How to start an Oracle Consolidation Case and PoC

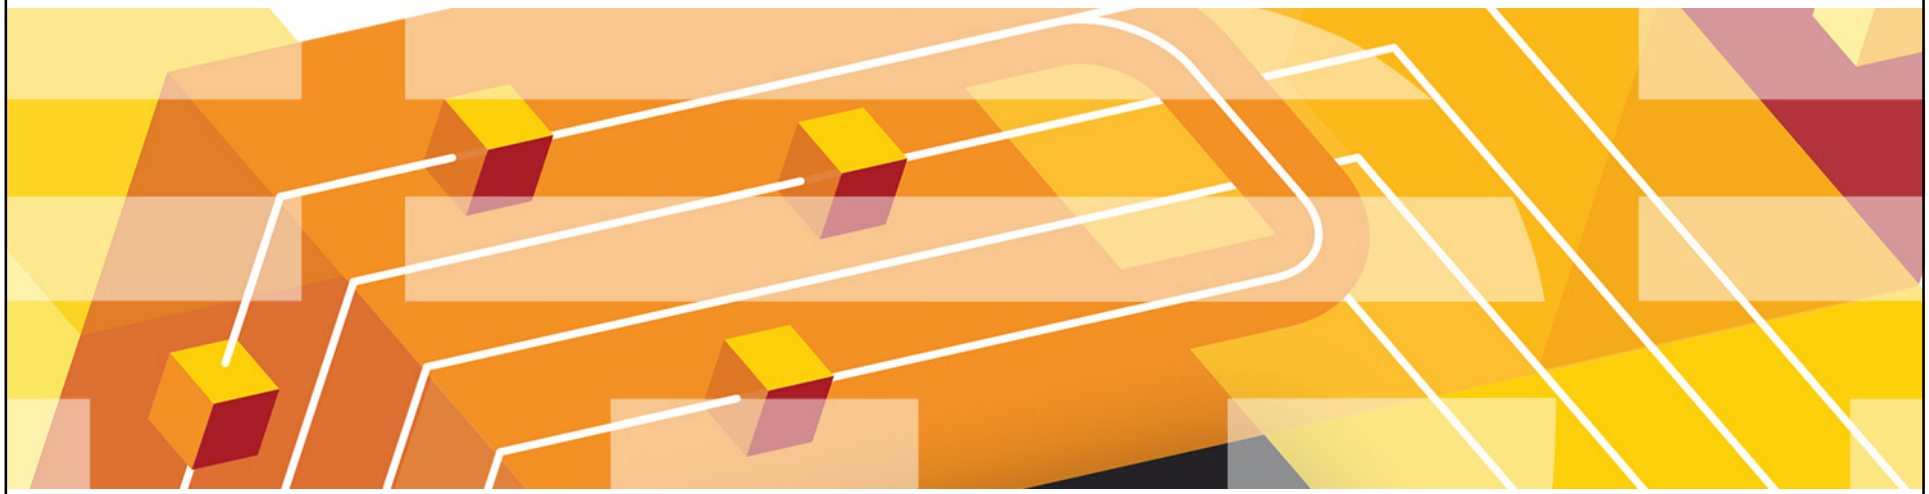

Wilhelm Mild
IBM Lab, Boeblingen, Germany
wilhelm.mild@de.ibm.com

Trademarks

- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
- This presentation contains trade-marked IBM products and technologies. Refer to the following Web site:

  http://www.ibm.com/legal/copytrade.shtml

Agenda

- Why Oracle on System z?
- Consolidation methodology
  - Scope of the project
  - Sizing (CPU and Memory)
- PoC phase
  - Preparation: what is needed
  - During the PoC: how to proceed
  - After the PoC: outcomes and next steps
- Real cases examples
- How we can help you?

Agenda

- Why Oracle on System z?
- Consolidation methodology
  - Scope of the project
  - Sizing (CPU and Memory)
- PoC phase
  - Preparation: what is needed
  - During the PoC: how to proceed
  - After the PoC: outcomes and next steps
- Real cases examples
- How we can help you?

The IBM & Oracle Relationship

Sustaining Partnership

- **Oracle 24 Yrs, PeopleSoft 21 Yrs, JD Edwards 32 Yrs, Siebel 10y**

Oracle is an IBM *"Integrated Account"* (2005)

- **Regular Executive Reviews – Global and Geographic**
- **Named Oracle Sponsor: Charles Phillips, Oracle's Co-President**

Over 19,000 Joint Customers Worldwide

- **Hardware and Software support via Apps Unlimited**

Vibrant Technology Relationship

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

Market Leading Services Practice

- **IBM's GBS is Oracle's #1 SI Partner (4900 Joint Projects!)**
- **9,500 skilled, of which 5,500 are dedicated to Oracle**

