

Oracle Databases on Linux for z Systems - PoC and beyond

Sam Amsavelu

samvelu@us.ibm.com

ISV & Channels Technical Sales - Oracle

IBM Advanced Technical Skills (ATS), America

Session Objective

- The objective of this session is to share our team experiences in conducting Oracle databases Proof of Concepts (POC) on System z Linux for more than hundred customer engagements

Step by step

- Define the POC scope
- Architect the environment
- Resources estimation
- Build the environment
- Validate the environment
- Conduct the POC

Define the POC scope

- Objectives
- Document the success criteria
- Methodologies to validate
- Milestones
- Teamwork

Define the POC scope

- **What are the main objectives for this POC?**
- Every player has their vision of the scope for the POC.
 - CTO, CFO, SA, DBA, Users....
- Validation of technology for
 - Server consolidation, Virtualization
 - Cloud
- High Availability, DR
 - Frequent outages in the current system
 - Adapting DR in the current Legacy system
- 99% ends with performance....
 - Except

Define the POC scope

- **Define and document the POC objectives**
 - Achievable targets
 - Everyone is busy
- **Sample objectives are:**
 - Payroll application should be able to run on z Systems
 - How a Oracle RAC environment with Data Guard can be implemented on z Systems
 - Implementing Oracle Databases as a Service on z Systems

Define the POC scope

- If a Payroll application objective is chosen then
- Define and document the **POC scopes** by asking the following questions
- Only Database layer or application layer (WebSphere or Weblogic) also have to be tested with Database layer?
 - What is the Oracle version?
 - Stand alone database or Grid / ASM / Oracle RAC?
 - Database size?
 - Memory requirements?
- For Application layer
 - Support on System z Linux -- Customer has to press the vendor
 - Other dependencies (Wintel servers)

Define the POC scope

- **Define and document the POC success criteria**
- Following are some examples:
 - RMAN backup jobs should be completed within three hours
 - Average screen returns should meet the SLA during peak 300 users connections to the system
 - Application Functionality
 - Maintenance and day to day activities
-

Define the POC scope

- **Define the methodologies to prove the success criteria**
- Real data and real environment provides real results
- Home grown test scripts
 - Stress test, q/a scripts
- IBM or other vendors automated testing tools
- Oracle RAT tool
 - Good choice to compare the Database time
 - When complicated application environment
- Any Synthetic tools
 - Swingbench
 - Be cautious
- POC is not a benchmark

Define the POC scope

- **Estimate and establish the milestones**
 - Network availability
 - Firewall issues
 - Space availability
 - Backup, dump, results
 - Build the System z Linux environment
 - Build Oracle Databases and migrate the data
 - Application environment setup
 - Testing environment setup
 - Preliminary testing
 - Rework
 - Final testing
 - Evaluation

Define the POC scope

- **Form a team**

- Point persons in each of the following
 - Network
 - Storage
 - DBA
 - System Administration Linux
 - System Administration zVM
 - Application
 - Testing

