Konsolidierung von Datenbanken auf System z am Beispiel von Oracle

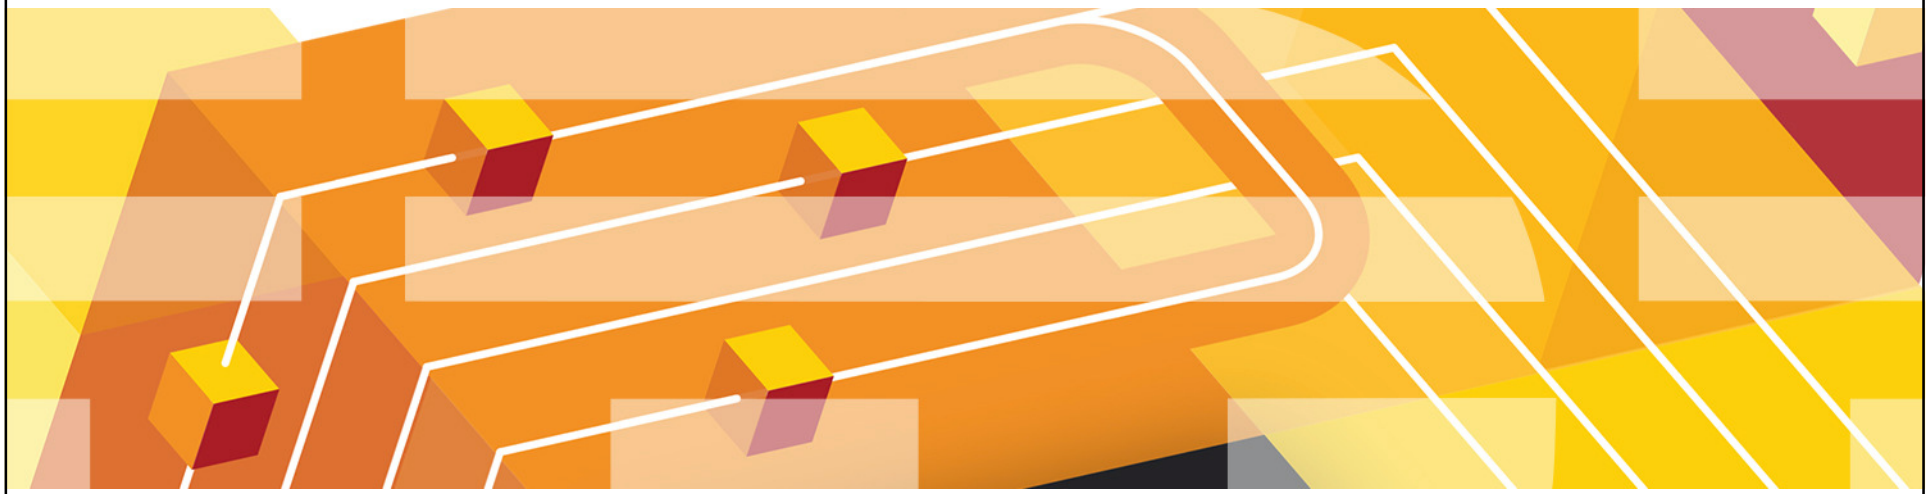

Wilhelm Mild
IBM Lab, Boeblingen, Germany
wilhelm.mild@de.ibm.com

© 2013 IBM Corporation

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

© 2013 IBM Corporation

Discussion Topics

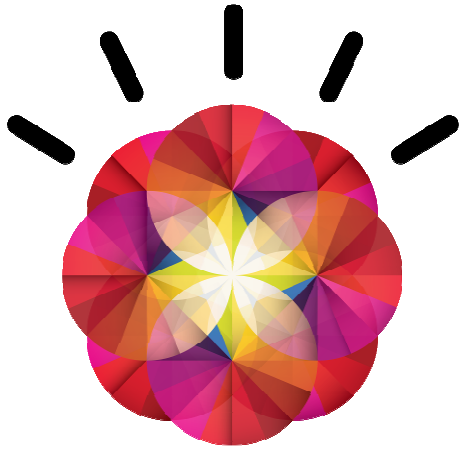

- **IBM & Oracle relationship**
- Why Oracle on Linux on zEnterprise
- Economics for consolidation
- Best practices for Oracle on Linux on System z
- Customer implementations

© 2013 IBM Corporation

The IBM & Oracle Relationship

Sustaining Partnership

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

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

- **Regular Executive Reviews – Global and Geographic**
- **Named Oracle Sponsor: Charles Phillips, Oracle's Co-President**
- **Dedicated IBM executive-led alliance team**