Unrivaled Customer Support Process

- **Dedicated On-Site Resources**
- **Significant Program Investments**

Value of Consolidation with Linux on System z

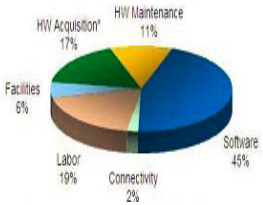

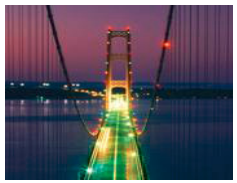

- Do more with less
  - Consolidate more servers, more networks, more applications, and more data with Linux running on z/VM
  - Achieve nearly 100% utilization of system resources nearly 100% of the time
  - Enjoy the highest levels of resource sharing, I/O bandwidth, and system availability
- Reduce costs on a bigger scale
  - Significant savings derived from reductions in server footprints, simplified infrastructure, lower software costs and a flexible and simplified infrastructure which is easy to manage
  - Consume less power and floor space
  - Save on software license fees
    - Consolidating from x86 servers to a single IFL could potentially reduce licensing costs by as much as 97%
  - Minimize hardware needed for business continuance and disaster recovery
- Manage growth and complexity, increase security
  - Exploit extensive z/VM facilities for life cycle management: provisioning, monitoring, workload mgmt, capacity planning, security, charge back, patching, backup, recovery, more...
  - Add hardware resources to an already-running system without disruption – the epitome of Smarter Infrastructure
  - Consolidation on a “scale up” machine like Linux on System z means fewer cables, fewer components to impede growth
- More flexibility, minimize lead time for new projects
  - Consolidating Oracle and Linux environments to a single System z server offers significant advantages in terms of flexibility
  - Rapid provisioning reduces lead time for new IT projects, helping to increase business agility

Why High-End Servers?

Utilization on x86 systems

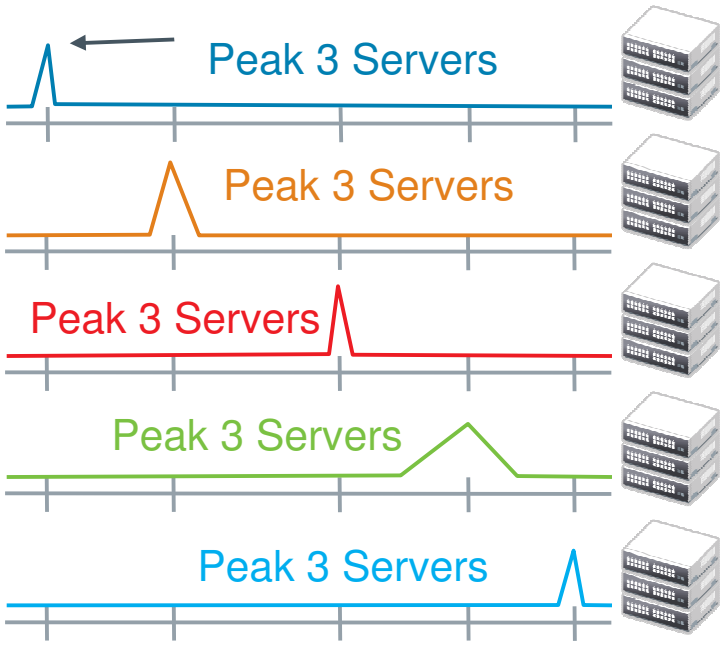

According to a study by Gartner, data centers that do not use virtualization have an average server CPU utilization rate of only 15%.

Mixed Utilization on IBM High End Servers

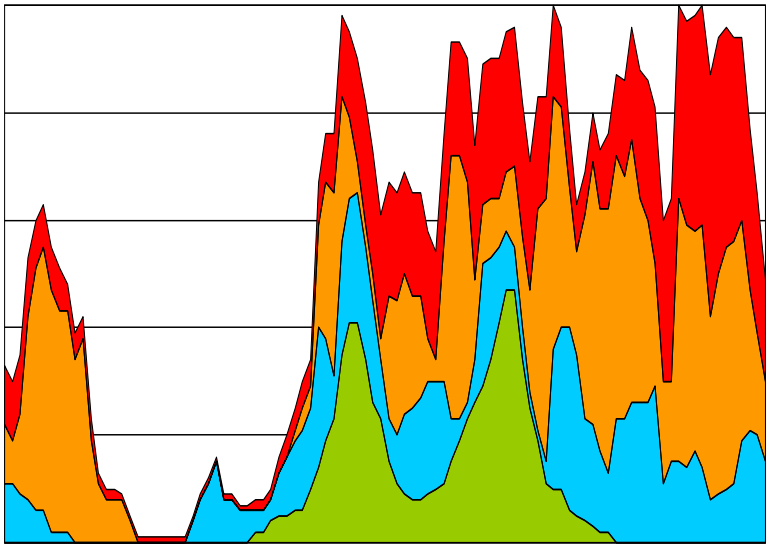

IBM High End Server: Up to 100% utilization

- Highly virtualized and shared resources
- Fewer servers, less power, cooling & admin
- Optimized use of SW assets

Reducing software cost through consolidation
Example: Oracle database

- License and annual Software Update License & Support is based on processor cores
- A “processor core factor” is applied to adjust for different technologies

ORACLE

Oracle Technology Global Price List
January 7, 2011
Software Investment Guide

	Processor Licence	Software Update Licence & Support
Oracle Database		
Standard Edition One	5,800	1,276.00
Standard Edition	17,500	3,850.00
Enterprise Edition	47,500	10,450.00
Personal Edition	-	-
Lite Mobile Server	23,000	5,060.00
Enterprise Edition Options:		
Real Application Clusters	23,000	5,060.00
Real Application Clusters One Node	10,000	2,200.00
Active Data Guard	10,000	2,200.00