- Method of communication and frequency

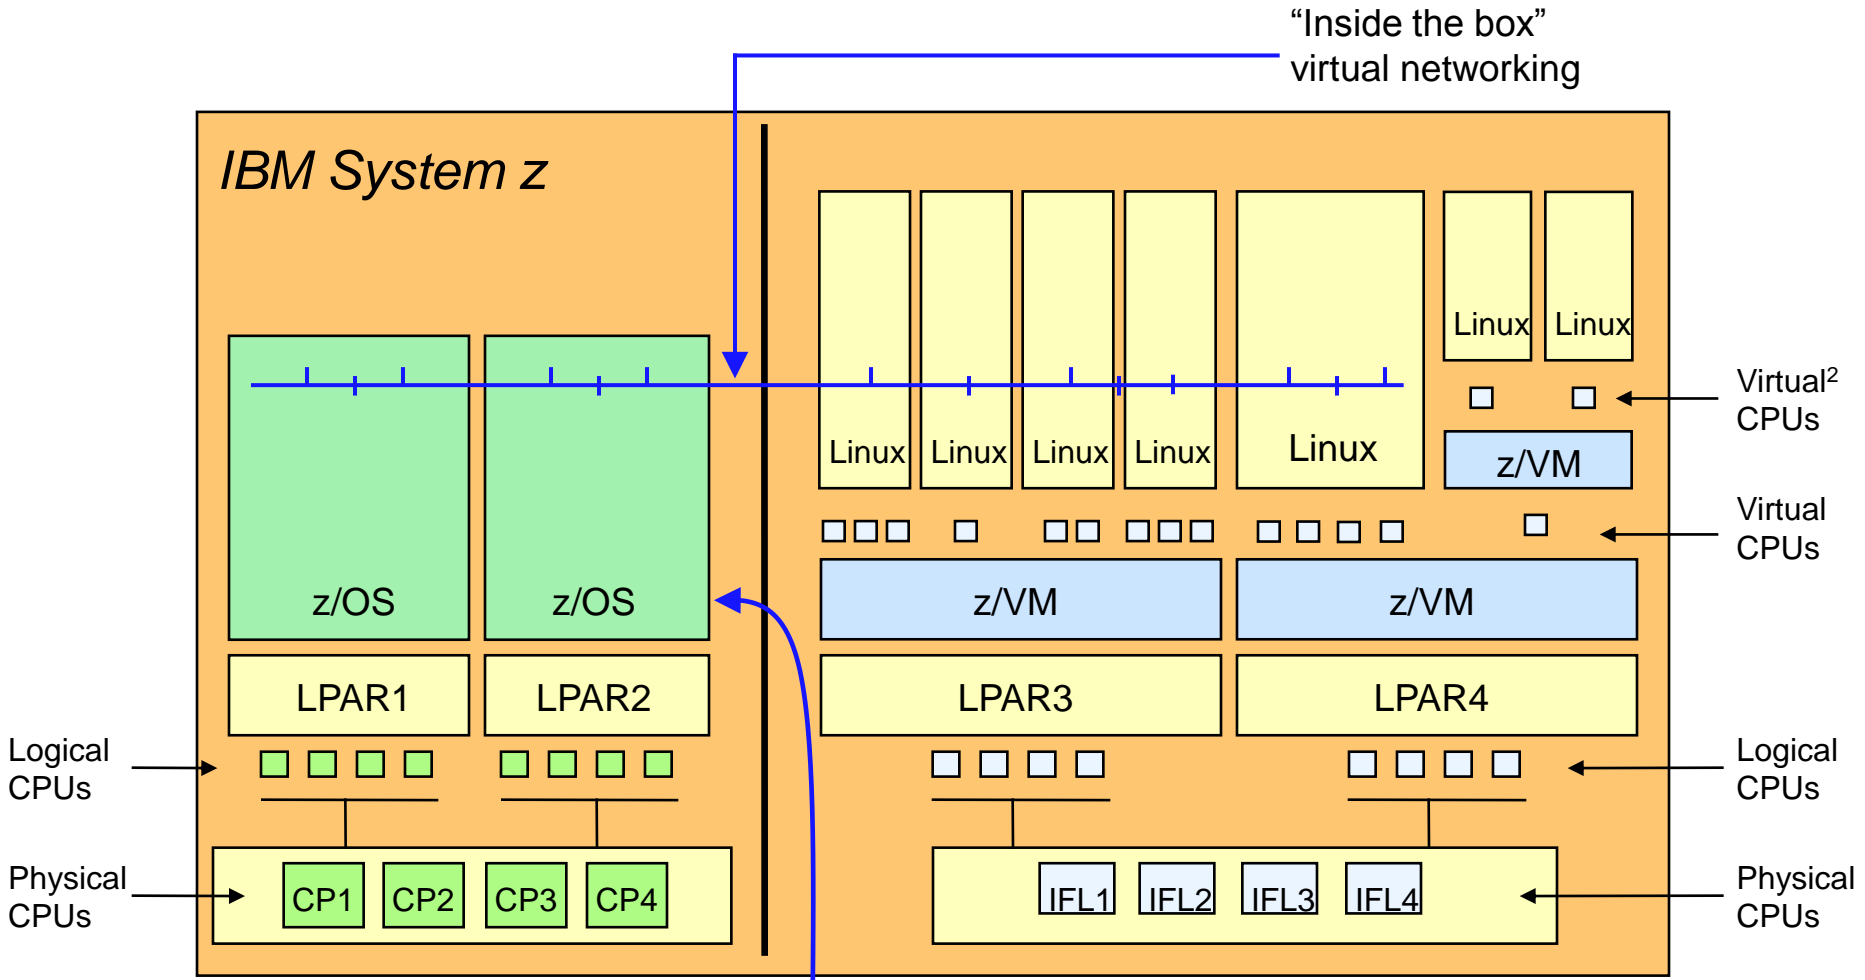
- Build a project plan

- Project plan may have many activities but not necessary to discuss each and every line items
- Just track maximum of five major milestones

