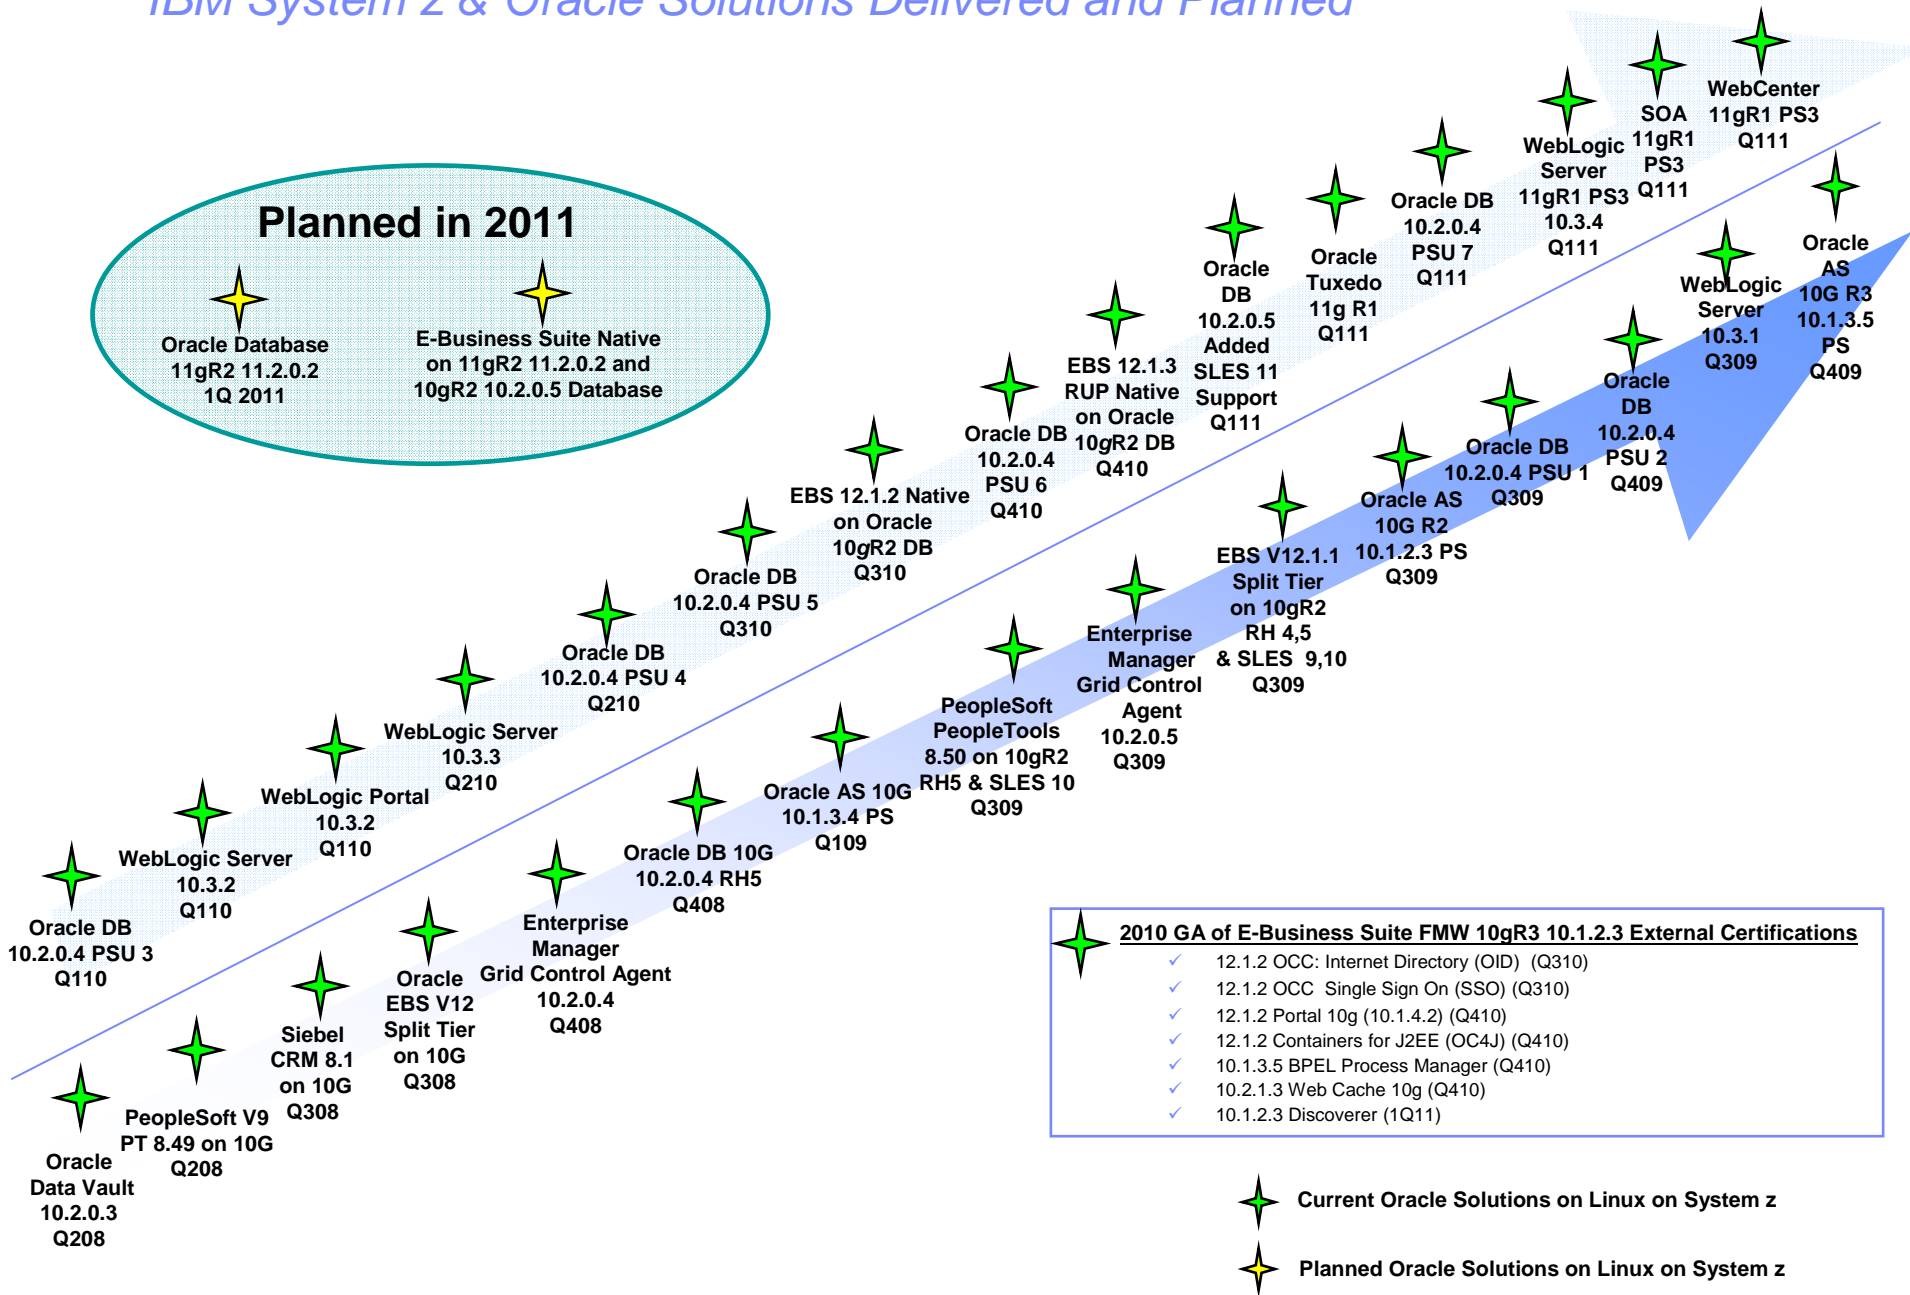
- Substantial investment in skills and resources
- Dedicated International Competency Center

Unrivalled Joint Customer Support Process

- Dedicated On-Site Resources
- Significant Program Investments

Linux on System Servers z with Dedicated Porting Teams

IBM System z & Oracle Solutions Delivered and Planned



Oracle Solutions Available Today on IBM System z Servers

IBM Data Server on DB2 z/OS and or Linux

Oracle DB Server on Linux

** Note: Multi-Platform "Split Tier" Configuration – Only the Database runs on System z Servers*

| | | | |
|--------------------------------|------------------------------|---|---|
| ERP & CRM Solutions | Oracle PeopleSoft Enterprise | * Version 9.0 & 9.1 /Tools 8.49 & 8.50 DB2 V8, V9 (Database and Batch Server Supported) | * Version 9.0 & 9.1 /Tools 8.49. & 8..50 Oracle 10gR2 |
| | Oracle Siebel Enterprise | * Version 8.0 & 8.1.1 DB2 V9.1 | * Version 8.0 & 8.1.1 Oracle 10gR2 |
| | Oracle E-Business Suite | | New!!! Version R12.1.3 Oracle 10gR2 |
| Banking & Insurance | Oracle Financial Services | | FLEXCUBE Retail Core Banking V2.2 Universal Banking (UBS) V10 Oracle 10gR2 on SLES9 |
| | Oracle Insurance | Documaker 11.4 DB2 8.2 & 9.1 z/OS | |
| | Oracle Cross Industry | | Oracle Policy Automation v10.1 Oracle 10gR2 on SLES10 |

Oracle Solutions Available on IBM System z Servers

Public Sector Solutions: Taxpayer registration, tax return processing, revenue collection and audit, Siebel CRM

IBM Data Server on DB2 z/OS or Linux

**Oracle
Enterprise Tax Management**

**Version 2.2
DB2 V8 & 9, WAS 6.1**

**Oracle
Siebel CRM for Public Sector**

**Version 8.2
DB2 V8 & 9**

BIEE Solutions: Data Source Only

IBM Data Server on DB2 z/OS or Linux

Oracle DB Server on Linux

**Oracle
Business Intelligence
Enterprise Edition**

***Version 10.1.3.4.1
DB2 V8.2 & and v9**

***Version 10.1.3.4.1
Oracle 10gR2**

** Note: Multi-Platform "Split Tier" Configuration) – Only the Database runs on System z Servers*

Oracle Server Technology Available or Planned For Linux on IBM System z

| | Oracle Solution | Version Available | Technology Status - Planned |
|---------------------------|---|--|--|
| Database | Oracle Database 10gR2 | Oracle DB 10g Release 2 10.2.0.4, 10.2.0.5 | |
| | Oracle Data Vault 10gR2 | Oracle Data Vault 10.2.0.4 | |
| | Oracle Database 11gR2 | | Oracle DB 11g Release 2 Planned – Q1 2011 |
| Fusion Middleware | Oracle FMW 10gR2/10gR3 Application Server | Oracle Application Server 10gRelease 2 10.1.2.3 10g Release 3 10.1.3.5 | |
| | Oracle FMW 11gR1 | WebLogic Server 10.3.2, 10.3.3 WebLogic Portal 10.3.2 SOA 11gR1 WebCenter 11gR1 Tuxedo 11gR1 | |
| Enterprise Manager | Oracle Enterprise Manager Agent | Oracle Enterprise Grid Control Agent 10.2.0.5 | |

Note: Oracle Database Components

- Oracle Real Application Clusters
- Oracle OLAP
- Oracle Spatial
- Oracle Label Security
- Oracle Partitioning
- Oracle Data Mining
- Oracle Advanced Security
- Oracle Data Guard



IBM zEnterprise System – Best-in-class systems and software technologies A “System of Systems” that unifies IT for predictable service delivery



IBM zEnterprise 196 (z196)

- Optimized to host large-scale database, transaction, and mission-critical applications
- The most efficient platform for large-scale Linux consolidation
- Capable of massive scale-up
- New easy-to-use z/OS V1.12

zEnterprise Unified Resource Manager

- Unifies management of resources, extending IBM System z qualities of service end-to-end across workloads
- Provides platform, hardware and workload management

zEnterprise BladeCenter Extension (zBX)

- Selected IBM POWER7 blades and IBM System x Blades* for tens of thousands of AIX and Linux applications
- High-performance optimizers and appliances to accelerate time to insight and reduce cost
- Dedicated high-performance private network

Oracle Licensing with IBM System z

- **An Integrated Facility Linux (IFL) is an engine or a core**
- **System z has an Oracle core factor of 1.0**
- **Most System z Oracle implementations / licenses are Oracle Core Based Pricing (vs User Based Pricing)**
- **Most System z Oracle implementations / licenses are Oracle Enterprise Edition (EE)**
 - All z9 EC, z10 EC and z196 IFLs must be Enterprise Edition
 - Can be Oracle RAC but is not required
- **Only the z10 BC is eligible for Oracle Standard Edition (SE)**
 - It is licensed by socket

Oracle's Dedicated Level 2 System z Support

Enhancement to existing Support for Linux on System z Servers

- Manager, Raimund Reng
- System z skilled and knowledgeable
- WW Support
- Level 2 support team
- Request z team connected when z environment problems
- Joint User/Oracle/IBM Conference Calls
- Webcast, March 16th, Oracle z support update



My Oracle Support Communities

- **Migration from the older Forum format**
- **Actively managed and moderated**
- **Encourages user posts**
- **Spotlight and highlight posts**
- **Specific community for System z customers**
- **Accessed via My Oracle Support (Metalink)**
 - ▶ **support.oracle.com, click on Community, Subscribe to and Click on zSeries Platforms**
- **Announcements will be made in Community Featured Section !**
 - ▶ **“Webinars” One hour “brownbag” type presentations given by Oracle Support**
 - ▶ **First one planned in December 2010**
 - ▶ **Topic will be ASM**
- **Join the Oracle zLinux community and help to shape the future:**
 - ▶ **Provide feedback**
 - ▶ **Exchange ideas**
 - ▶ **Get answers**
 - ▶ **Expand networks**
 - ▶ **Share successes**