Prices in USA (Dollar)

Oracle documentation: http://www.oracle.com/us/corporate/pricing/technology-price-list-070617.pdf

ORACLE

Oracle Processor Core Factor Table
Effective Date: March 16, 2009

AMD Opteron Models 13XX, 23XX, 24XX, 41XX, 61XX, 83XX, 84XX or earlier Multicore chips	0.5
Intel Xeon Series 56XX, Series 65XX, Series 75XX, or earlier Multicore chips	0.5
IBM POWER6	1.0
IBM POWER7	1.0
IBM POWER7+	1.0
IBM System z (z10 and earlier)	1.0
All Other Multicore chips	1.0

Oracle documentation: http://www.oracle.com/us/corporate/contracts/processor-core-factor-table-070634.pdf

IBM documentation: http://www-01.ibm.com/software/lotus/passportadvantage/pvu_licensing_for_customers.html

When is an Oracle Consolidation paying out

- starting with 2 Server (RAC) installation

- Real customer situation
- For an Installation of Oracle (RAC) starting with 2 servers
  - Severs with 6 Cores - 2 X 6 = 12 Cores
  - Oracle Enterprise Licenses
  - RAC Feature
- Replacement with z114 - much cheaper and effective
  - workload could be handled with 2 IFLs
- Price saving over 3 years:
  - almost one million Euro savings

Agenda

- Why Oracle on System z?
- Consolidation methodology
  - Scope of the project
  - Sizing (CPU and Memory)
- PoC phase
  - Preparation: what is needed
  - During the PoC: how to proceed
  - After the PoC: outcomes and next steps
- Real cases examples
- How we can help you?

Method proposed for consolidation projects

- these steps are independent of the platform - not System z specific

- **1 - Scope of the project delimitation – feasibility study**
  - Gather information on existing environment (servers, applications, network)
  - Fit for Purpose study
  - Prioritize the non-functional requirements (RAS, scalability, performance, management)
  - Assess the skills
  - Can follow a Cost & Value study
- **2 - Architecture design and sizing exercise**
  - Select applications and servers to be consolidated (check support !)
  - Definition of the targeted architecture – Physical model development
  - Collection of performance and monitoring data from current distributed environment
  - Initial sizing exercise in collaboration with IBM Techline
  - First planning of the project
- **3 - Proof of Concept**
  - Functional Validation
  - Performance and Sizing Validation (if benchmark)
  - Targeted architecture validation
  - zLight can be a good option for a PoC
- **4 - Pre-production tests**
  - Validation in the real environment
  - Environment health check before production (LPAR, z/VM, Linux, Middleware)
  - Skill transfer phase
- **5 - Put to production**
  - Iterative put to production
  - Monitor the system to tune it accordingly

Agenda

- Why Oracle on System z?
- Consolidation methodology
  - Scope of the project
  - Sizing (CPU and Memory)
- PoC phase
  - Preparation: what is needed
  - During the PoC: how to proceed
  - After the PoC: outcomes and next steps
- Real cases examples
- How we can help you?

How to gather information about Oracle environment

- For System information
  - List of the servers models with details (constructor, model, CPU, cores, processor…)
- For workload information
  - Need to have an idea of the type of workload, if we have no information we take DB production
- For CPU information
  - Best is to have vmstats, collected 1 or several days during a relevant period. Collect interval should be at most 10min or less, if possible, with either
    - VMSTAT
    - SAR data
    - NMON
  - If not possible to get the vmstats we need an estimation of CPU utilization during the peak period
- For memory information (see details on p. 17)
  - To determine the quantity of about SGA and PGA sizes and memory used :
    - AWR reports
  - About number of concurrent user connections:
    - at the Linux level or AWR reports

CPU and Memory work different on System z than distributed systems - more effective and less invasive