Architect the POC Environment

- Take out the mystery
- High Availability discussions
- Storage
- Network
- Concept architecture
-

Take out the mystery



IFL processors have no impact on z/OS license fees

Terminology

- 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 (i.e. virtualization)
- 'Linux on z' is a better term and is less confusing in the Oracle space
- SLES and Red Hat have Oracle certified distributions for Loz
 - SLES 10 and 11
 - Red Hat 6.x

Oracle Databases on System z Linux

- Oracle Databases on System z Linux works same way they work on any other platform.
- It installs and is administered to the same way as it is on Linux on Intel.
- The differences are in configuring Linux and z/VM.
- From a DBA perspective, once they get an IP address, userid and pw for Linux they are good to go.

Define the High Availability

- **Do we need RAC or not on System z**
- Depends on your availability requirements
 - Planned outages
 - Maintenance (OS, Database, application)
 - UPS testing etc
 - Unplanned outages
 - Hardware failures
 - Software abend????? Is there a term?
 - Oracle
 - Linux
 - Scalability
 - Vertical
 - Horizontal
- Everyone will say they want 24 X 7 but is it real???

Oracle storage requirements on System z

- **Requires space for the following**
 - Oracle, Linux binaries
 - Oracle Data files
 - Oracle temp, undo table spaces, redo and archive logs
 - Dump files
- **Supported by**
 - ECKD (3390)
 - FCP/SCSI
 - Flash

Oracle / Linux binaries

- You can place binaries on ECKD / minidisks
 - Easy to clone
 - Utilize established backup and recovery mechanisms that the System z customer already has in place
- FCP (SCSI) also fine
 - Performance difference for binaries are minimal
- Once the system is in production you may explore a common binary code approach
 - Easy to maintain
 - But lose your flexibility
 - Same binary level for all the applications, difficult to patch a specific database
 - Availability issues
 - lose the code disk, lose all the databases availability

Oracle data files (DASD or FCP)

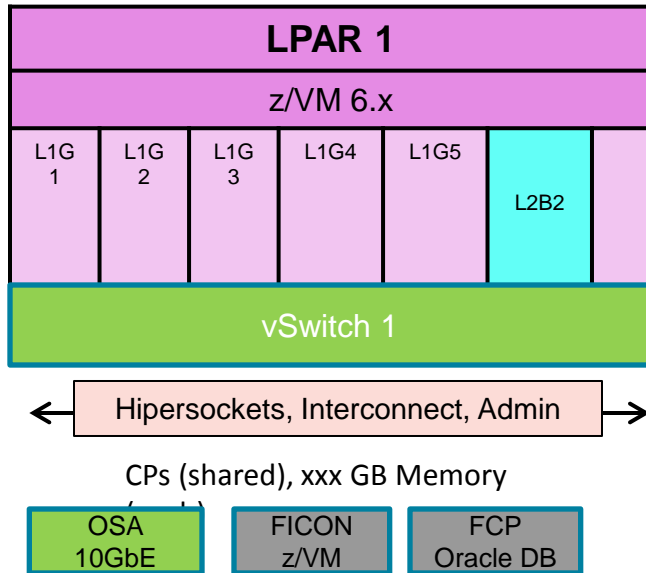
- FCP (SCSI) offers better throughput and performance
 - Use Linux multipathing for SCSI
- ECKD uses less CPU per transaction
 - We get good performance with
 - HyperPav in the DASD subsystem and HyperPav driver support in Linux distribution SLES 11 update 1 and RHEL 6
 - Multipathing is handled by z/VM