Customers running Oracle on Linux on IBM System z

- ✓ **Hundreds of customers running Oracle on Linux on IBM System z**
 - ✓ **Various sizes and deployments**
 - ✓ **Across industries**
 - ✓ **Active volunteer led System z Oracle User Group (www.zseriesoraclesig.org)**
- ✓ **Small System z Customer Example**
 - ✓ **Oil & Gas Industry Services Provider**
 - ✓ **Serves 4,200 Companies, 44,000 Users, \$80B in transaction detail yearly**
 - ✓ **Was Windows, Dell, Linux**
 - ✓ **ISSUES – Rapid Company Growth, Server Sprawl, Cost Control, Hardware Outages**
 - ✓ **SOLUTION – z10BC, 3 IFL's, 24 GB**
 - ✓ **SLES10, Oracle 10g EE**
 - ✓ **Databases: 7 production, 400 GB – 3 TB, 7 virtual servers/database**
- ✓ **Medium Size System z Customer Example**
 - ✓ **Scientific Equipment Retailer**
 - ✓ **Oracle Database on mostly UNIX (also z890 legacy, was considering 'move off mainframe')**
 - ✓ **ISSUES – Availability on UNIX, growth (without additional footprints), strong UNIX team**
 - ✓ **POC – Linux on z, revealed 3X performance increase over UNIX**
 - ✓ **SOLUTION – z10BC, 7 IFL's, 76GB**
 - ✓ **Global DW, Oracle DB on Loz**
 - ✓ **Migrated two more databases onto Loz**
 - ✓ **COGNOS, home grown applications accessing the databases**
 - ✓ **In process of upgrading for additional Oracle database workloads**

Customers running Oracle on Linux on IBM System z

- ✓ **Large System z Customer Example**
 - ✓ Large Government Installation
 - ✓ 100 IFL's, z10 Oracle RAC Environment across 2 z10's with Oracle ASM
 - ✓ 35TB Database and 45TB Flash Recovery Area
 - ✓ Project is getting very high I/O throughput inserting 5.79 billion records in a 7 hour window and updating another 320 million records (exceeds 5 year SLA's)
- ✓ **Large System z Customer Example**
 - ✓ Leading Systems Integrator and IT Consulting Firm
 - ✓ Z990 x 2, Z9 S54 x 4, z10 E64 x 1 (192GB to 256GB per box)
 - ✓ 32 IFL's per z990, 54 per z9, 64 per z10
 - ✓ All Linux
 - ✓ 5 LPAR's per CEC (4 for Oracle, 1 for Management)
 - ✓ 16 shared IFL's per LPAR, 45 GB Memory per LPAR
 - ✓ 4 nodes RAC running on same CEC with Hipersocket interconnect
 - ✓ 2,000 – 3,000 transactions per second at peak
 - ✓ Response time less than 1 sec (threshold 5 sec)
 - ✓ DB Size – >5 TB for online & ~50 TB for DWH
 - ✓ Benefits – TCO, Extreme high availability, scalability
 - ✓ Planning z196 upgrade



International zSeries Oracle SIG

- **Independent User Organization**
 - President, Mike Zechman
 - Worldwide user participation
 - Oracle and IBM Participation

- **Annual Conference**
 - Next is April 10 – 14, 2011, Orlando (Oracle Collaborate)

- **Communicates requirements to Oracle and IBM**

- **Longest running still active Oracle User Group**

- **Website www.zseriesoraclesig.org**
 - Presentations, Links, Bulletin Board



Information Sources

- <http://www.oracle.com/ibm>
 - Oracle IBM Partner Relationship
- <http://otn.oracle.com>
 - Oracle Select “Downloads”
- <http://www.vm.ibm.com/perf/tips>
 - General z/VM Performance & Tuning Tips, Capacity planning
- <https://support.oracle.com>
 - Oracle Support Webpage (My Oracle Support)
- <http://www-124.ibm.com/developerworks/oss/linux390/index.shtml>
 - Lot’s of information on Linux for zSeries, IBM DeveloperWorks
- <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.ibm.com/redbooks>
 - SG24-6482-00 Experiences with Oracle Database 10g on Linux for zSeries
 - SG24-7191-00 Experiences with Oracle 10gR2 Solutions on Linux for System z
 - SG24-7573-00 Using Oracle Solutions on Linux on System z
 - SG24-7634-00 Experiences with Oracle Solutions on Linux for IBM System z



Linux on IBM System z

Experiences Implementing Oracle Database in a Linux on IBM System z Environment

Tom Kennelly
IBM Technical Specialist, NA
Oracle on System z

kennelly@us.ibm.com

Trademarks

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

| | | | |
|--|---|------------------------------------|----------------|
| AIX* | Geographically Dispersed Parallel Sysplex | OS/390* | VSE/ESA |
| APPN* | HiperSockets | Parallel Sysplex* | VTAM* |
| CICS* | HyperSwap | PR/SM | WebSphere* |
| DB2* | IBM* | Processor Resource/Systems Manager | z/Architecture |
| DB2 Connect | IBM eServer | RACF* | z/OS* |
| DirMaint | IBM e(logo)server* | Resource Link | z/VM* |
| DRDA* | IBM logo* | RMF | z/VSE |
| Distributed Relational Database Architecture | IMS | S/390* | zSeries* |
| e-business logo* | InfoPrint* | Sysplex Timer* | |
| ECKD | Language Environment* | System z | |
| Enterprise Storage Server* | MQSeries* | System z9 | |
| ESCON* | Multiprise* | TotalStorage* | |
| FICON* | NetView* | Virtualization Engine | |
| GDPS* | On demand business logo | VSE/ESA | |

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of Intel Corporation in the United States, other countries, or both.

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 Linus Torvalds in the United States, other countries, or both.

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.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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

- Objectives
- Terminology
- Oracle DB on Linux on z (Loz) through the eyes of availability
- Getting started - Sizing
- Proof of Concept
 - z/VM, Linux, Oracle DB, disk
 - Load the database and test
 - Performance evaluation
- Production Readiness
- Oracle Applications
- Summary of PoC

Objectives

- Awareness of Oracle DB on Linux on z
- How to get started with Oracle DB in a systematic way
- How to run a competent Proof of Concept (PoC)
- How to measure the PoC's performance
- How to make the environment production ready

Presentation based on Oracle DB 10gR2. 11gR2 early drop code has been delivered to a few customers

zLinux - It is really Linux on z


- zLinux is an IBM term for running Linux on z
- zLinux is NOT a special distribution of Linux created by IBM.
- Linux can run natively on System z or under z/VM
- Linux on z (Loz) is a better term and is less confusing in the Oracle space
- SLES and Red Hat have distributions for Loz

What is an IFL (Integrated Facility for Linux)?

- An IFL is a specialty engine on a System z
- Oracle uses the term core for purposes of pricing
- From an Oracle perspective an **IFL = a core**
- Yes, the System z has quad core processors BUT an IFL represents one core.

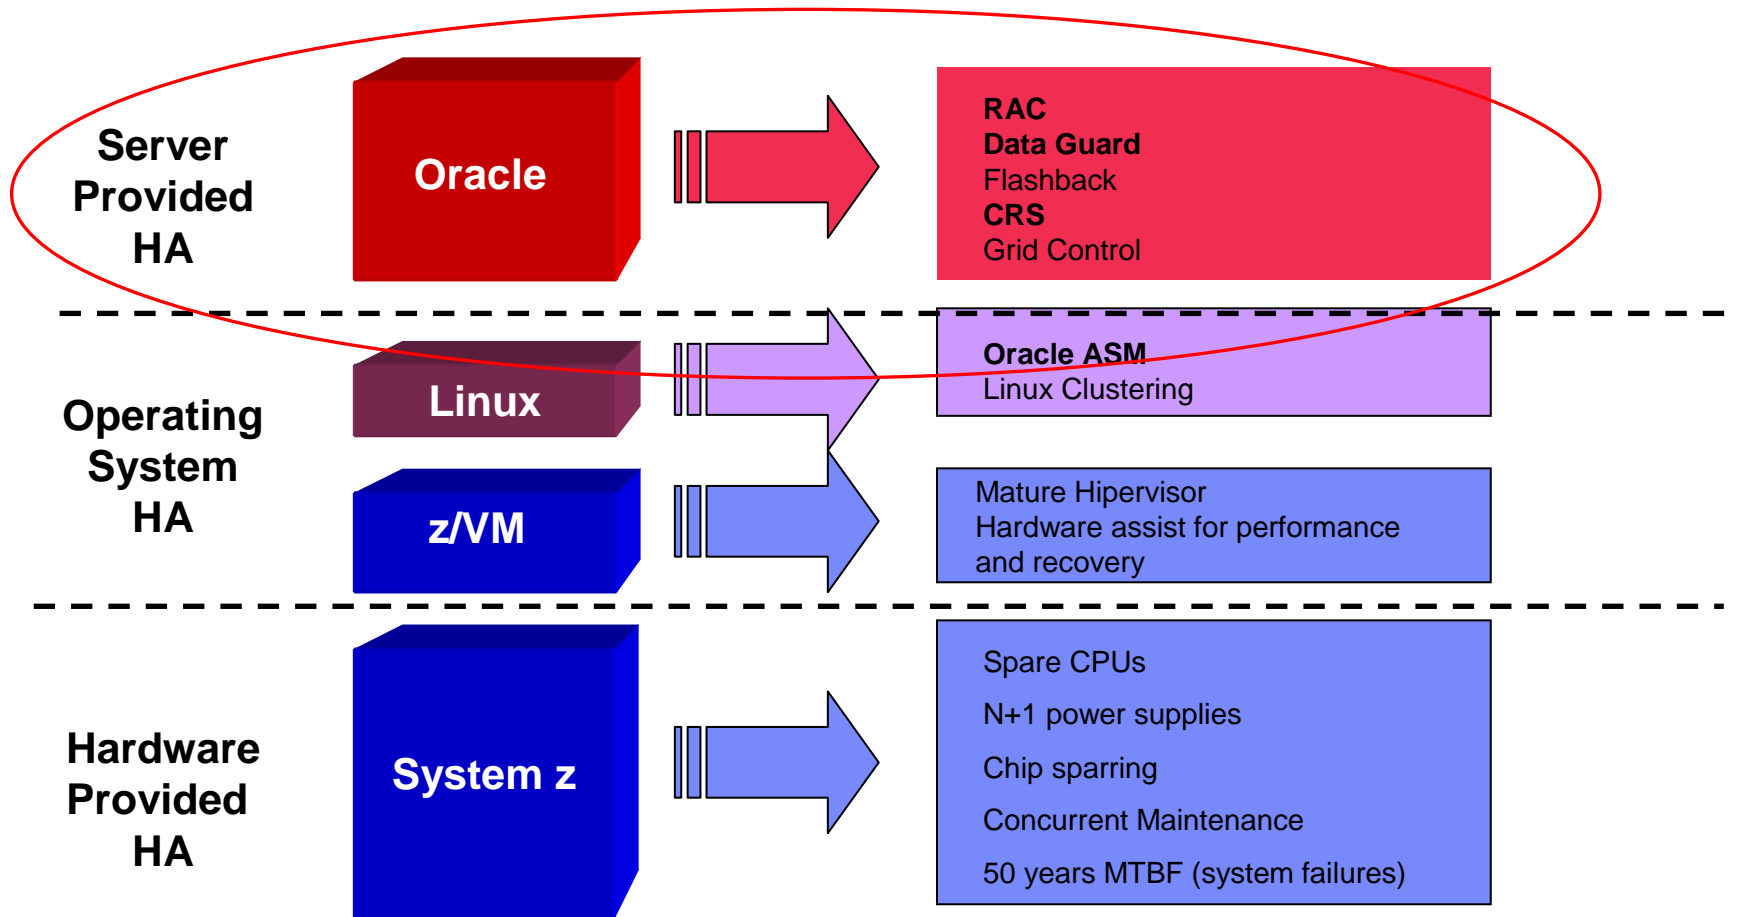


Linux on IBM System z

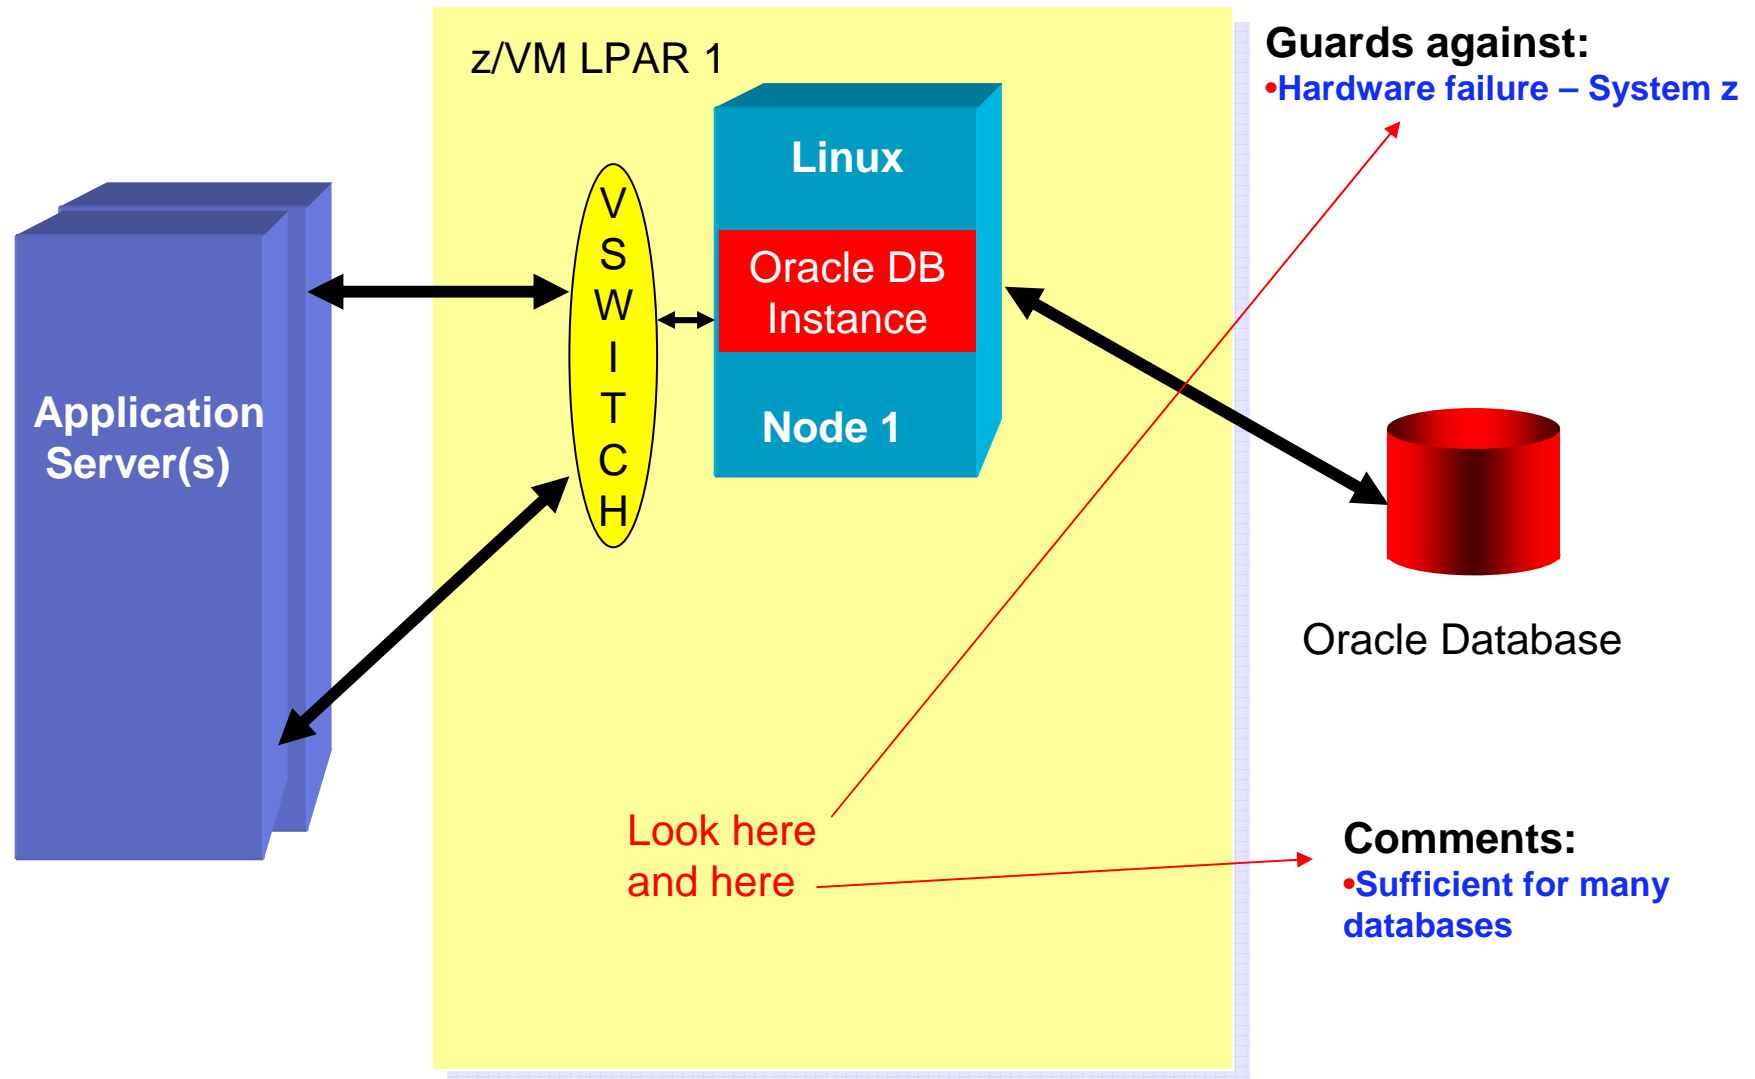
Oracle DB on Linux on z (Loz) through the
eyes  of availability