Oracle Consolidation on System z study : methodology example

Row	Defined Sequence Number	User Desired Sequence	Ex-clude (1)	Application Name	Vendor	Server Hardware Description and Lookup Identification	Note: Fractional values must be less than 1.00 # OEM Servers Enter #	Result	Peak Utilization Case-1 Default Values 90.0%	Case-2 65.0%	Workload Assignment No.	Description
1	1	1		d-intellinx	IBM	BladeCenter HS22V Xeon E5649 Hex Core 2.53GHz (2ch/12co)	0.08	0.08	20.0%	65.0%	33	DB: Production
2	2	2		dw-bo-t	IBM	BladeCenter HS22V Xeon E5649 Hex Core 2.53GHz (2ch/12co)	0.17	0.17	20.0%	65.0%	33	DB: Production
3	3	3		IFNPROD	IBM	BladeCenter HS21 Xeon 5150 Dual Core 2.66GHz (1ch/2co)	1.00	1.00	50.0%	65.0%	33	DB: Production
4	4	4		IFNT	IBM	BladeCenter HS22V Xeon E5649 Hex Core 2.53GHz (2ch/12co)	0.17	0.17	10.0%	65.0%	33	DB: Production
5	5	5		INTELLINX	IBM	System x3850 (3U) Xeon EM64T 3.66GHz 1MB (4ch/4co)	1.00	1.00	50.0%	65.0%	33	DB: Production
6	6	6		twindba1	IBM	BladeCenter HS22V Xeon L5638 Hex Core 2.0GHz (2ch/12co)	1.00	1.00	10.0%	65.0%	33	DB: Production
7	7	7		managegrid	IBM	BladeCenter HS22V Xeon E5645 Hex Core 2.4GHz (2ch/12co)	0.33	0.33	20.0%	65.0%	33	DB: Production
8	8	8		ORANT1	IBM	BladeCenter HS21 XM Xeon E5345 Quad Core 2.33GHz (2ch/8co)	1.00	1.00	10.0%	65.0%	33	DB: Production
9	9	9		ORANT5	IBM	BladeCenter HS22V Xeon E5649 Hex Core 2.53GHz (2ch/12co)	0.33	0.33	20.0%	65.0%	33	DB: Production
10	10	10		ORAST1	IBM	BladeCenter HS22V Xeon E5649 Hex Core 2.53GHz (2ch/12co)	0.08	0.08	50.0%	65.0%	33	DB: Production
11	11	11		ORAST3	IBM	BladeCenter HS22V Xeon E5645 Hex Core 2.4GHz (2ch/12co)	0.17	0.17	10.0%	65.0%	33	DB: Production
12	12	12		OraTest1	IBM	BladeCenter HS22V Xeon L5640 Hex Core 2.26GHz (2ch/12co)	0.17	0.17	20.0%	65.0%	33	DB: Production
13	13	13		OraTest3	IBM	BladeCenter HS22V Xeon X5675 Hex Core 3.06GHz (2ch/12co)	0.08	0.08	30.0%	65.0%	33	DB: Production
14	14	14		OraTest5	IBM	BladeCenter HS22V Xeon L5640 Hex Core 2.26GHz (2ch/12co)	0.17	0.17	20.0%	65.0%	33	DB: Production
15	15	15		PCTHCON	IBM	BladeCenter HS22V Xeon X5675 Hex Core 3.06GHz (2ch/12co)	0.08	0.08	10.0%	65.0%	33	DB: Production
16	16	16		pdw-boxi	IBM	BladeCenter HS21 XM Xeon E5420 Quad Core 2.5GHz (2ch/8co)	1.00	1.00	10.0%	65.0%	33	DB: Production
17	17	17		PGL	IBM	BladeCenter HS22V Xeon X5650 Hex Core 2.66GHz (2ch/12co)	0.17	0.17	10.0%	65.0%	33	DB: Production
18		18	1	pisrdb1					40.0%		33	
19		19	1	pisrdb2					50.0%		33	
20	18	20		poralnx01	IBM	BladeCenter HS22V Xeon X5650 Hex Core 2.66GHz (2ch/12co)	1.00	1.00	30.0%	65.0%	33	DB: Production
21	19	21		poralnx02	IBM	BladeCenter HS22V Xeon X5650 Hex Core 2.66GHz (2ch/12co)	1.00	1.00	30.0%	65.0%	33	DB: Production
22	20	22		poralnx03	IBM	BladeCenter HS22V Xeon X5650 Hex Core 2.66GHz (2ch/12co)	1.00	1.00	30.0%	65.0%	33	DB: Production
23	21	23		PRIORITY	IBM	BladeCenter HS22V Xeon X5650 Hex Core 2.66GHz (2ch/12co)	0.17	0.17	10.0%	65.0%	33	DB: Production
24	22	24		ptm-oradb1.ext	IBM	BladeCenter HS22V Xeon E5649 Hex Core 2.53GHz (2ch/12co)	0.17	0.17	40.0%	65.0%	33	DB: Production
25	23	25		qaora1	IBM	BladeCenter HS22V Xeon E5649 Hex Core 2.53GHz (2ch/12co)	0.33	0.33	40.0%	65.0%	33	DB: Production
26	24	26		qaoralnx1	IBM	BladeCenter HS22V Xeon X5650 Hex Core 2.66GHz (2ch/12co)	0.17	0.17	30.0%	65.0%	33	DB: Production
27	25	27		qasrv1	IBM	BladeCenter HS22V Xeon X5670 Hex Core 2.93GHz (2ch/12co)	0.08	0.08	30.0%	65.0%	33	DB: Production
28		28	1	risrdb1					10.0%		33	
29		29	1	risrdb2					10.0%		33	
30	26	30		saoralnx1	IBM	BladeCenter HS22V Xeon X5650						
31	27	31		storalnx1	IBM	BladeCenter HS22V Xeon X5650						
32	28	32		storalnx2	IBM	BladeCenter HS22V Xeon X5650						
33	29	33		storalnx3	IBM	BladeCenter HS22V Xeon X5650						
34	30	34		TGL	IBM	BladeCenter HS22V Xeon E5649						
35		35	1	tisrsdb2								
36	31	36		toralnx1	IBM	BladeCenter HS22V Xeon X5650						
37	32	37		pmove2prod	IBM	BladeCenter HS22V Xeon E5645						
38		38	1	TSYSSDB2								
39	33	39		ttm-oradb1.ext	IBM	BladeCenter HS22V Xeon E5649						
40	34	40		pemquestpaora	IBM	BladeCenter HS22V Xeon E5645						
41		41	1	Tmobidb								
42	35	42		tsysdba1	IBM	BladeCenter HS21 XM Xeon E53						
43		43	1	Pmobidb								
44		44	1	STmobidb								
45		45	1	QAmobidb								
46	36	46		TPRIORITY	IBM	BladeCenter HS22V Xeon X5650						