Network

- Very important
- Firewall rules
- Connectivity
- Bandwidth

POC Concept Configuration



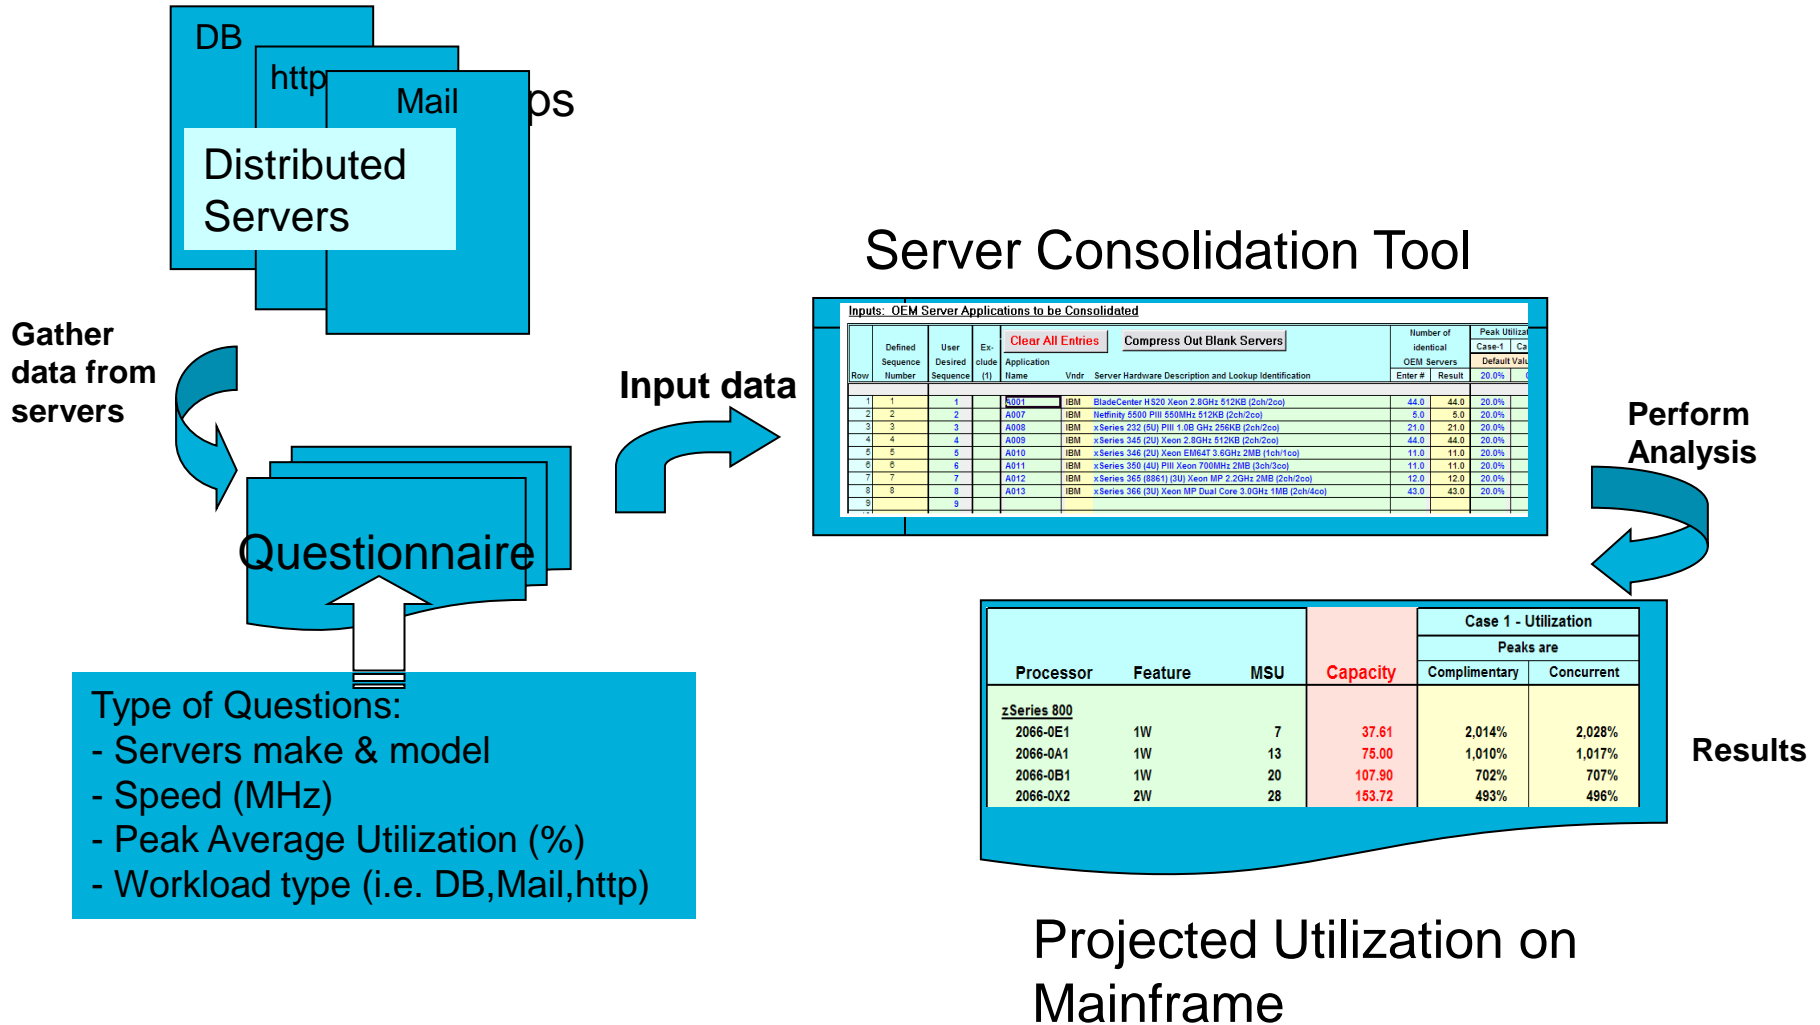
How many Linux Guests?
Application servers ?
How many for DB?