aka – Begin with the end in mind

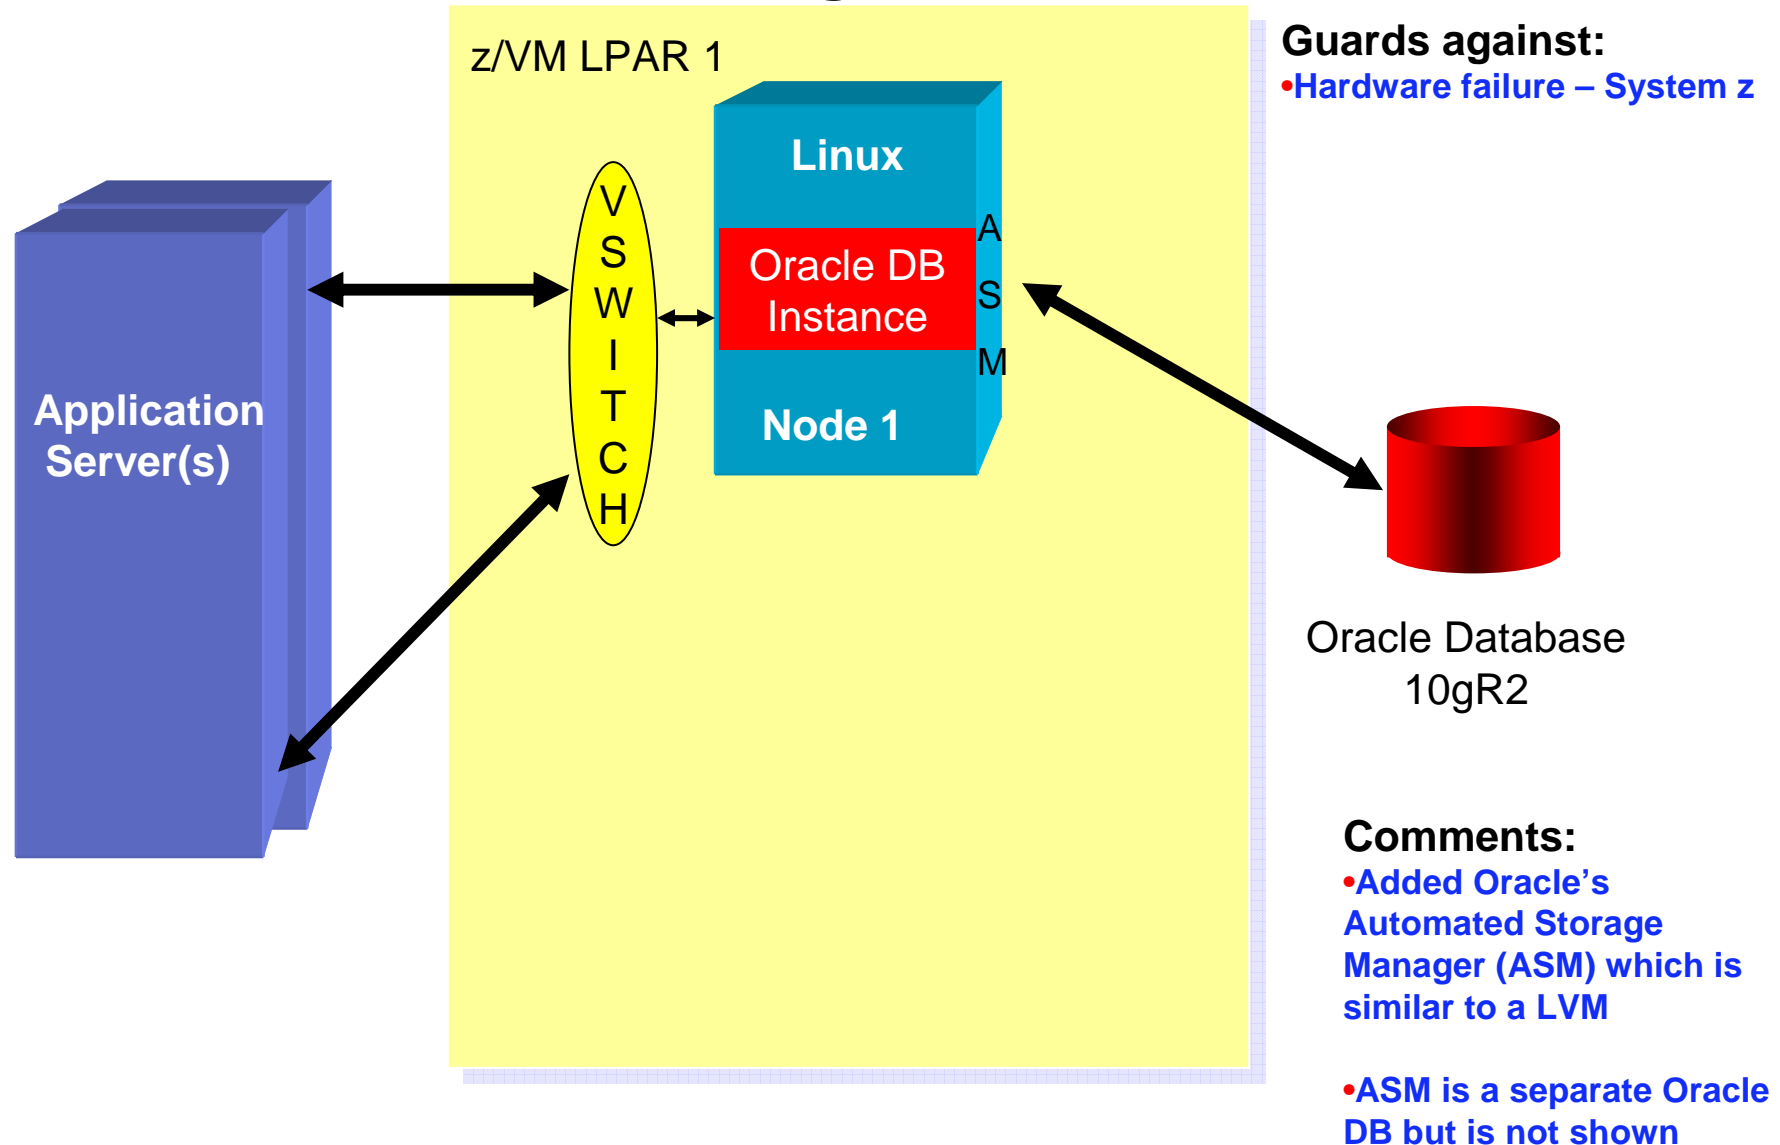
Building Blocks of HA for Oracle DB on Linux for System z