Summary of Servers to be Consolidated

Servers	Chips	Cores	Applications
36	35	172	36

Processor	Feature	Utilization for Case 1 < Complementary 0%	40.0%	Peaks 70.0%	Concurrent > 100%
IBM z196 IFL					
2817-7xx I9	9W IFL	60%	105%	139%	174%
2817-7xx I10	10W IFL	55%	96%	126%	157%
2817-7xx I11	11W IFL	50%	88%	116%	144%
2817-7xx I12	12W IFL	46%	81%	107%	133%
2817-7xx I13	13W IFL	43%	75%	99%	124%
2817-7xx I14	14W IFL	40%	70%	93%	116%
2817-7xx I15	15W IFL	38%	66%	87%	109%
2817-7xx I16	16W IFL	36%	62%	82%	102%
2817-7xx I17	17W IFL	34%	59%	78%	97%

Example of memory sizing for Oracle

- Standard Memory estimation = **sum of**:
  - Memory required for **Linux Kernel**: 512 MB
  - Memory required for **Oracle SGA**: As per DBA calculation (AWR Report)
  - Memory required for **Oracle PGA**: As per DBA calculation (AWR Report)
  - Memory required for **Oracle ASM**: 256 MB to 512 MB (If ASM is used)
  - Memory required for **additional agents** like OEM, Tivoli etc., as needed by the application
  - **Linux Overhead** requirements: 5 % of the total memory

- Oracle documents suggest the following about the memory requirements for dedicated user connections :
  - They depend on what is being done in the client connections and how many client connections are being used
  - On average, dedicated connections use 4.5MB per connection (assuming you are using default values for workarea parameters
  - See PGA advisor to check if the PGA size is optimum (avoid multipass)

Example of memory sizing for Oracle (user connections)

- For dedicated connections you can have an idea of the quantity in AWR, section Instance Activity Stats – Absolute Values

Instance Activity Stats - Absolute Values

- Statistics with absolute values (should not be diffed)

Statistic	Begin Value	End Value
session uga memory	12,618,216	13,589,576
opened cursors current	34	30
logons current	28	28
session uga memory max	502,328,936	1,945,112,232
session pga memory	133,131,664	130,427,632
session pga memory max	156,921,232	154,217,200
session cursor cache count	1,378	1,744

Agenda

- Why Oracle on System z?
- Consolidation methodology
  - Scope of the project
  - Sizing (CPU and Memory)
- PoC phase
  - Preparation: what is needed
  - During the PoC: how to proceed
  - After the PoC: outcomes and next steps
- Real cases examples
- How we can help you?

Proof of Concept preparation (1/3)

- Design the final architecture and review it with System z experts
- Determine the scope of the PoC - Verifiy all the involved components are supported!
- Determine the success criteria
  - Take performance data on the source platform if you need to do comparison tests
- Determine Hardware configuration
  - Server
    - Model
    - Partitionning
    - IFL (number, shared, dedicated…)
  - Network
  - Storage
    - Server
    - Type of disks (ECKD, SCSI)
- Determine software configuration for z/VM (if used), Linux, Oracle
  - Licenses
  - Versions
  - Patchsets levels
  - For Oracle, Critical Patch Update Advisories are available at the following location:
    Oracle Technology Network:
    http://www.oracle.com/technetwork/topics/security/alerts-086861.html

Best practices: use the latest release and level of patchset to avoid any known bug!

Proof of Concept preparation (2/3)

- **Make sure all the skills needed are available!**
- Set up the hardware
- Install z/VM and perftoolkit (if part of the PoC)
- Install Linux
- Test your I/O subsystem with Orion tool (Befor Oracle installation, because the writing test will erase the data on the disks)
- Install Oracle
  - Use RPM checker prior to installation: download the appropriate RPM checker from the bottom of the My Oracle Support (MOS) Note 1306465.1
  - Oracle DB installation is identical on System z and on distributed platforms
  - Oracle Entreprise Manager is identical

Best practices: Be careful with prerequisites for Oracle Installation!

Proof of Concept preparation (3/3)

- **Determine the success criteria before the test start – revalidate them with all the stakeholders**
- Apply best practices, among them don't forget:
  - If using ext3 then verify Oracle init.ora has the following settings:
    - filesystemio_options = setall (direct I/O)
    - disk_asynch_io=true

    to eliminate Linux double caching which wastes storage and CPU resources
  - Calibrate I/O with Oracle Enterprise Manager
  - Collect statistics at Oracle level
    - EXEC DBMS_STATS.gather_schema_stats('soe', granularity => 'ALL', cascade => true, options => 'GATHER', degree => x);

    (Where x is number of CPU * 2)
  - Increase the size of the redologs for Oracle (50 MB by default, most of time too small)
    - alter database add logfile ('/logs/swing_log1.log') size 10G;

Agenda

- Why Oracle on System z?
- Consolidation methodology
  - Scope of the project
  - Sizing (CPU and Memory)
- PoC phase
  - Preparation: what is needed
  - During the PoC: how to proceed
  - After the PoC: outcomes and next steps
- Real cases examples
- How we can help you?