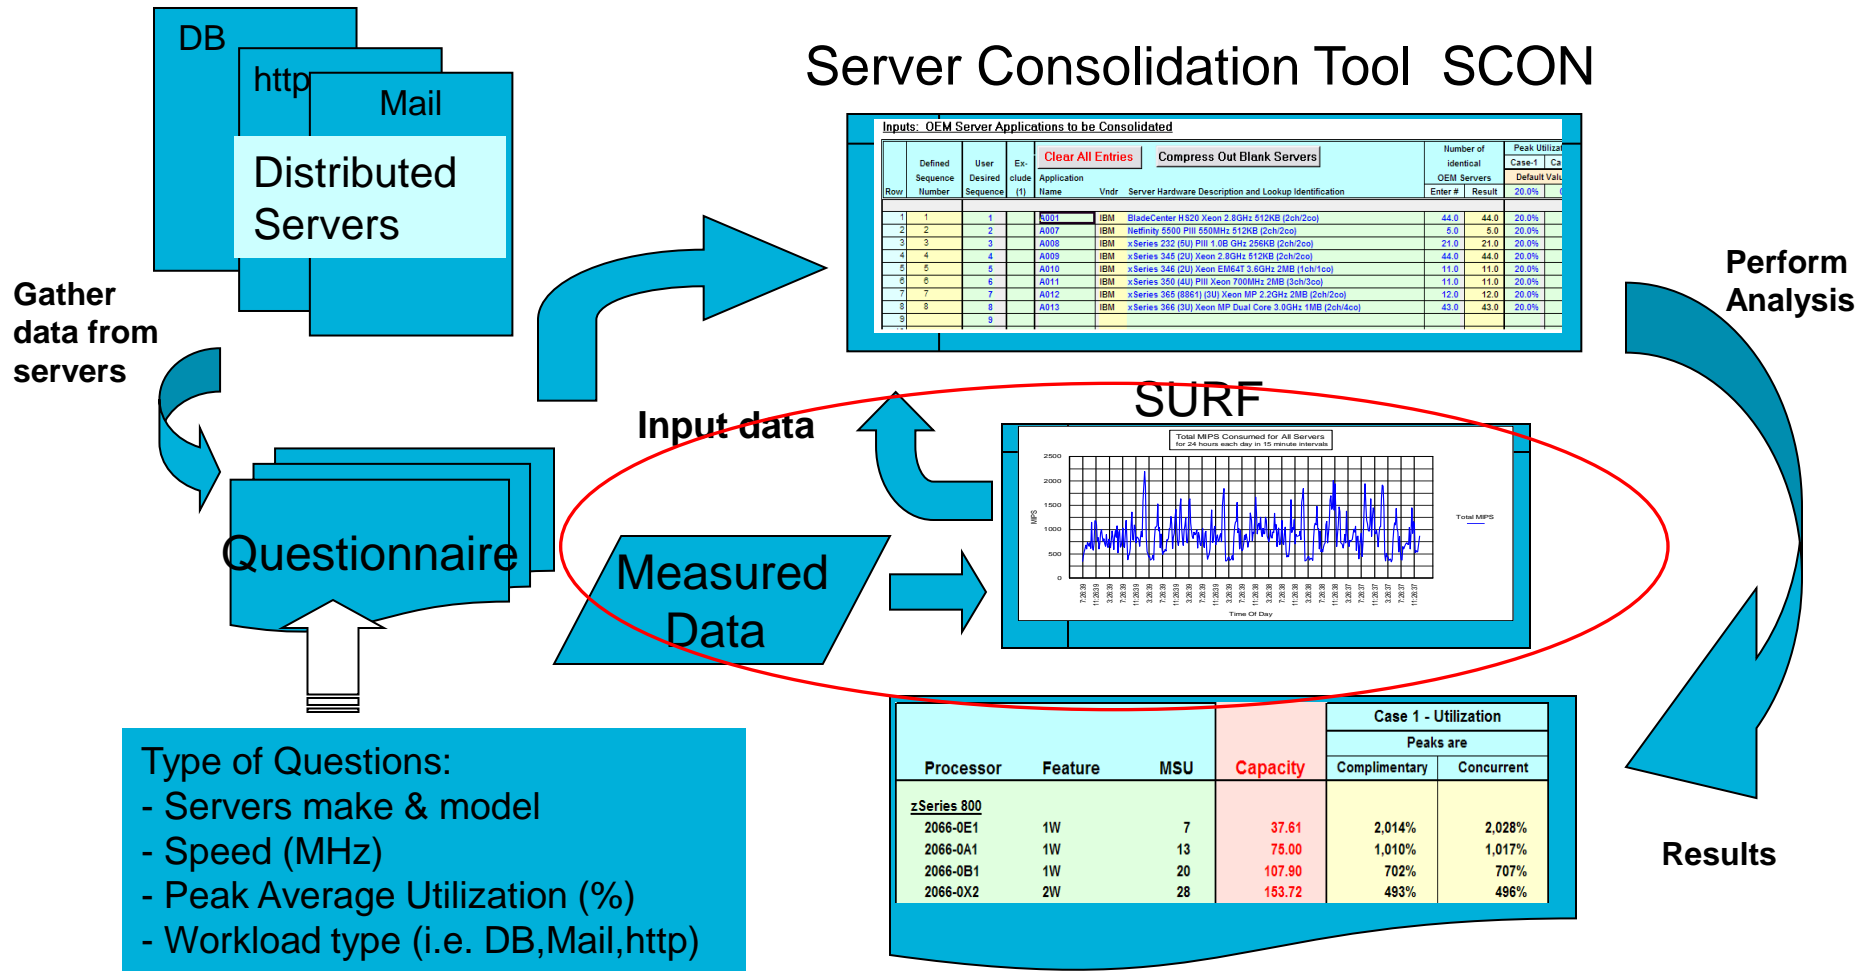
Resources Estimation

- Sizing process
- IFL
- Memory
-

System z Linux Server Consolidation Sizing Process - SCON




System z Linux Server Consolidation Sizing Process – SCON with SURF



Projected Utilization on Mainframe

Oracle Database Memory sizing for deployment on System z

- Obtain Oracle SGA and PGA sizes from all database instances
- Calculate individual guest storage requirements (assume MB):
 - Sum of (optimized) SGA and PGA settings +
 - 256 MB for ASM +
 - 512 MB for Linux +
 - 512 MB for Oracle Enterprise Manager 12c agent (if used) +
 - threads and process memory +
 - 10%*
- *Consider implementing Linux HugePages for large SGAs to avoid PageTable growth



Build the environment

- z/VM LPAR
 - SSI or Stand alone
 - Latest PMR
- Linux
 - Distribution
- Virtual CPU allocation
- Memory
- Databases