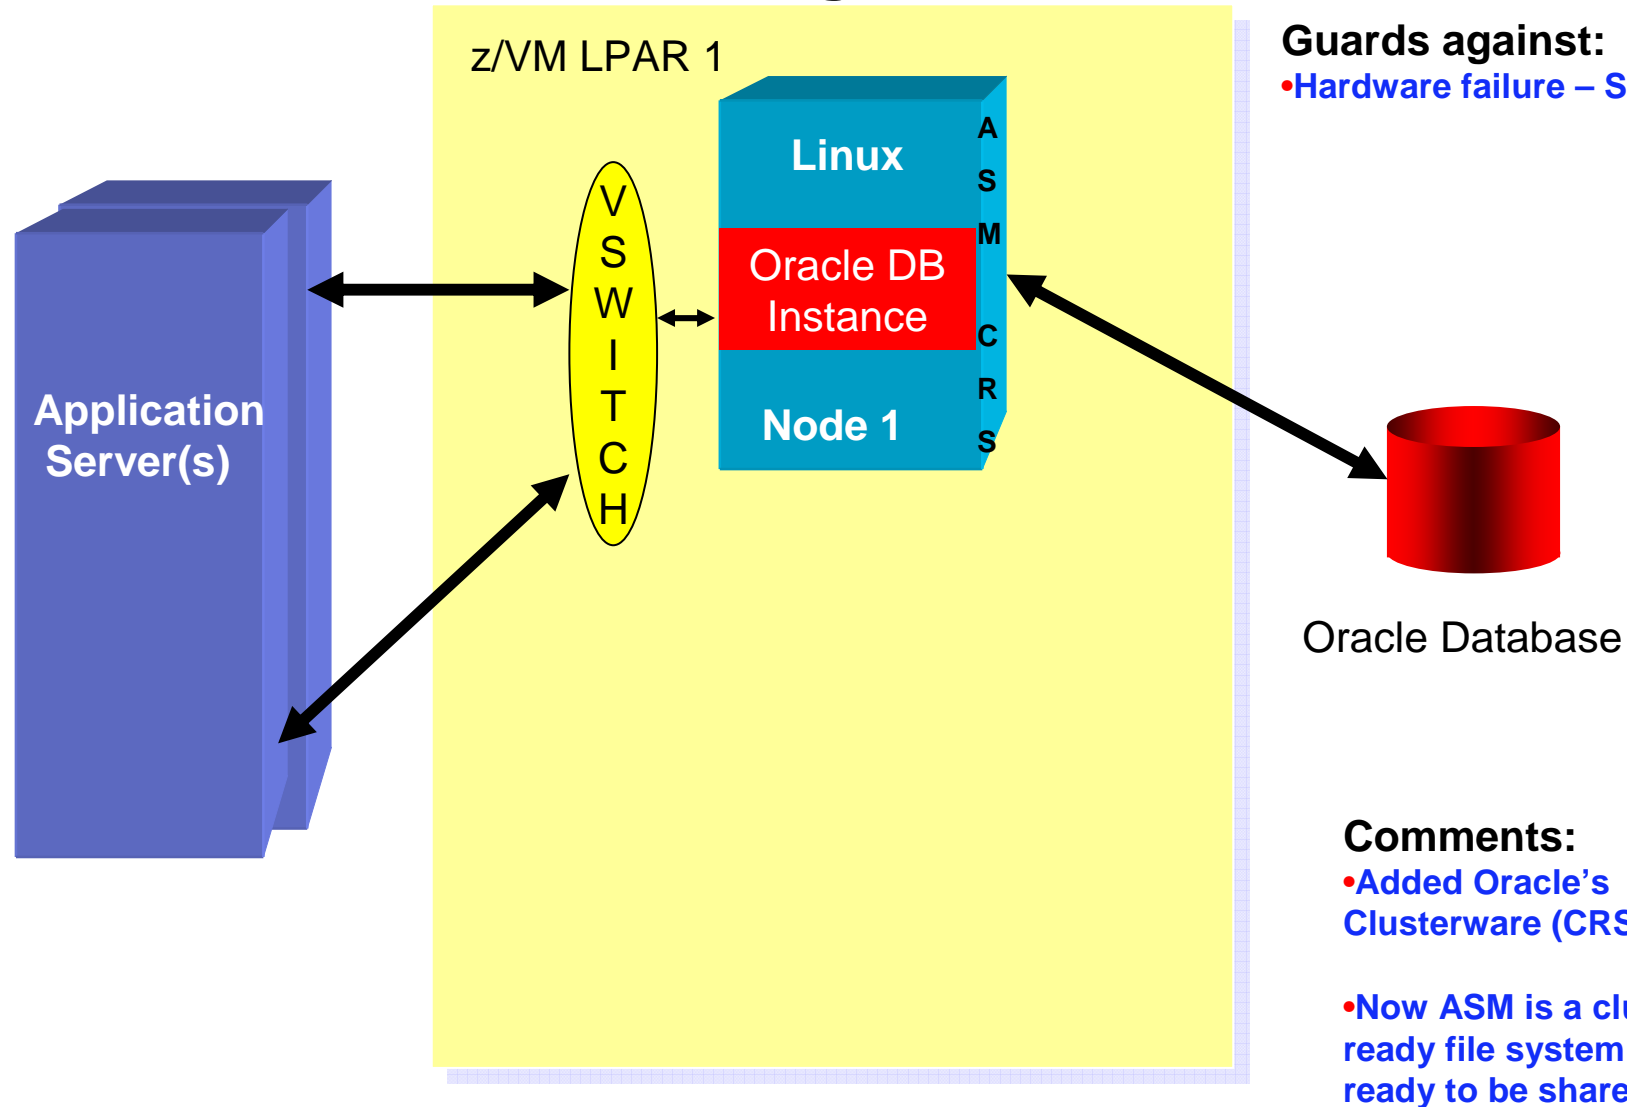
Oracle Database without Oracle MAA



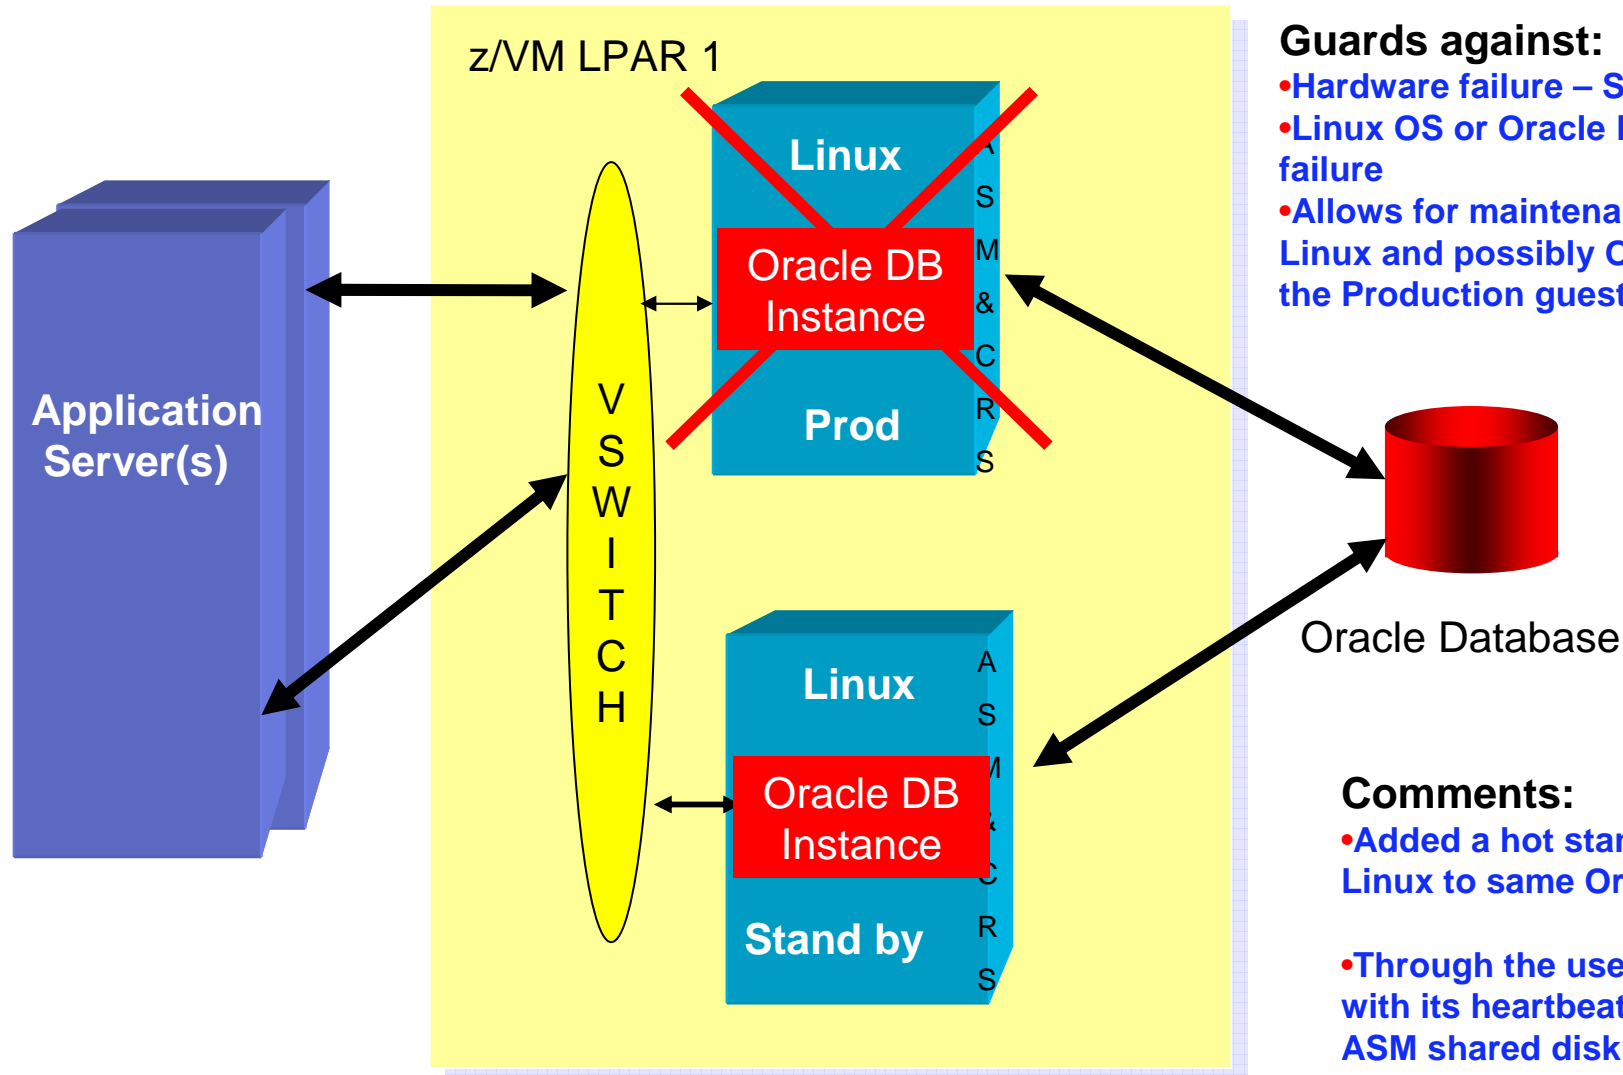
Oracle Database - building Oracle MAA



Oracle Database - building Oracle MAA



Oracle Database - building Oracle MAA - Hot standby



Guards against:

- Hardware failure – System z
- Linux OS or Oracle DB failure
- Allows for maintenance to Linux and possibly Oracle in the Production guest

Oracle Database

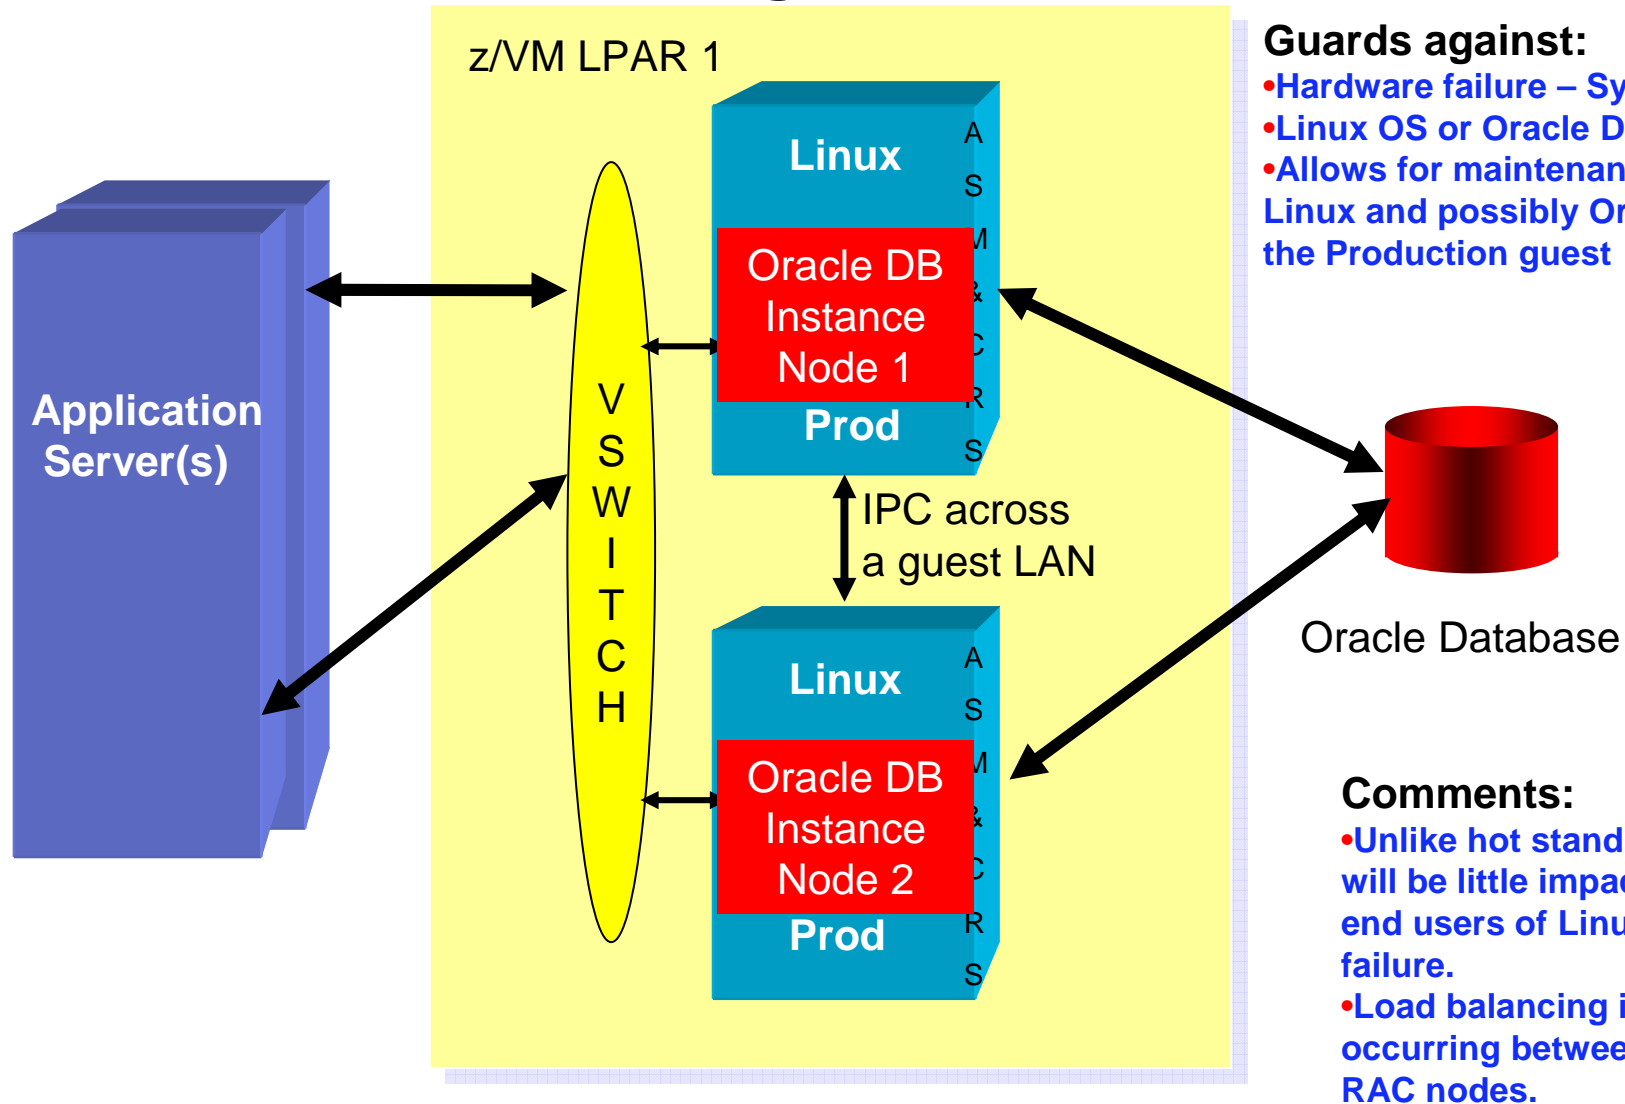
Comments:

- Added a hot standby Linux to same Oracle DB.
- Through the use of CRS, with its heartbeat, and ASM shared disk storage does a failover to the stand by Linux guest.