During the PoC

- Make sure all the skills needed are available!
- Remind the success criteria before the test start – revalidate them with all the stakeholders
- Chose a rigorous approach to store the tests results
- Monitor your system at all levels, for example:
  - PerformanceToolKit to monitor z/VM
  - Nmon to monitor the Linux guests
  - SADC and IOstat to monitor the Linux guests in details
  - TPC to monitor the Storage Subsystem
  - Oracle Enterprise Manager DB console to monitor Oracle Database
- Keep a trace of all the results of your tests, with the changes you made (one change at a time!)

- **Document all changes made during PoC**

Agenda

- Why Oracle on System z?
- Consolidation methodology
  - Scope of the project
  - Sizing (CPU and Memory)
- PoC phase
  - Preparation: what is needed
  - During the PoC: how to proceed
  - After the PoC: outcomes and next steps
- Real cases examples
- How we can help you?

After the PoC

- Write down a document to keep a trace of the PoC
  - For this specific case (used parameters, workload optimization…)
  - For reuse for other cases!
- Present and explain the results to the customer
- Discuss the next steps
  - Additional NFR like HA or DR (often not considered during the PoC)
  - Sizing validation
  - Further functions to tests
  - Migration considerations
  - Put to production

Best practices / Return of experience

- Project management
  - Need to have an accurate statement of work
    - Description of what is expected
    - Who is doing what
  - Need to have a dedicated project manager for
    - Preparation
    - PoC
    - Results presentation and explantion
- Technical issues
  - Use best practices to set up your systems/software
  - Use the last level of patches for each component
- Skills
  - If the PoC is done at customer site we need to make sure all the skills will be available (no bottleneck during the PoC!)

Agenda

- Why Oracle on System z?
- Consolidation methodology
  - Scope of the project
  - Sizing (CPU and Memory)
- PoC phase
  - Preparation: what is needed
  - During the PoC: how to proceed
  - After the PoC: outcomes and next steps
- Real cases examples
- How we can help you?

Real case example 1: IT Service provider

PoC at Customer: without Lab involvement - challenges and long

- Context
  - This IT service provider has a lot of Oracle DB on distributed systems
  - They had some experience with Linux on System z (just for test)
  - They wanted to be able to quickly deploy new Oracle servers
  - They wanted to test their own infrastructure (« background task »)

- During the PoC
  - They asked for help for installation and documentation
  - They experienced errors during the installation

- Results and lessons learned from the experience
  - No planning, no dedicated resources =>it took a long time to take a decision (several months)
  - Customer in production now => could have been much faster
  - Trend is to go towards a « Cloud » environment

Real case example 2: Travel and Transportation

PoC at customer: with joined expertise Boeblingen and Montpellier

- Context
  - This System z customer (legacy) wanted to leverage their System z infrastructure
  - They had more than 250 Oracle databases
  - They wanted to be proven that:
    - A Linux would have no impact on their production environment
    - Oracle DB was running fine on Linux on System z

- PoC description
  - Statement of work IBM/Customer
  - Set up phase
  - Test phase
  - Results delivery phase and next steps

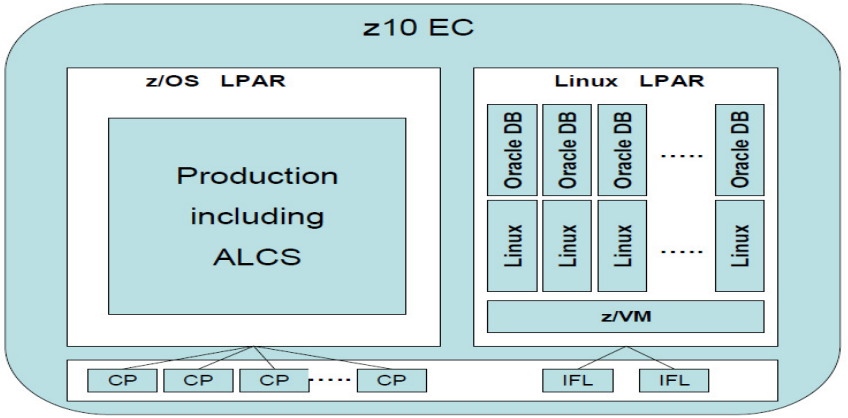

- Results
  - The Poc was done at their site with the help of BOE people (on site) and Mop people (remotely)
  - After this first step, a workshop was done at their site for HA/DR with Oracle DB on System z environment
  - Customer now in production

Real case example 3: Public Sector, Government

PoC at a Lab: Leveraging joined expertise Boeblingen and Montpellier

- Context
  - System z existing customer with strong knowledge on Linux on System z
  - They wanted to consolidate Oracle DB from Intel to System z
  - Some very critical applications needed to be at least as faster on System z as Intel to carry on consolidation
- During the PoC
  - The PoC took place in BOE with MOP support as well, and last a short period (days)
  - All the team (IBM local team, BOE, MOP and customer team) worked together
  - As soon as issues arose they were corrected immediately
- Results and return of experience
  - In most of the cases, after tuning, most of the test cases were in favor of System z
  - This PoC was key to continue the Oracel consolidation project
  - Customer is in production now

Some Reference Customers
Linux on System z with Oracle

Bank of New Zealand

Agenda

- Why Oracle on System z?
- Consolidation methodology
  - Scope of the project
  - Sizing (CPU and Memory)
- PoC phase
  - Preparation: what is needed
  - During the PoC: how to proceed
  - After the PoC: outcomes and next steps
- Real cases examples
- How we can help you?