-

CPU

- Estimate the IFLs needed for the workloads (IBM Techline)
- Each LPAR is going to be assigned with logical processors
- Each Linux guest is going to be assigned with virtual processors
- Make sure that number of virtual processors defined for a Linux guest does not exceed the logical processors available to the LPAR
- It is our experience that Oracle performs better with more than one virtual processor
- Oracle recognizes automatically how many virtual processors (CPUs) are assigned to the guest
- Using z/VM we can control how much CPU resources can be allocated to each guest
 - Hard and soft limit

Oracle Database Memory sizing for deployment on System z

- Apply a z/VM memory over-commit factor such as 1.2 or 1.5 for production
 - For development you can be more aggressive
- System z memory = real memory for guests + memory for z/VM

- Swap considerations
 - The total of swap space and the guest memory relates to how much virtual memory is available for Linux to handle
 - If that is not enough for Linux then obviously will end up with out-of-memory limitations.

- The recommended best practice is to allocate two VDISKs and some disk for swap according to your comfort level and application behavior

- Monitor the VDISK and if it starts to swap into the second one, it is time to increase the memory to the guest

Linux pre-reqs

- Oracle has provided a “rpm checker” module to validate if the Linux guest has all the required rpms installed
- Also make sure that VNC Server or any graphical display
- Follow the Redbooks and Oracle support notes
- ext3 for SLES and ext4 for RHEL 6

Oracle database setup DBA is our friend...

- Many months of production experience into a days of POC
- Need DBA help to tune the system
 - Any SQL Profiles?
 - Any plan adjustments?
 - Any specific initialization parameters?
 - Automatic memory management?
 - If all the Oracle related system customization are recreated?
- Segregate data and log files devices
- Size adequately storage for temp and undo table spaces

Methodology to migrate the data

- Have enough space
- Connectivity and network bandwidth
- Oracle tools available for migration
 - DataPump (expdp / impdp)
 - Transportable tablespaces
 - Transportable databases

Validate the environment

- Tools

Use the right tools for the relevant component

- zVM tools
 - Performance tool kit
 - Velocity
- Linux tools
 - sysstat package
 - sar, top, vmstat, mpstat, iostat
- Tivoli products
- Oracle tools
 - Many of the Oracle Database tools require licenses from Oracle
 - Database tools can be executed in command line or through Oracle Enterprise Manager
 - Oracle also provides some performance / diagnostic tools available for downloading via Oracle 'MyOracleSupport' web
 - Oracle tools are transparent across other distributed platforms

Use the right tools for the relevant component

- Oracle tools available in the product
 - AWR
 - ADDM
 - ASH
 - Advisors
 - ORION
- Tools downloadable from MyOracleSupport (MOS)
 - SQLT
 - LTOM
 - RDA
 - OSWatcher

References

- SQLT (SQLTExplain)
 - SQLT helps to expedite SQL tuning process
 - Refer to MOS Note 215187.1
- LTOM (Light Onboard Monitor)
 - LTOM provides real-time automatic problem detection and data collection
 - Refer to MOS Note 352363.1
- RDA (Remote Diagnostic Agent)
 - The data captured provides Oracle Support with a comprehensive picture of the customer's environment which aids in problem diagnosis
 - Refer to MOS Note 414966.1
- OSWatcher
 - Gathers system profile data – vmstat, top, iostat etc works like ‘sar’
 - Refer to MOS Note 461053.1

Oracle data files (DASD or FCP)

- Verify I/O Performance with Oracle Orion/Calibrate tool
 - I/O simulator simulates reads and writes without creating a database
 - Now included with Oracle install code
 - Get your I/O subsystem tuned ahead of time



Information resources

- <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
 - REDP-4788-00 Installing Oracle 11gR2 RAC on Linux on System z
 - **Experiences with Oracle 11gR2 on Linux on System z**
 - <http://www.redbooks.ibm.com/redpieces/abstracts/sg248104.html>
- <http://linuxmain.blogspot.com>
- <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