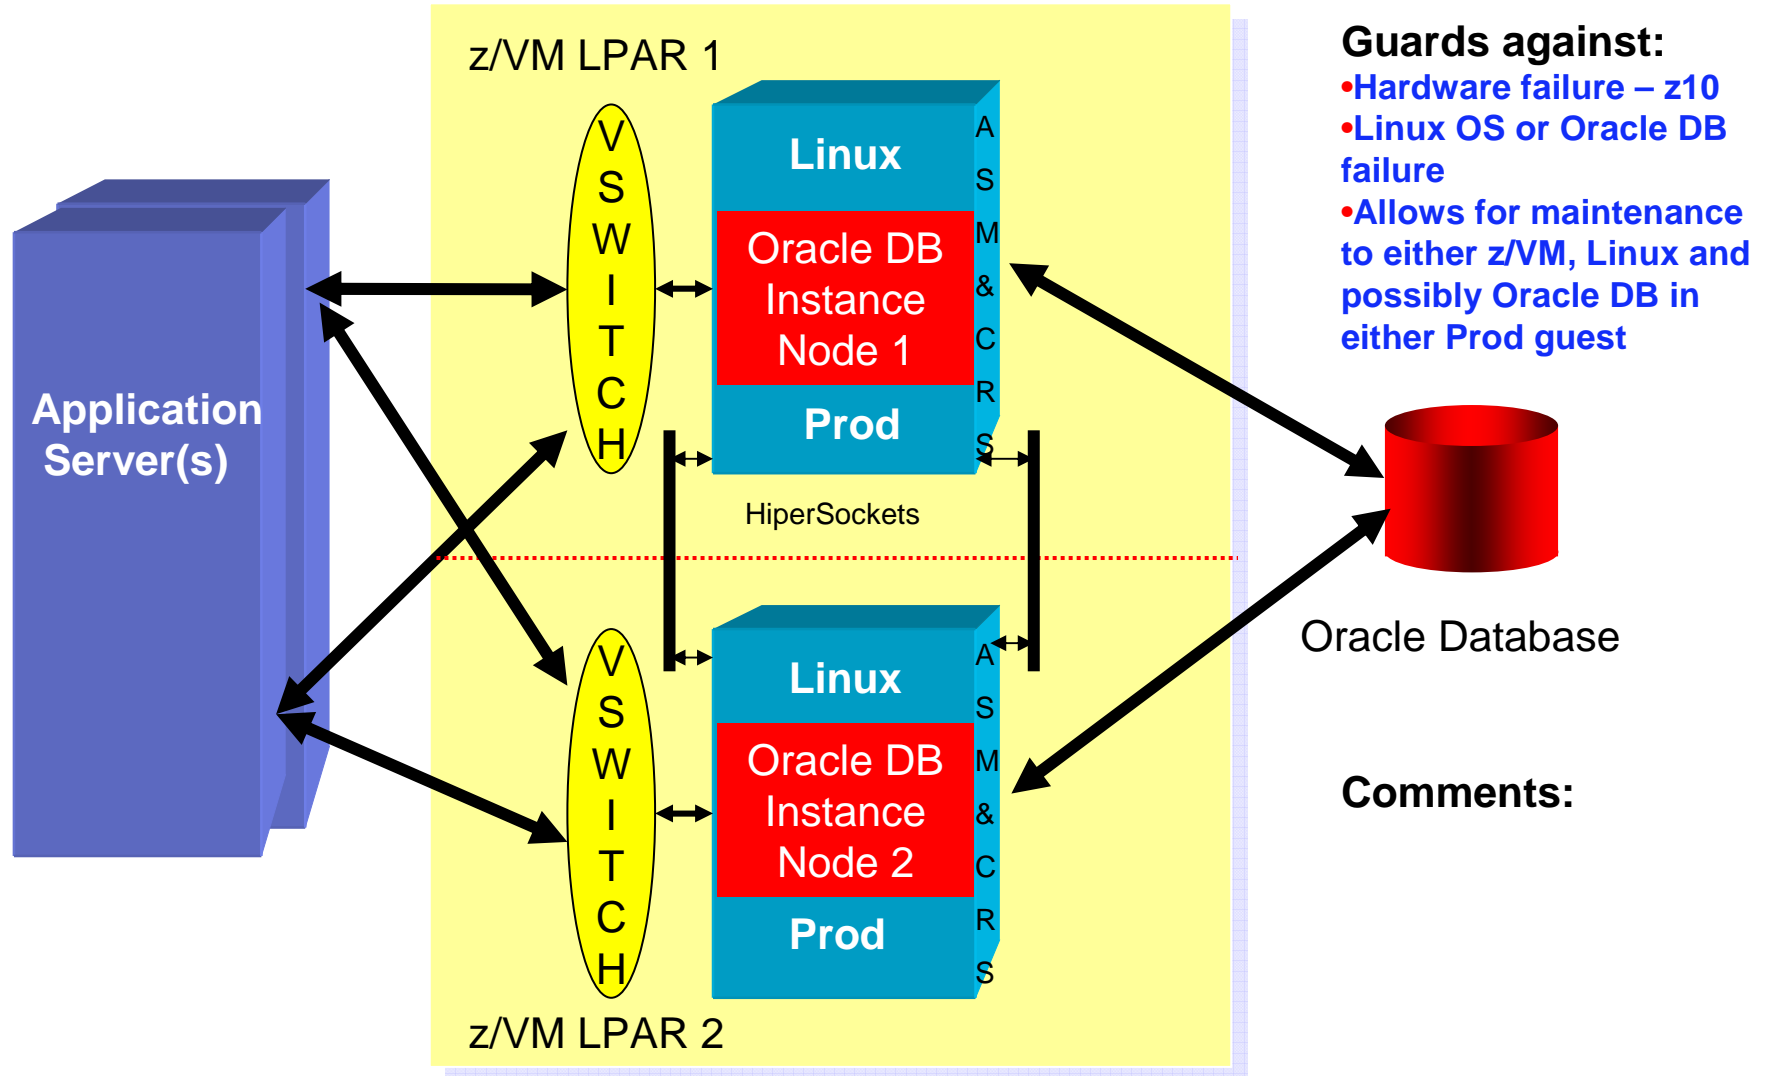
Oracle Hot Stand By Approach Comments

- Can also be accomplished across LPARs using HiperSockets connections.
- Can be accomplished across different System z platforms using appropriate network connectivity.
- Only allowed between Oracle databases using the same binaries (i.e. Linux on z in this case)
- An outage that can affect users can be of a short duration (short duration?)

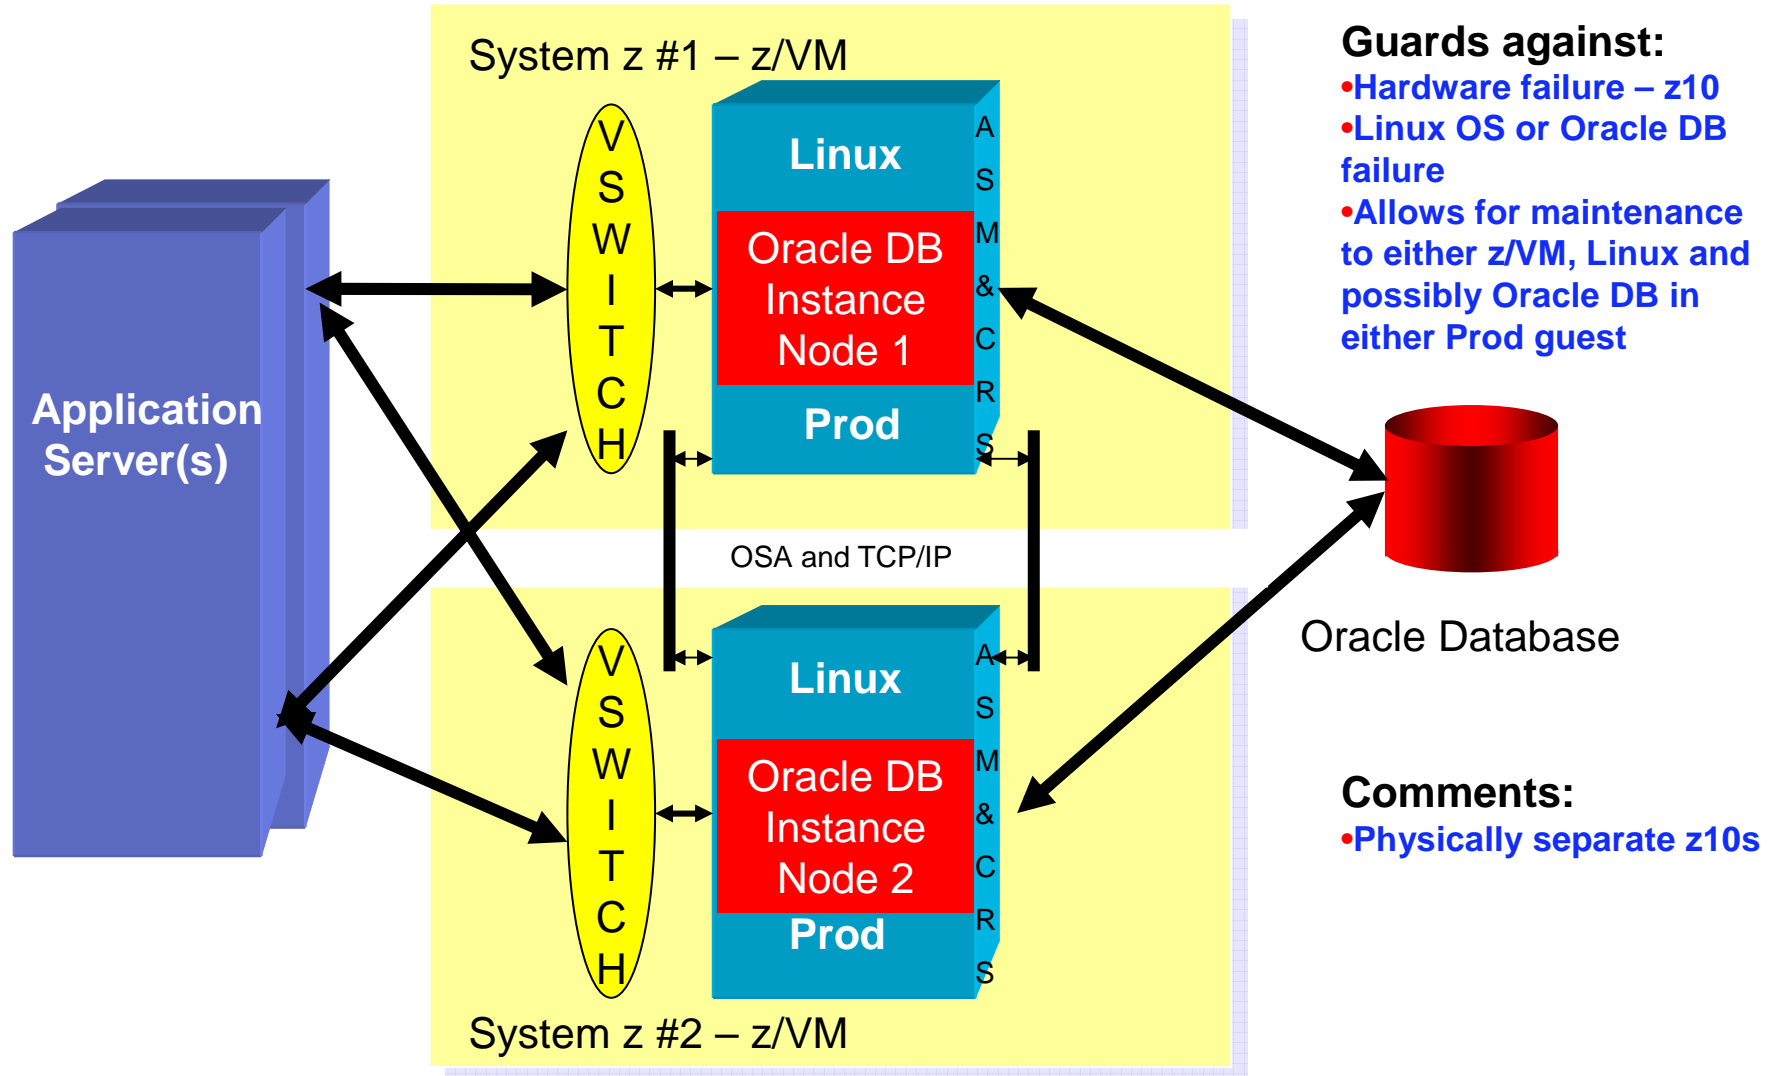
Oracle Database building Oracle MAA - RAC



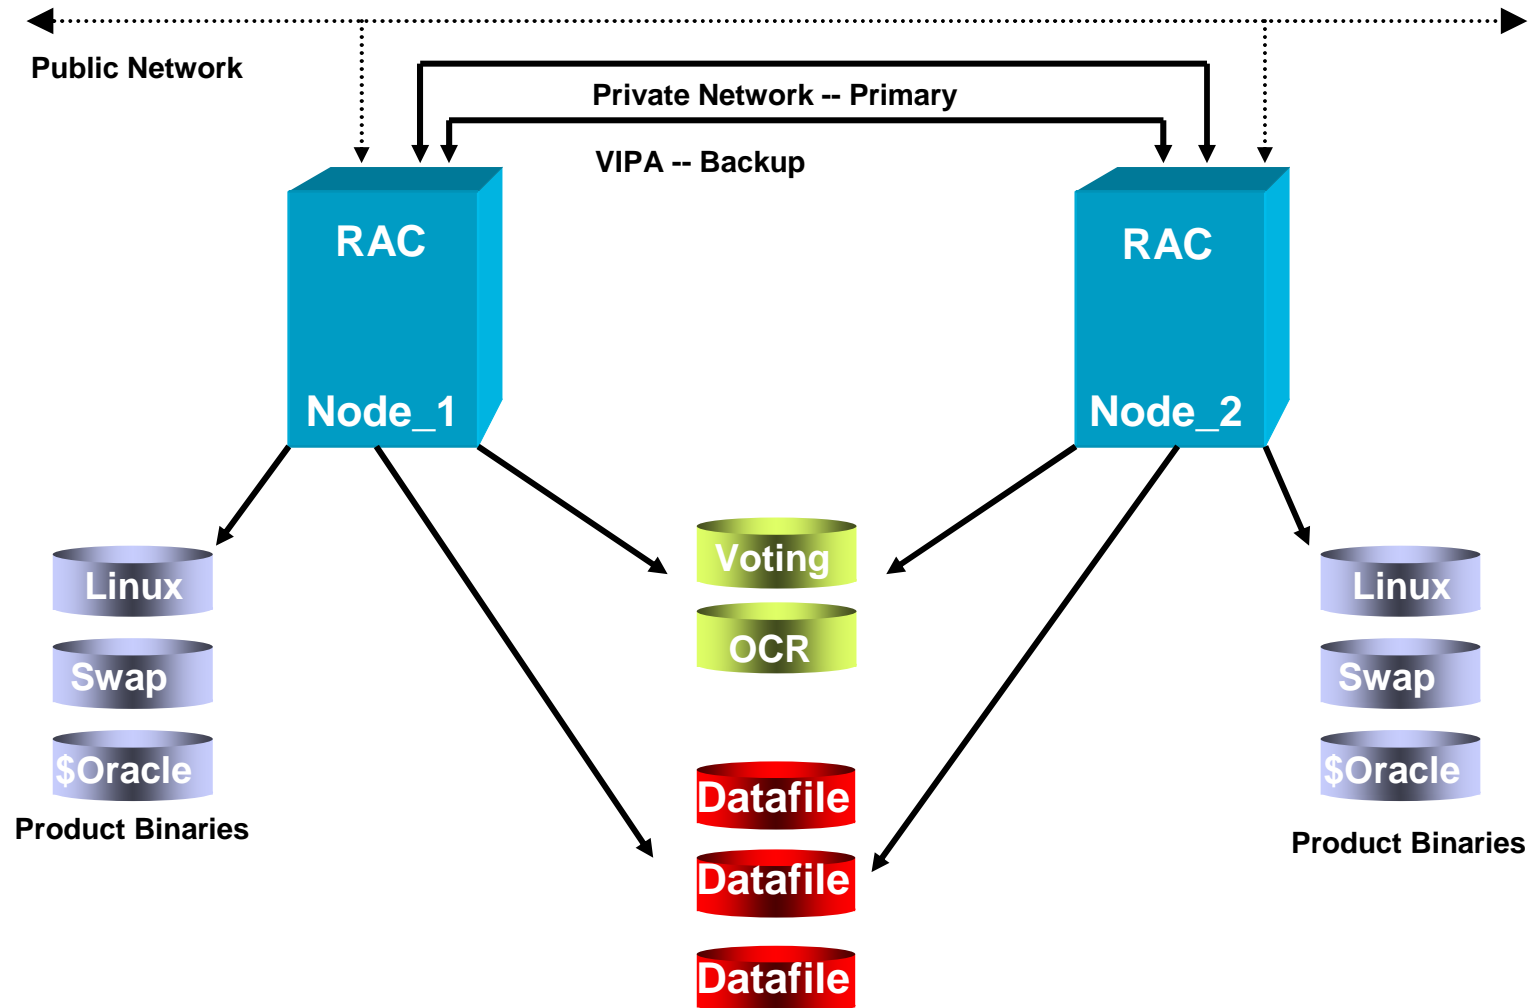
Oracle Database - building Oracle MAA - RAC



Oracle Database - building Oracle MAA - RAC



Overview of Major RAC Components



Oracle RAC as an HA Solution

- RAC implies a HA Solution
 - RAC provides high availability for database instances

- Have you taken into account single points of failure for:
 - Disk failures?
 - IPC Interconnect failures?
 - Are the servers on the same electrical circuit?
 - Are the servers under the same sprinkler?
 - If the nodes are in a different building, is it a single cable run?
 - Did you do appropriate capacity planning for a node or multiple node failures?