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 Practice**

Unrivaled Customer Support Process

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

© 2013 IBM Corporation

IBM Oracle Competence Centers

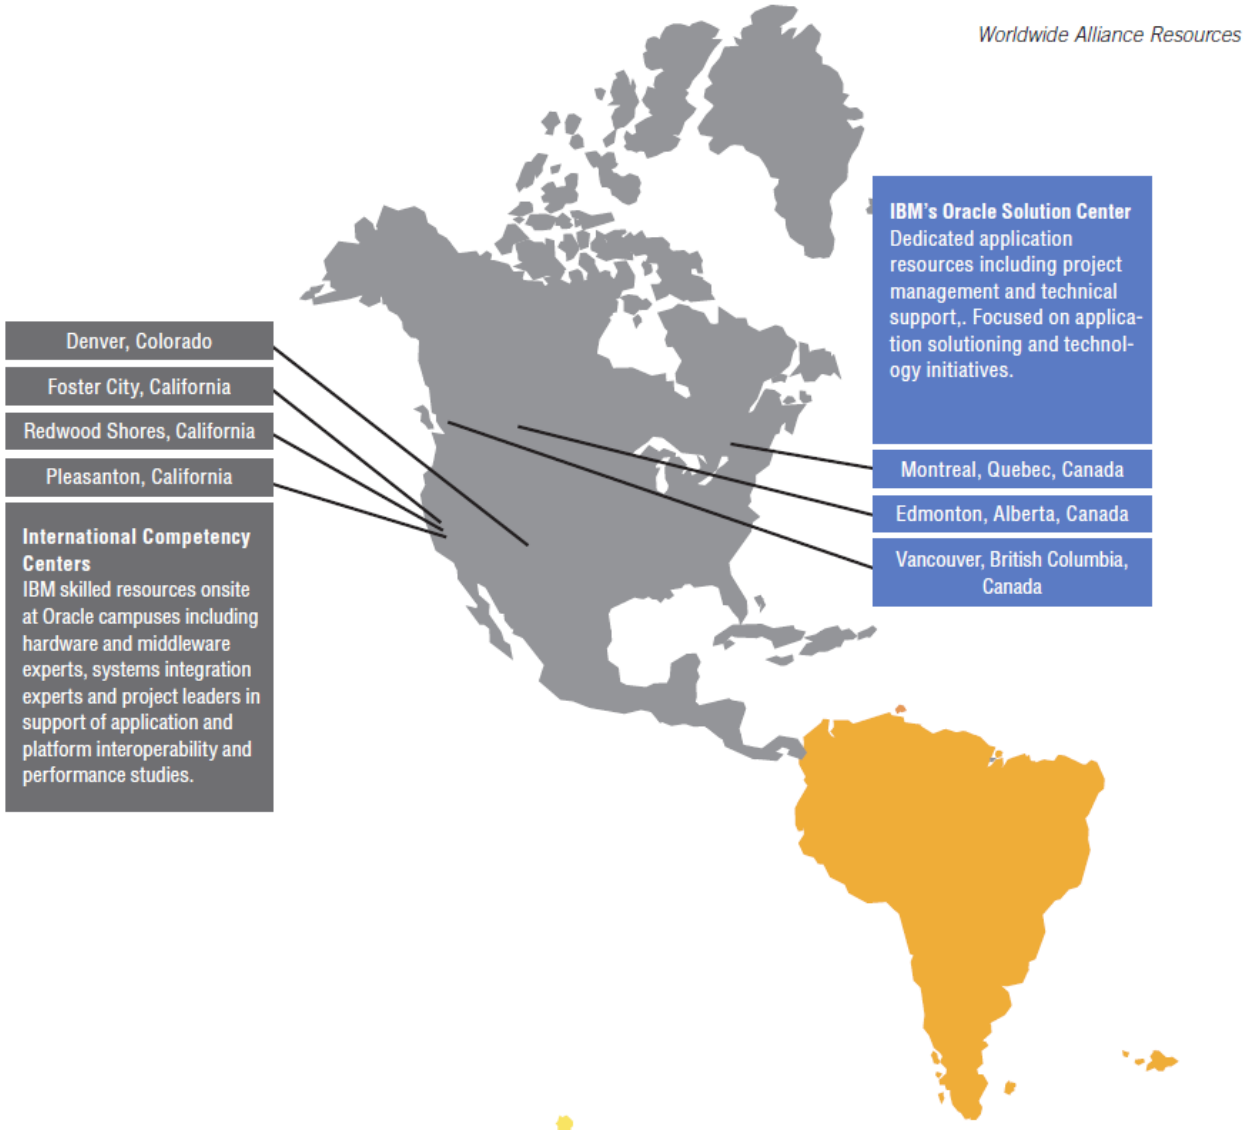

© 2013 IBM Corporation

IBM Oracle Competence Centers

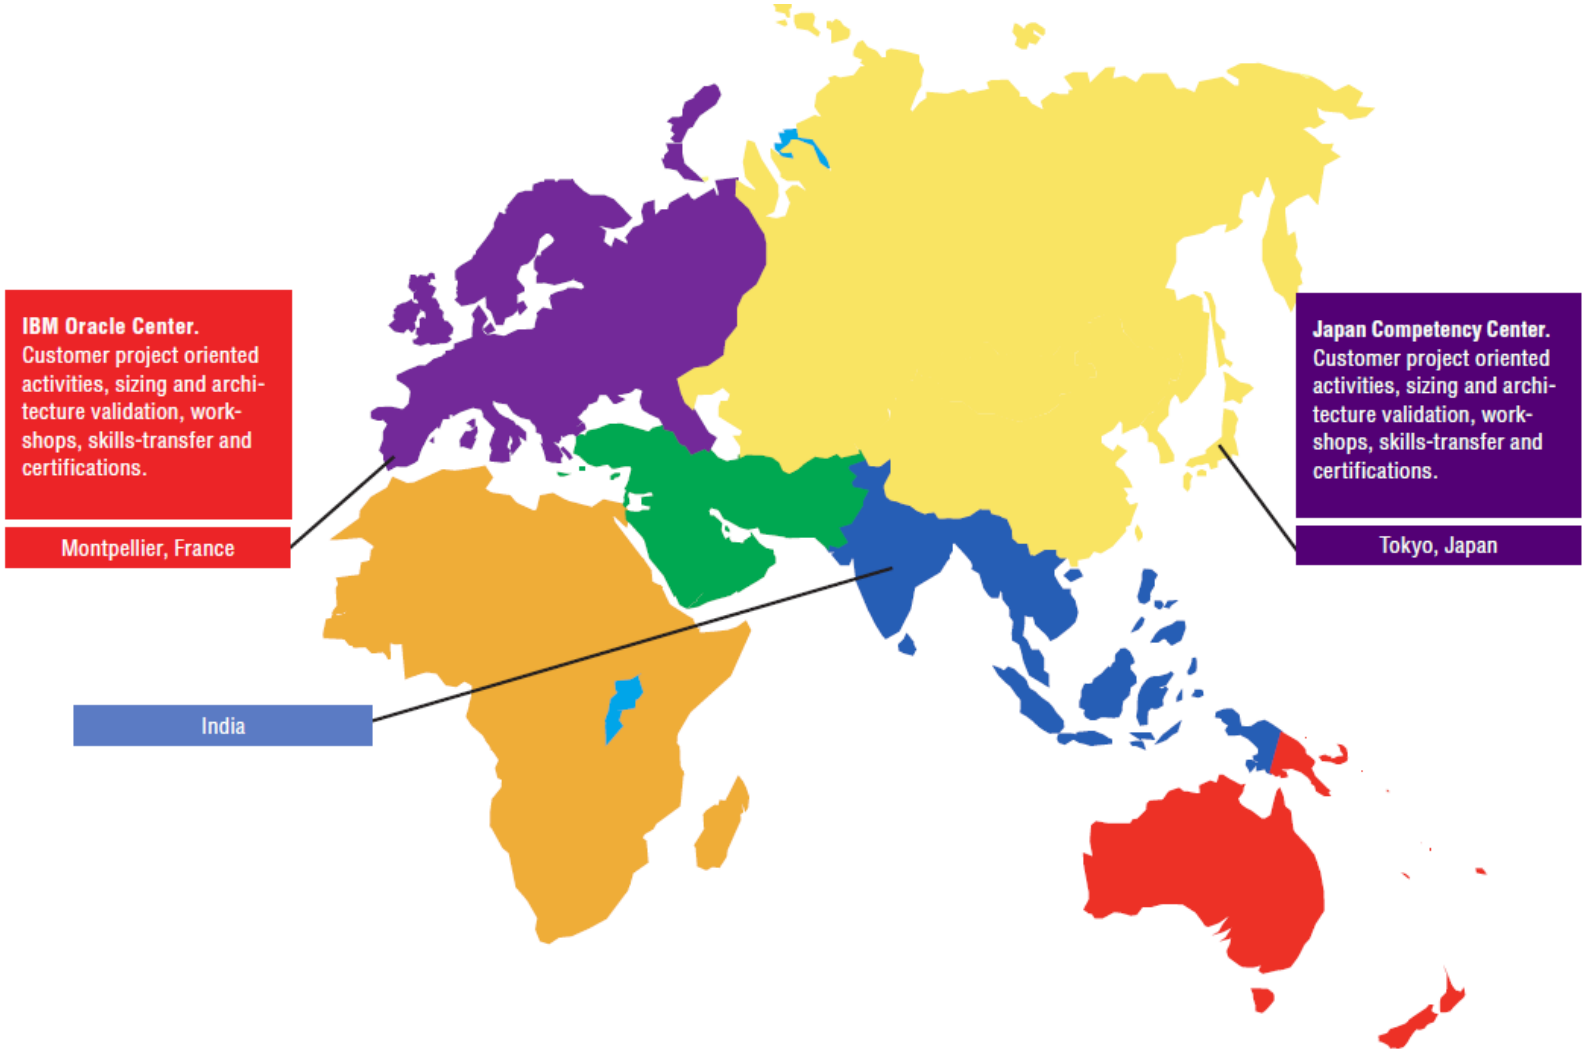

© 2013 IBM Corporation

The IBM Oracle Center (IOC)

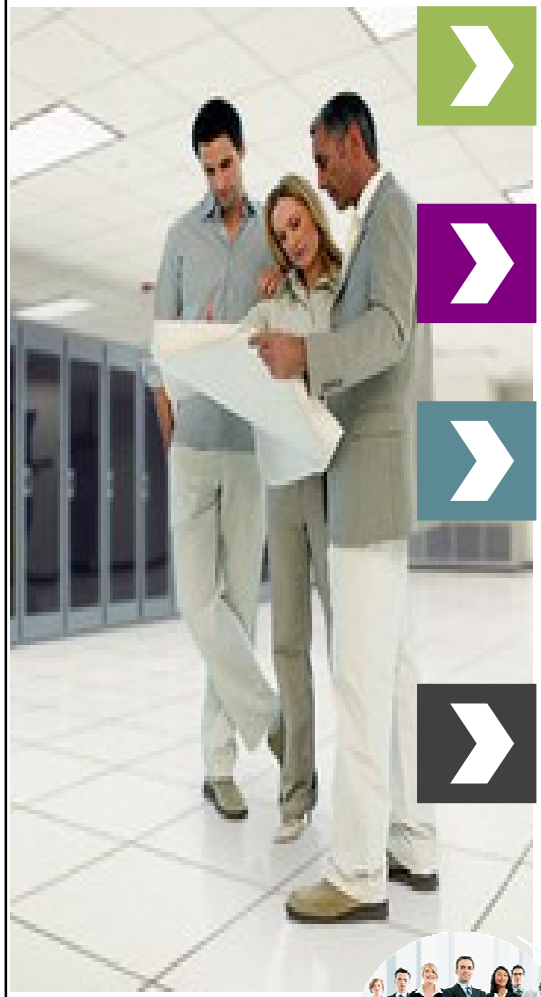

OUR MISSION

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: ioc@fr.ibm.com