My Oracle Support Links

- Note 1306465.1 Getting Started - 11gR2 Grid Infrastructure, SI(Single Instance), ASM and DB (IBM: Linux on System z)
- Note 1470834.1 - Requirements for Installing Oracle 11gR2 on RHEL 6 on IBM: Linux on System z (s390x)
- Note 1290644.1 - Requirements for Installing Oracle 11gR2 on SLES11 on IBM: Linux on System z (s390x) Also review note: OHASD fails to start on SuSE 11 SP2 on IBM: Linux on System z [ID 1476511.1]
- Note 1308859.1 Requirements for Installing Oracle 11gR2 on SLES 10 on IBM: Linux on System z (s390x)
- Note 1306889.1 Requirements for Installing Oracle 11gR2 on RHEL 5 on IBM: Linux on System z (s390x)
- Note 1086769.1 - Ensure you have prerequisite rpms to install Oracle Database and AS10g(midtier) on IBM: Linux on System z (s390x)
- Note 1377392.1 How to Manually Configure Disk Storage devices for use with Oracle ASM 11.2 on IBM: Linux on System z)
- Note 1400185.1 How to Upgrade Oracle Restart i.e. Single Node Grid Infrastructure/ASM from 11.2.0.2 to 11.2.0.3
- Note 1276058.1 Oracle GoldenGate Best Practices: Instantiation from an Oracle Source Database
- Note 1413787.1 How to completely remove 11.2 Grid Infrastructure, CRS and/or Oracle Restart - IBM: Linux on System z

- Note 259301.1 CRS and 10g Real Application Clusters
- Note 268937.1 Repairing or Restoring an Inconsistent OCR in RAC
- Note 239998.1 10g RAC How to clean up after a failed CRS Install
- Note 220970.1 RAC Frequently Asked Questions Topic

- Note 420382.1 Requirements for Installing Oracle 10gR2 RDBMS on RHEL 4 on zLinux (s390x).
- Note 431443.1 Requirements for Installing Oracle 10gR2 RDBMS on SLES 9 zLinux (s390x)
- Note 1082253 Requirements for Installing Oracle 10gR2 RDBMS on SLES 10 zLinux (s390x)
- Note 741646.1 Requirements for Installing Oracle 10gR2 RDBMS on RHEL 5 on zLinux (s390x).
- Note 415182.1 DB Install Requirements Quick Reference - zSeries based Linux .

- Note 741146.1 Installing Standalone Agent 10.2 on Linux on z

System z Social Media

- System z official Twitter handle:
 - [@ibm_system_z](#)
- Top Facebook pages related to System z:
 - [Systemz Mainframe](#)
 - [IBM System z on Campus](#)
 - [IBM Mainframe Professionals](#)
 - [Millennial Mainframer](#)
- Top LinkedIn Groups related to System z:
 - [Mainframe Experts Network](#)
 - [Mainframe](#)
 - [IBM Mainframe](#)
 - [System z Advocates](#)
 - [Cloud Mainframe Computing](#)
- YouTube
 - [IBM System z](#)



- Leading Blogs related to System z:
 - [Evangelizing Mainframe \(Destination z blog\)](#)
 - [Mainframe Performance Topics](#)
 - [Common Sense](#)
 - [Enterprise Class Innovation: System z perspectives](#)
 - [Mainframe](#)
 - [MainframeZone](#)
 - [Smarter Computing Blog](#)
 - [Millennial Mainframer](#)

References

- Refer to the following Redbook for the kernel parameters
- The following redbook will be handy : "Experiences with Oracle 11gR2 on Linux on System z":
 - <http://www.redbooks.ibm.com/abstracts/sg248104.html?Open>

Trademarks

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

AIX*	FICON*	IMS	Power7*	Redbooks*	WebSphere*
BladeCenter*	IBM*	InfiniBand	PowerHA	RMF	zEnterprise*
CICS*	IBM (logo)*	Lotus*	Power Systems	System x*	z/OS*
Cognos*	GDPS*	MQSeries*	PowerVM	System z*	z/VM*
DataPower*	Geographically Dispersed Parallel Sysplex	Parallel Sysplex*	PR/SM	System z10*	z/VSE*
DB2*	HiperSockets	POWER*	PureSystems	Tivoli*	
DS8000*	HyperSwap	POWER4*	Rational*		

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel Speed Step, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Windows Server and the Windows logo are trademarks of the Microsoft group of countries.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

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

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Oracle and Java are registered trademarks of Oracle and/or its affiliates

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

* Other product and service names might be trademarks of IBM or other 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.

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

- Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT").
- No other workload processing is authorized for execution on an SE.
- IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.
- IBM trademark information can be found at Copyright and trademark information