- Your availability is as solid as your planning for any platform on which you implement a RAC solution
 - If you plan well, it is a very Highly Available software solution

Deploying RAC for High Availability

- RAC – Real Application Clusters
 - Active/Passive configuration
 - One node processes work
 - The other node waits for the first node to fail
 - Active/Active configuration
 - All nodes process work
 - If any node fails the cluster is re-mastered.
 - Besides availability, RAC can be used for workload distribution
 - All work does not have to go through all nodes
 - Deploy
 - In the same LPAR for test/dev applications
 - Across LPARs for LPAR maintenance or software failures (most common implementation)
 - Across CECs when taking entire systems down is a “common” occurrence

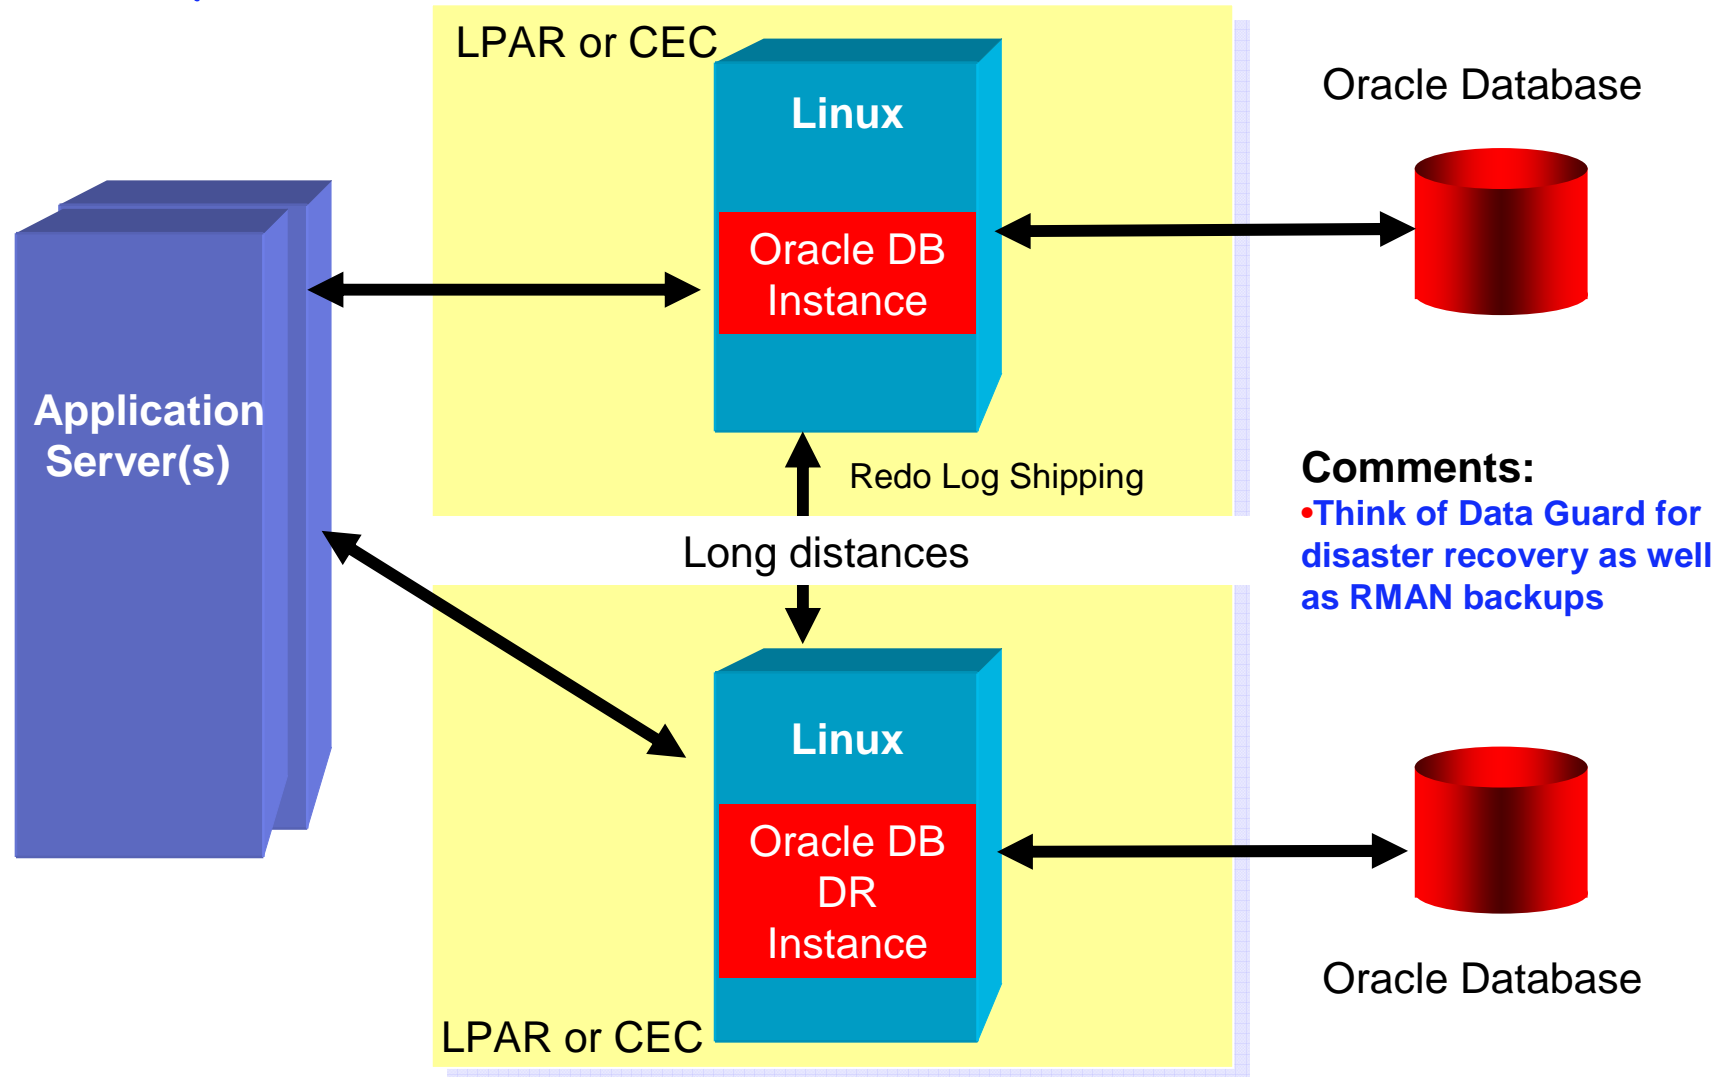
Oracle Standby and Replication Solutions for Disaster Recovery

- Standby – replication to standby database
 - Oracle Data Guard
 - Uses redo log shipping for log apply or SQL Apply
 - Less data transmitted than replication
 - Sync or async
 - Various configurations of logical and physical standby databases

 - Both production and standby databases must be installed from same CD/DVD
 - Support for heterogeneous systems not supported yet
 - Both systems must match for endian, chip set and headers

 - Data Guard generally deployed between CECs

Standby Database - Data Guard



High Availability with Oracle on Linux for System z

- System z - most highly available platform on the planet
 - Attention to detail over decades of engineering
 - Fault Tolerant (HA) design
 - Elimination of single points of failure
 - Driving to 100 years MTBF
- Oracle Maximum Availability Architecture
 - Best Practices based on Oracle technology
 - Best HA/DR in distributed database technology (Forrester, Oct 2006)
 - Spans all Oracle products
 - Constantly evolves with new releases
- Synergistic
 - Continue on your path with Grid using System z
 - Develop a Grid strategy for Oracle on Linux for System z
 - Take advantages of the HA/DR features of IBM and Oracle technologies

Information Sources

- <http://www.ibm.com/redbooks>
 - SG24-6482-00 Experiences with Oracle Database 10g on Linux for zSeries
 - SG24-7191-00 Experiences with Oracle 10gR2 Solutions on Linux for System z
 - SG24-7573-00 Using Oracle Solutions on Linux on System z
 - SG24-7634-00 Experiences with Oracle Solutions on Linux for IBM System z