IBM Oracle Competence Centers

IBM Oracle Competence Centers

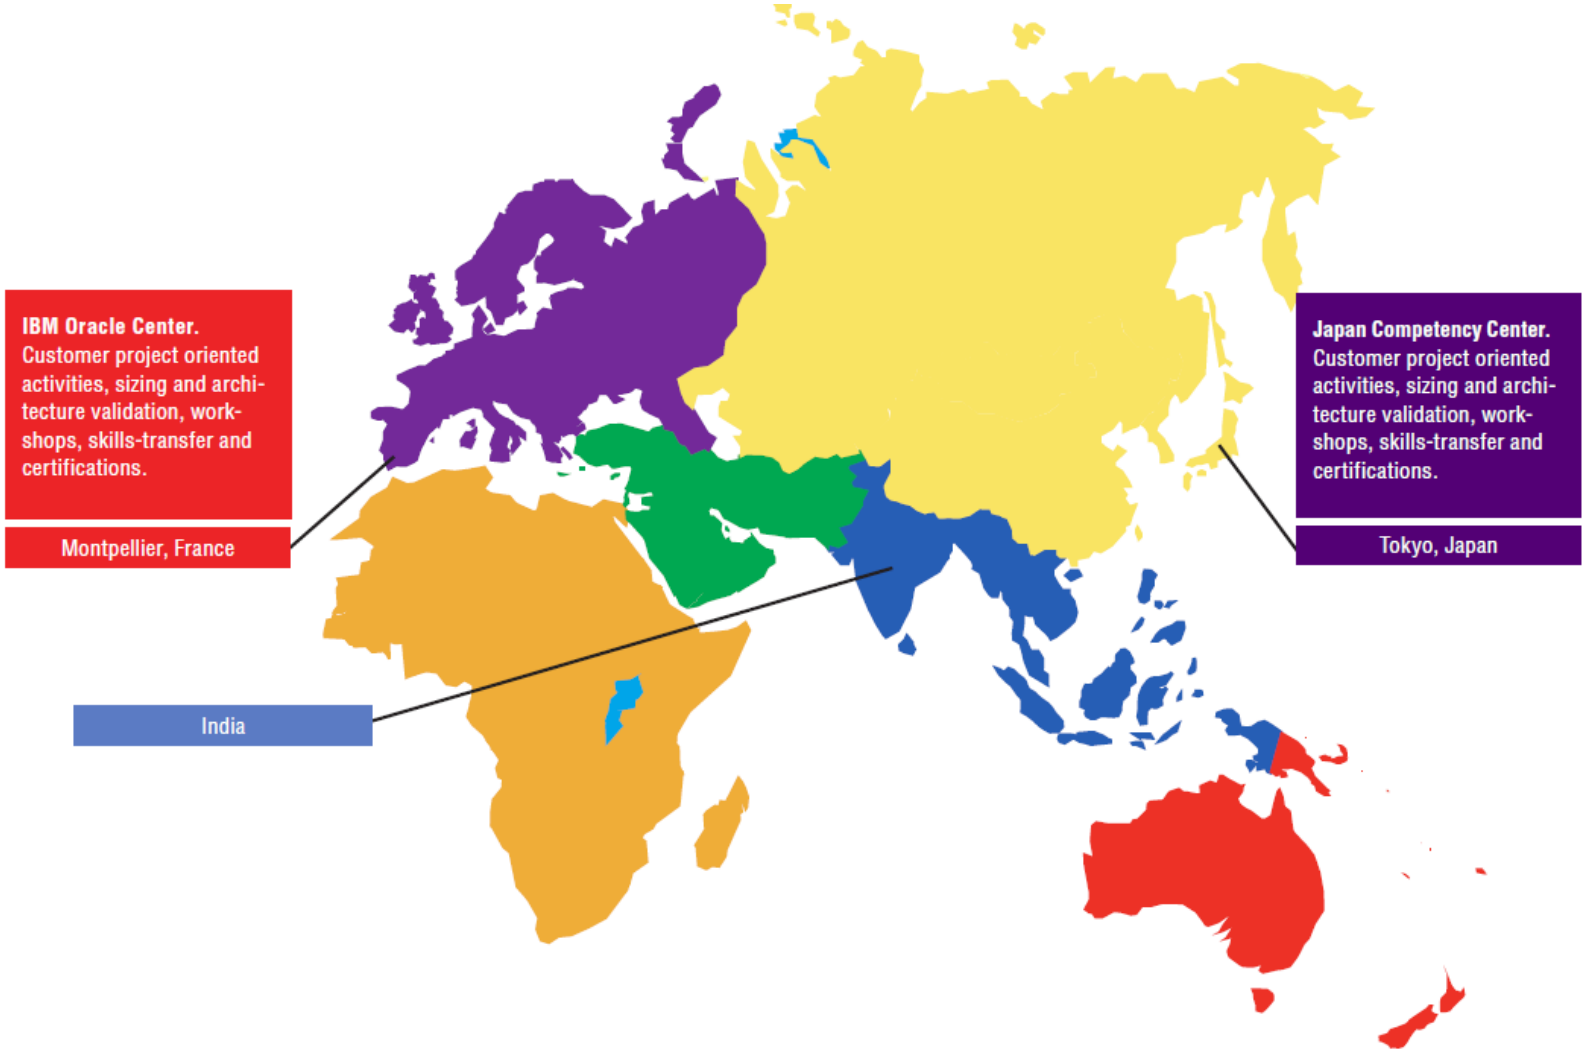

The IBM Oracle Center (IOC)

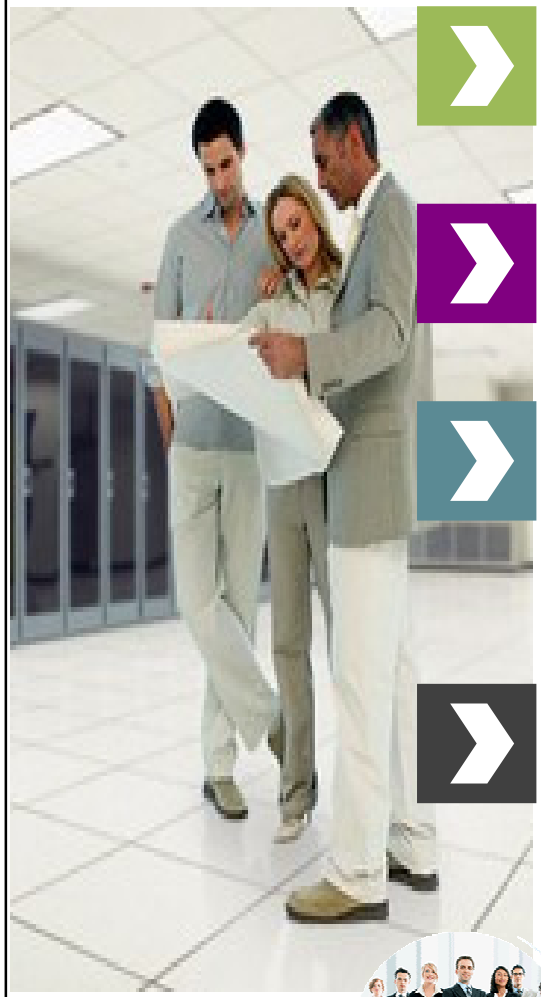

OUR MISSION

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

OUR STRENGH

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

Contact: ibmoracl@us.ibm.com

Oracle and Linux on System z – IBM & Oracle working together

- Linux on System z is Oracle's platform for the mainframe
- Oracle database 11g R2 available on Linux on System z (since 1Q2011)

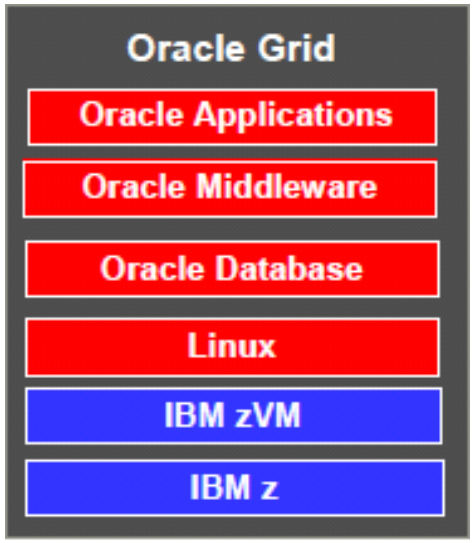

http://www.redbooks.ibm.com/abstracts/sg247191.html

Resources

- RedBooks
  - Experiences with Oracle Solutions on Linux for System z

http://www.redbooks.ibm.com/redbooks/pdfs/sg247634.pdf

- Using Oracle Solutions on Linux for System z

http://www.redbooks.ibm.com/redbooks/pdfs/sg247573.pdf

- DeveloperWorks Linux on system z
  - Tuning Hints and Tips

http://www.ibm.com/developerworks/linux/linux390/perf/index.html

- Databse Tuning for Linux on System z

http://www.ibm.com/developerworks/linux/linux390/perf/tuning_database.html

Acknowledgement

This presentation is a joint
slide deck build by
Helene Grosch from IBM Oracle Center in
Montpellier, IBM France
and Wilhelm Mild, Boeblingen Lab, IBM Germany

धन्यवाद
Hindi

多謝
Traditional Chinese

ขอบคุณ
Thai

Спасибо
Russian

Gracias
Spanish

Thank You
English

شكرا
Arabic

Merci
French

Obrigado
Brazilian Portuguese

Bedankt
Nederlands

多谢
Simplified Chinese

Danke
German

நன்றி
Tamil

ありがとうございました
Japanese

감사합니다