- <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 System z

- <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

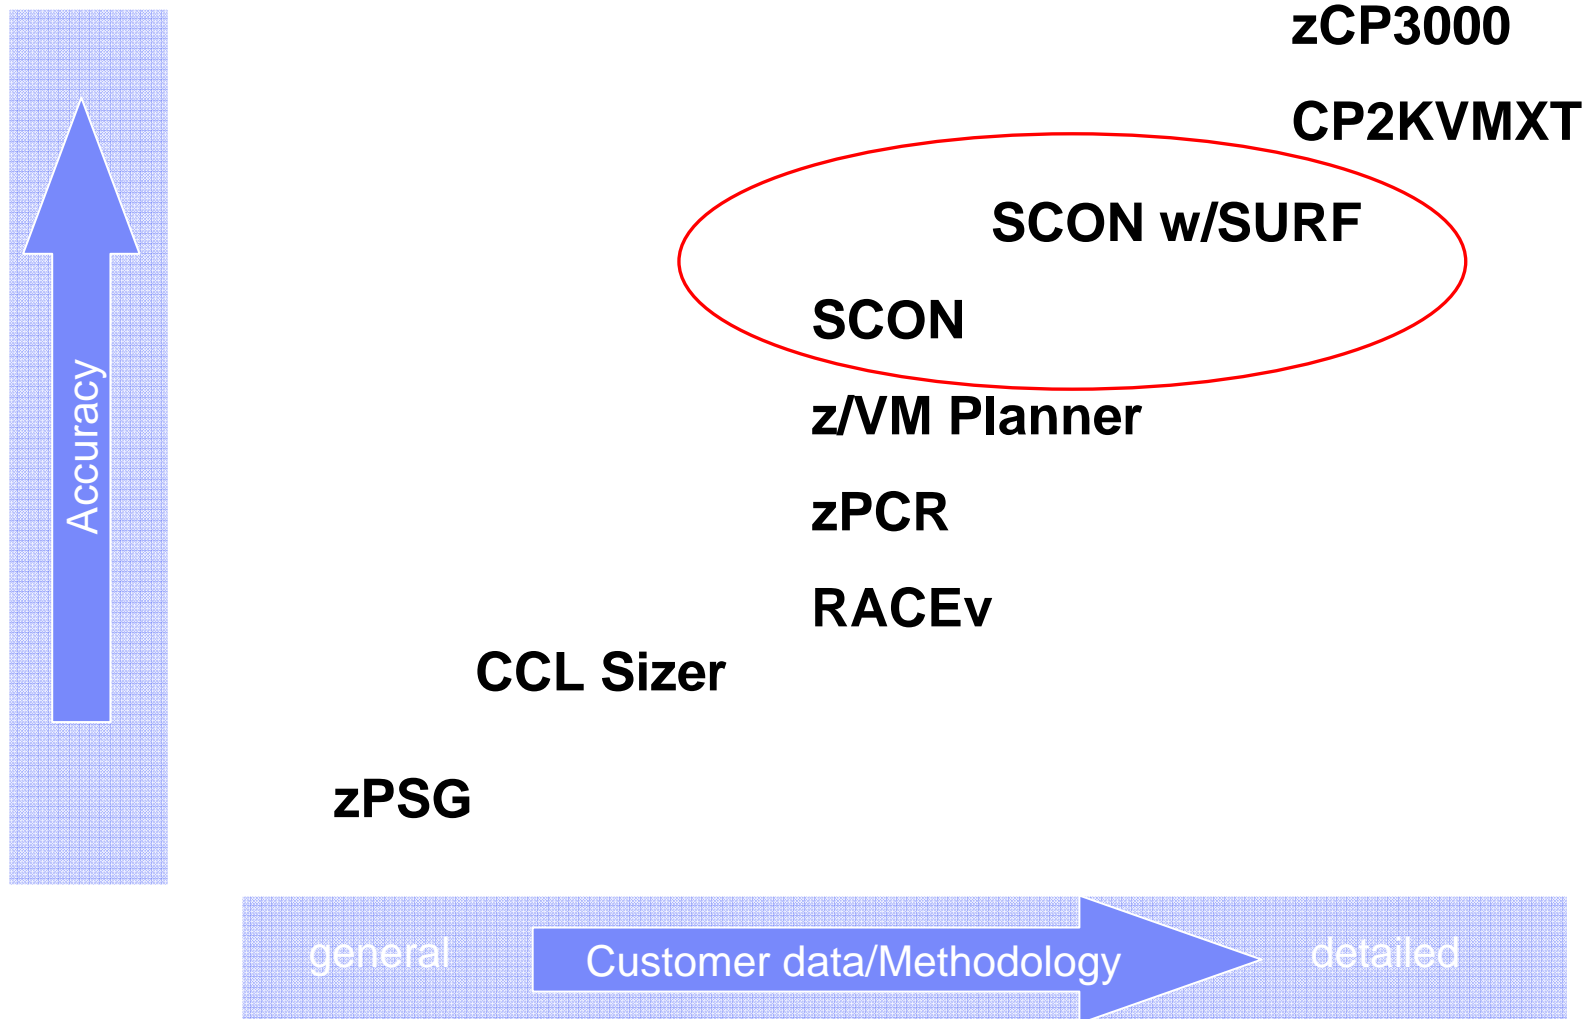


Linux on IBM System z

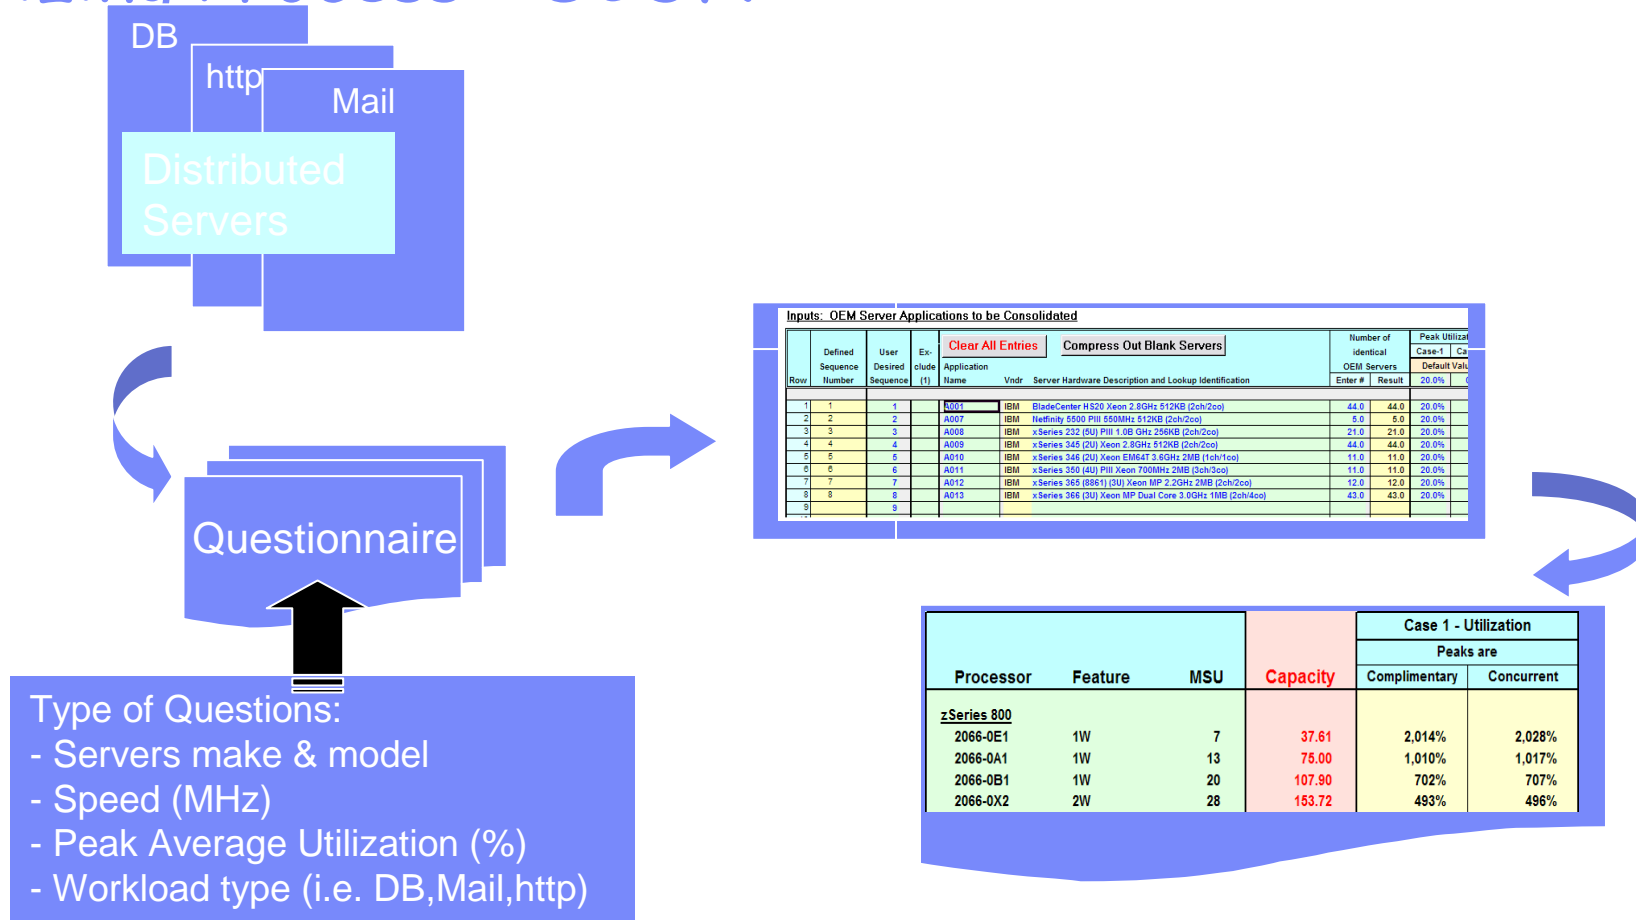
Sizing - the most important step

For PoC or full production

System z Sizing Tools



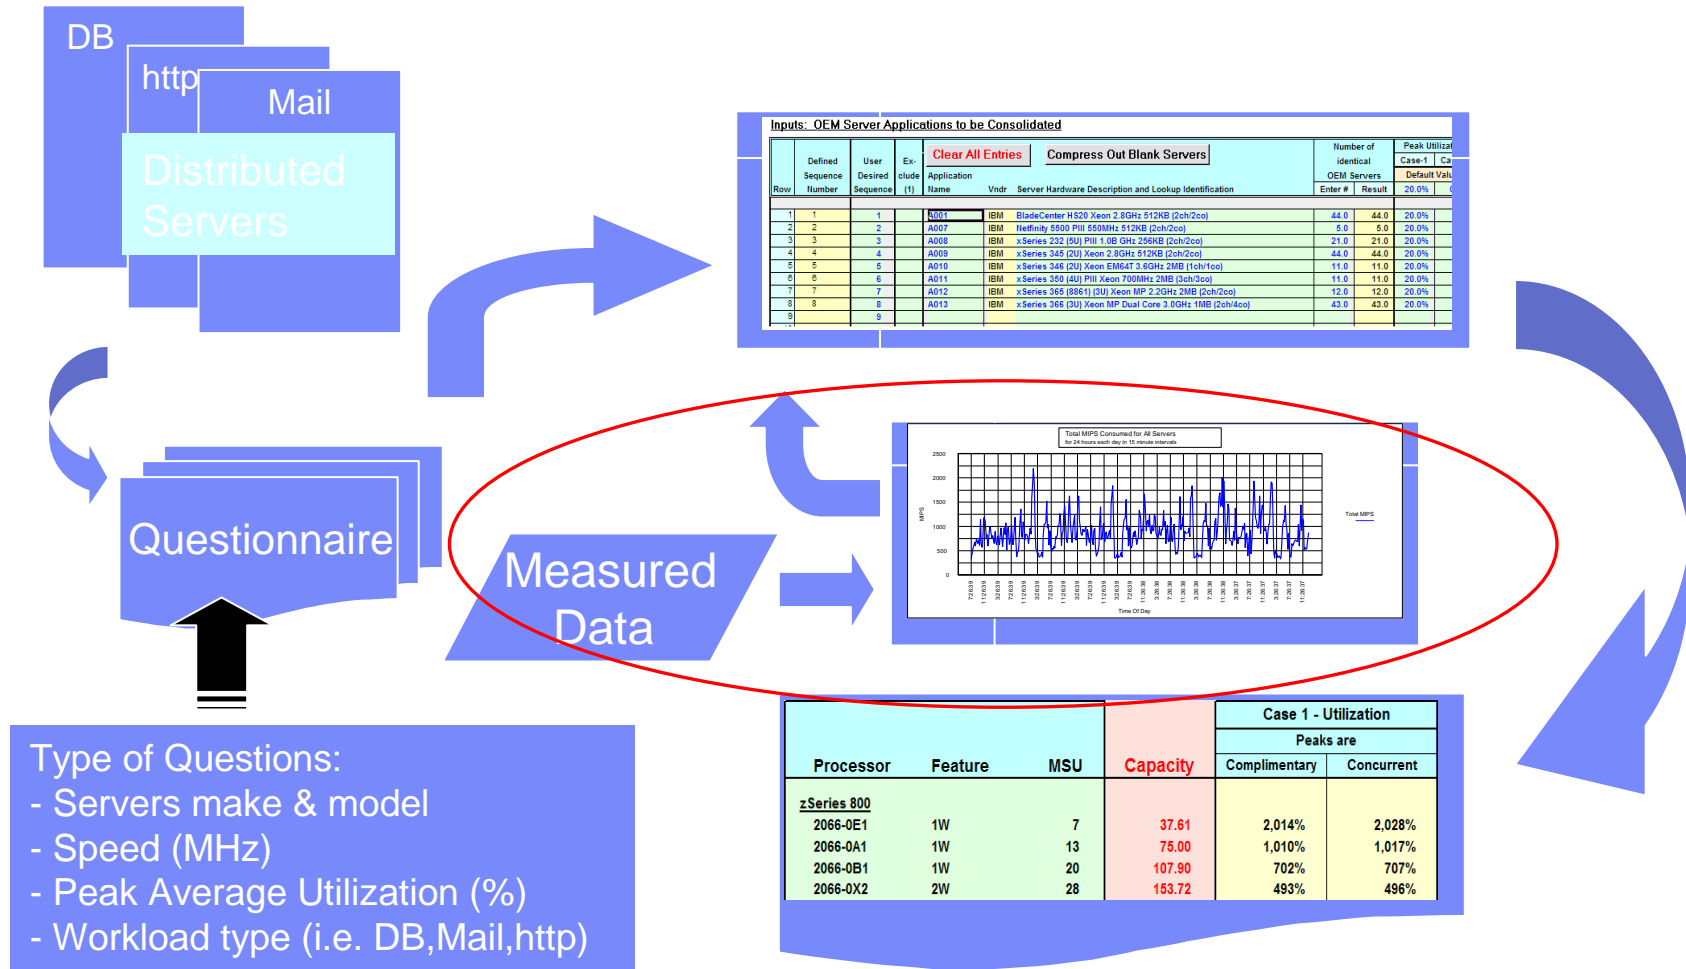
Mainframe Linux Server Consolidation Sizing Process - SCON



Type of Questions:

- Servers make & model
- Speed (MHz)
- Peak Average Utilization (%)
- Workload type (i.e. DB,Mail,http)

Mainframe Linux Server Consolidation Sizing Process - SCON with SURF




Type of Questions:

- Servers make & model
- Speed (MHz)
- Peak Average Utilization (%)
- Workload type (i.e. DB,Mail,http)

Inputs: OEM Server Applications to be Consolidated

| Row | Defined Sequence Number | User Desired Sequence | Exclude (1) | Application Name | Vndr | Server Hardware Description and Lookup Identification | Number of identical OEM Servers | | Peak Utilizat | |
|-----|-------------------------|-----------------------|-------------|------------------|------|---|---------------------------------|--------|---------------|----|
| | | | | | | | Enter # | Result | Case-1 | Ca |
| 1 | 1 | 1 | | A001 | IBM | BladeCenter H520 Xeon 2.8GHz 512KB (2ch/2co) | 44.0 | 44.0 | 20.0% | |
| 2 | 2 | 2 | | A007 | IBM | Netfinity 5500 PIII 850MHz 512KB (2ch/2co) | 5.0 | 5.0 | 20.0% | |
| 3 | 3 | 3 | | A008 | IBM | xSeries 232 (SU) PIII 1.08 GHz 256KB (2ch/2co) | 21.0 | 21.0 | 20.0% | |
| 4 | 4 | 4 | | A009 | IBM | xSeries 345 (2U) Xeon 2.8GHz 512KB (2ch/2co) | 44.0 | 44.0 | 20.0% | |
| 5 | 5 | 5 | | A010 | IBM | xSeries 346 (2U) Xeon EM64T 3.0GHz 2MB (1ch/1co) | 11.0 | 11.0 | 20.0% | |
| 6 | 6 | 6 | | A011 | IBM | xSeries 355 (4U) PIII Xeon 700MHz 2MB (2ch/2co) | 11.0 | 11.0 | 20.0% | |
| 7 | 7 | 7 | | A012 | IBM | xSeries 365 (8851) (SU) Xeon MP 2.2GHz 2MB (2ch/2co) | 12.0 | 12.0 | 20.0% | |
| 8 | 8 | 8 | | A013 | IBM | xSeries 366 (SU) Xeon MP Dual Core 3.0GHz 1MB (2ch/4co) | 43.0 | 43.0 | 20.0% | |
| 9 | 9 | 9 | | | | | | | | |



| Processor | Feature | MSU | Capacity | Case 1 - Utilization | |
|--------------------|---------|-----|----------|----------------------|------------|
| | | | | Complimentary | Concurrent |
| zSeries 800 | | | | | |
| 2066-0E1 | 1W | 7 | 37.61 | 2,014% | 2,028% |
| 2066-0A1 | 1W | 13 | 75.00 | 1,010% | 1,017% |
| 2066-0B1 | 1W | 20 | 107.90 | 702% | 707% |
| 2066-0X2 | 2W | 28 | 153.72 | 493% | 496% |

Oracle DB Memory sizing

- Obtain Oracle SGA and PGA sizes from all database instances
 - Prefer Advisory sizes from multiple AWR reports.
- Calculate guest(s) virtual storage size (assume MB):
(SGA + PGA) + 256 MB for ASM + 512 MB for Linux* **
- Assume the sum all of the guest virtual sizes for production equals p and the sum of all guest virtual sizes for dev/qa/training equals t.
Real memory for guests = $p/.66 + t/ (.33)$ for z/VM memory over commit
 - Assumes multiple guests are involved. Not correct for a single guest
- System z memory = real memory for guests + memory for z/VM and expanded storage.

*Increase estimate when Oracle SGA is large and there are expected to be hundreds of dedicated server connections

** A large overall virtual storage requirement may result in larger Page Tables in Linux which require storage

PGA Memory Advisory from an AWR report

PGA Memory Advisory

- When using Auto Memory Mgmt, minimally choose a pga_aggregate_target value where Estd PGA Overalloc Count is 0

| PGA Target Est (MB) | Size Factr | W/A MB Processed | Estd Extra W/A MB Read/ Written to Disk | Estd PGA Cache Hit % | Estd PGA Overalloc Count |
|---------------------|------------|------------------|---|----------------------|--------------------------|
| 896 | 0.13 | 148,138.91 | 182,994.64 | 45.00 | 1,297 |
| 1,792 | 0.25 | 148,138.91 | 173,054.91 | 46.00 | 1,197 |
| 3,584 | 0.50 | 148,138.91 | 30,487.16 | 83.00 | 0 |
| 5,376 | 0.75 | 148,138.91 | 30,487.16 | 83.00 | 0 |
| 7,168 | 1.00 | 148,138.91 | 29,701.39 | 83.00 | 0 |
| 8,602 | 1.20 | 148,138.91 | 12,032.42 | 92.00 | 0 |
| 10,035 | 1.40 | 148,138.91 | 12,032.42 | 92.00 | 0 |
| 11,469 | 1.60 | 148,138.91 | 12,032.42 | 92.00 | 0 |
| 12,902 | 1.80 | 148,138.91 | 12,032.42 | 92.00 | 0 |
| 14,336 | 2.00 | 148,138.91 | 12,032.42 | 92.00 | 0 |
| 21,504 | 3.00 | 148,138.91 | 12,032.42 | 92.00 | 0 |
| 28,672 | 4.00 | 148,138.91 | 12,032.42 | 92.00 | 0 |
| 43,008 | 6.00 | 148,138.91 | 12,032.42 | 92.00 | 0 |
| 57,344 | 8.00 | 148,138.91 | 12,032.42 | 92.00 | 0 |

It appears that the allocated memory of 7,168 MB is twice as large as required.

SGA Target Advisory from an AWR report

SGA Target Advisory

| SGA Target Size (M) | SGA Size Factor | Est DB Time (s) | Est Physical Reads |
|---------------------|-----------------|-----------------|--------------------|
| 2,304 | 0.25 | 317,428 | 14,503,025 |
| 4,608 | 0.50 | 282,694 | 11,631,530 |
| 6,912 | 0.75 | 270,413 | 10,965,119 |
| 9,216 | 1.00 | 263,535 | 10,396,434 |
| 11,520 | 1.25 | 258,791 | 10,003,449 |
| 13,824 | 1.50 | 255,418 | 9,725,864 |
| 16,128 | 1.75 | 252,915 | 9,517,935 |
| 18,432 | 2.00 | 252,150 | 9,454,517 |

It appears that the allocated memory of 9,216 MB might be reasonable.

Threads for dedicated servers

Decide on number of dedicated threads and multiply that by 4.5 MB for required real memory to include in guest sizing.

The logons current below may give a hint about number of threads.

Instance Activity Stats - Absolute Values

- Statistics with absolute values (should not be diffed)

| Statistic | Begin Value | End Value |
|----------------------------|-------------|-----------|
| session cursor cache count | 20,573 | 21,027 |
| opened cursors current | 186 | 91 |
| workarea memory allocated | 870,391 | 3,575 |
| logons current | 124 | 30 |

Obvious comments for sizing

- Garbage in, garbage out.
- Choose appropriate time frames that represent reasonable capacity usage
- Do not make guesses about the sizing input
- We must get the CPU capacity, I/O subsystem, and the memory at the correct levels before any testing starts
- Engage a System z - Oracle specialist to assist with sizing



Linux on IBM System z

Proof of Concept (PoC)

PoC part 1

- Engage a System z - Oracle specialist to assist with PoC planning
- Attend education
- Obtain IFLs and memory as per the sizing process
 - No zIIPs, zAAPs or CP's for this environment
 - Choose I/O subsystem (ECKD or SCSI)
- Install z/VM and it's performance tools
- Install Linux
 - Choose certified levels of SUSE or Red Hat
Go to support.oracle.com
 - Verify required Oracle modules have been installed
RPM checkers available
- Use Orion to validate the I/O subsystem even before a Oracle database is installed
 - Performs Oracle like I/O

Storage - Testing with ORION - 1

ORION Simulates Oracle reads and writes, without having to create a database and helps to isolate I/O issues. When a database is optimally configured you can expect to get up to 95% of the throughput of Orion.

```
./orion_zlinux -run oltp -testname mytest -num_disks 2 -duration 30 -simulate raid0
```

```
ORION VERSION 11.2.0.0.1
```

```
Commandline: -run oltp -testname mytest -num_disks 2 -duration 30 -simulate raid0
```

```
This maps to this test: Test: mytest
```

```
Small IO size: 8 KB Large IO size: 1024 KB
```

```
IO Types: Small Random IOs, Large Random IOs
```

```
Simulated Array Type: RAID 0 Stripe Depth: 1024 KB
```

```
Write: 0% Cache Size: Not Entered
```

```
Duration for each Data Point: 30 seconds
```

```
Small Columns:, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 32, 34, 36, 38, 40
```

```
Large Columns:, 0 Total Data Points: 22
```

```
Name: /dev/dasdq1 Size: 2461679616
```

```
Name: /dev/dasdr1 Size: 2461679616
```

```
2 FILES found.
```

```
Maximum Small IOPS=5035 @ Small=40 and Large=0
```

```
Minimum Small Latency=0.55 @ Small=2 and Large=0
```

Storage - Testing with ORION - 2

```
-run oltp -testname mytest -num_disks 2 -duration 30 -simulate raid0
```

This maps to this test:

Test: mytest

Small IO size: 8 KB Large IO size: 1024 KB

IO Types: Small Random IOs, Large Random IOs

Simulated Array Type: RAID 0 Stripe Depth: 1024 KB

Write: 0%

Cache Size: Not Entered

Duration for each Data Point: 30 seconds

Small Columns: 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 32, 34, 36, 38, 40

Large Columns: 0

Total Data Points: 22

Name: /dev/sda1 Size: 10737401856

Name: /dev/sdb1 Size: 10737401856

2 FILES found.

Maximum Small **IOPS=24945** @ Small=24 and Large=0

Minimum Small Latency=0.60 @ Small=12 and Large=0

Download - <http://www.oracle.com/technology/software/tech/orion/index.html>

Storage - Testing with Orion - 3

- Be careful of the options you choose. **The writes are destructive.**
- Perform Orion testing BEFORE installing the Oracle database to validate the I/O subsystem



OR



Moving data is like moving water – must have adequate flow throughout

AWR - I/O statistics

| Tablespace | Filename | Reads | Av Reads/s | Av Rd(ms) | Av Blks/Rd | Writes | Av Writes/s | Buffer Waits | Av Buf Wt(ms) |
|------------|----------|------------|------------|-----------|------------|-----------|-------------|--------------|---------------|
| | | 10,790,012 | 222 | 12.28 | 2.25 | 4,845,015 | 100 | 411 | 54.38 |
| | | 10,311,731 | 212 | 11.87 | 2.37 | 4,758,474 | 98 | 591 | 63.46 |
| | | 2,030,575 | 42 | 22.91 | 1.23 | 2,551,704 | 53 | 3,857 | 141.84 |
| | | 1,190,077 | 24 | 27.58 | 1.21 | 1,477,830 | 30 | 2,897 | 23.93 |
| | | 1,143,880 | 24 | 19.50 | 1.18 | 1,593,814 | 33 | 2,904 | 87.73 |



PoC part 2

- Install Oracle database – 10gR2
 - Consider starting with Oracle ASM versus LVM ext3 files
 - If using ext3 the verify Oracle init.ora has
`filesystemio_options = setall`
`disk_asynch_io=true`
to eliminate Linux double caching which wastes storage and CPU resources
- Create appropriate disk multipathing
 - Different for SCSI and ECKD
 - Consider running Orion again to test multipathing
- Load database from another Oracle database source
 - Use transportable tablespace or database for metadata when endian formats are the same
<http://en.wikipedia.org/wiki/Endian>
 - Additional steps, like rman conversions, are required for unlike endian formats
 - Import/export may be required when source database is older than 10gR2
 - **Recreate statistics for optimizer use**

Endian formats

```

SQL> COLUMN PLATFORM_NAME FORMAT A32;
SQL> SELECT * FROM V$TRANSPORTABLE_PLATFORM;
PLATFORM_ID PLATFORM_NAME                                ENDIAN_FORMAT
-----
1 Solaris[tm] OE (32-bit)                                Big
2 Solaris[tm] OE (64-bit)                                Big
7 Microsoft Windows IA (32-bit)                          Little
10 Linux IA (32-bit)                                      Little
6 AIX-Based Systems (64-bit)                              Big
3 HP-UX (64-bit)                                          Big
5 HP Tru64 UNIX                                           Little
4 HP-UX IA (64-bit)                                       Big
11 Linux IA (64-bit)                                      Little
15 HP Open VMS                                           Little
8 Microsoft Windows IA (64-bit)                          Little
9 IBM zSeries Based Linux                               Big
13 Linux x86 64-bit                                       Little
16 Apple Mac OS                                           Big
12 Microsoft Windows x86 64-bit                          Little
17 Solaris Operating System (x86)                          Little
18 IBM Power Based Linux                                   Big
20 Solaris Operating System (x86-64)                      Little
19 HP IA Open VMS                                         Little

```


PoC part 3

- Run PoC testing
 - Collect performance data by enabling:
 - z/VM Performance Toolkit
 - Note that you must now think about virtualization versus dedicated resources
 - sar and iostat data from the Linux on z guest(s)
 - AWR reports from the Oracle database
 - Review performance reports
 - z/VM
 - Understand CPU, memory, and paging consumption for the LPAR
 - Review virtual machine consumption of resources
 - Evaluate I/O performance (ECKD only)
 - Verify VDISK usage
 - Linux using sar and iostat
 - CPU, memory, swapping, and I/O performance for each guest
 - Oracle AWR report
 - I/O performance
 - SGA and PGA usage via automatic memory management (see previous chart)
 - Normal DBA tuning review
 - Review for performing SQL
 - Locking
- Rerun PoC if changes are made
 - Does the PoC validate the initial sizing?

PoC part 4

- Think in terms of virtualization – different mind set
 - Does that Oracle database require all of that memory it has in the non-virtualized environment?
 - Should you have a active/passive/stand by setup in the same z/VM?
 - Optimize use of resources
 - Did the guests get properly prioritized with respect to other guests?
 - What workloads are peaking at the same time
 - CPU peak
 - Memory load
 - I/O subsystem
 - DBA's, Linux admins, and z/VM sys progs must work as a team

AWR - other statistics

Operating System Statistics

Top 5 Timed Events

| Event | Waits | Time(s) | Avg Wait(ms) | % Total Call Time | Wait Class |
|--------------------------|-----------|---------|--------------|-------------------|------------|
| db file sequential read | 6,073,284 | 74,443 | 12 | 35.5 | User I/O |
| CPU time | | 64,668 | | 30.8 | |
| log file sequential read | 173,131 | 8,093 | 47 | 3.9 | System I/O |
| log file parallel write | 189,657 | 3,668 | 19 | 1.7 | System I/O |
| gc current grant 2-way | 2,697,994 | 2,469 | 1 | 1.2 | Cluster |

| Statistic | Total |
|------------------------|----------------|
| BUSY_TIME | 6,905,787 |
| IDLE_TIME | 1,288,223 |
| IOWAIT_TIME | 886,823 |
| NICE_TIME | 4,077 |
| SYS_TIME | 268,498 |
| USER_TIME | 6,493,552 |
| LOAD | 4 |
| RSRC_MGR_CPU_WAIT_TIME | 0 |
| PHYSICAL_MEMORY_BYTES | 33,711,116,288 |
| NUM_CPUS | 4 |

SQL ordered by Elapsed Time

- Resources reported for PL/SQL code includes the resources used by all SQL statements called by the code.
- % Total DB Time is the Elapsed Time of the SQL statement divided into the Total Database Time multiplied by 100

| Elapsed Time (s) | CPU Time (s) | Executions | Elap per Exec (s) | % Total DB Time | SQL Id | SQL Module | SQL Text |
|------------------|--------------|------------|-------------------|-----------------|------------|------------|-------------------------------------|
| 180,654 | 58,111 | 12 | 15054.53 | 86.10 | 5000000000 | PLACCUP | BEGIN ... |
| 88,004 | 35,905 | 174 | 505.77 | 41.94 | 5000000000 | PLACCUP | BEGIN :1 := ... |
| 25,374 | 3,294 | 167 | 151.94 | 12.09 | 5000000000 | PLACCUP | INSERT ALL WHEN ... |
| 16,124 | 2,939 | 174 | 92.67 | 7.68 | 5000000000 | PLACCUP | INSERT INTO ... |
| 12,080 | 5,048 | 3,519 | 3.43 | 5.76 | 5000000000 | PLACCUP | INSERT INTO ... |
| 8,754 | 4,475 | 167 | 52.42 | 4.17 | 5000000000 | PLACCUP | UPDATE ... |
| 8,313 | 1,293 | 167 | 49.78 | 3.96 | 5000000000 | PLACCUP | INSERT INTO ... |
| 6,177 | 1,484 | 167 | 36.99 | 2.94 | 5000000000 | PLACCUP | INSERT INTO ... |
| 5,545 | 2,357 | 15,590,673 | 0.00 | 2.64 | 5000000000 | PLACCUP | SELECT ... |
| 3,590 | 216 | 163 | 22.02 | 1.71 | 5000000000 | PLACCUP | INSERT INTO ... |
| 3,275 | 1,682 | 167 | 19.61 | 1.56 | 5000000000 | PLACCUP | UPDATE ... |

[Back to SQL Statistics](#)
[Back to Top](#)

SQL ordered by CPU Time

- Resources reported for PL/SQL code includes the resources used by all SQL statements called by the code.
- % Total DB Time is the Elapsed Time of the SQL statement divided into the Total Database Time multiplied by 100

| CPU Time (s) | Elapsed Time (s) | Executions | CPU per Exec (s) | % Total DB Time | SQL Id | SQL Module | SQL Text |
|--------------|------------------|------------|------------------|-----------------|------------|------------|-------------------------------------|
| 58,111 | 180,654 | 12 | 4842.56 | 86.10 | 5000000000 | PLACCUP | BEGIN ... |
| 35,905 | 88,004 | 174 | 206.35 | 41.94 | 5000000000 | PLACCUP | BEGIN :1 := ... |
| 5,048 | 12,080 | 3,519 | 1.43 | 5.76 | 5000000000 | PLACCUP | INSERT INTO ... |
| 4,475 | 8,754 | 167 | 26.79 | 4.17 | 5000000000 | PLACCUP | UPDATE ... |
| 3,294 | 25,374 | 167 | 19.73 | 12.09 | 5000000000 | PLACCUP | INSERT ALL WHEN ... |
| 2,939 | 16,124 | 174 | 16.89 | 7.68 | 5000000000 | PLACCUP | INSERT INTO ... |
| 2,357 | 5,545 | 15,590,673 | 0.00 | 2.64 | 5000000000 | PLACCUP | SELECT ... |
| 1,682 | 3,275 | 167 | 10.07 | 1.56 | 5000000000 | PLACCUP | UPDATE ... |
| 1,484 | 6,177 | 167 | 8.89 | 2.94 | 5000000000 | PLACCUP | INSERT INTO ... |
| 1,293 | 8,313 | 167 | 7.75 | 3.96 | 5000000000 | PLACCUP | INSERT INTO ... |
| 216 | 3,590 | 163 | 1.33 | 1.71 | 5000000000 | PLACCUP | INSERT INTO ... |



Linux on IBM System z

Production Readiness

Production Readiness

- Did the PoC validate the initial sizing
 - If not, attempt to resize or use PoC information as the basis
- Did the PoC test the availability requirements established during the requirements phase (i.e., Oracle MAA)
 - Standalone DB
 - Active/Passive
 - RAC with Active/Active
 - Use of multiple physical z10 machines
 - Data Guard for DR
- Is there sufficient IFL capacity, memory, and I/O for production
 - Are you ready to measure capacity usage over the long term.
- Are the latest Oracle patches applied

Summary

- Sizing in advance is important; really, it is.
- Think virtually with regard to allocation of resources
- PoCs with smaller and less important Oracle databases might be a good start
- Oracle database on Loz can handle ERP sized databases
- Oracle database on Loz can handle data warehouse applications
- IBM and Oracle continue to work together and invest in improving the Oracle on Loz solution