© 2013 IBM Corporation

• Global Client Center, IBM Germany Research & Development Lab

IBM Executive Briefing Center (STG, SWG, GTS, GBS)

- Tailored Client Briefings and Workshops (local–remote–on-site)
- IT Conference and Business Shows support in Europe

WW STG Design Center

- Architecture and Design Support (local–remote–on-site)
- Innovation and Exploration Workshops

STG System Center

- STG Technical Support Center, SWG zTEC and IIC function
- Technical Consulting, Technical Workshops, Technology Demos, POCs, Scalability Tests, ISV Enablement

http://tmcceurope.de.ibm.com

Executive Briefing Center

PoC/Demo Portal

Access to System z Resources

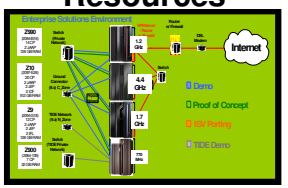

• Technical Expertise from IBM Developmentm, Lab Boeblingen

STG Lab Services

Providing key skills for post-sales delivery and deployments
Jump-Start' Support for z/BX, 'Client to Production'
Lab Services, System Networking and other areas

labserbb@de.ibm.com

Request for Technical Support

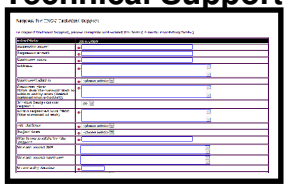

© 2013 IBM Corporation

Discussion Topics

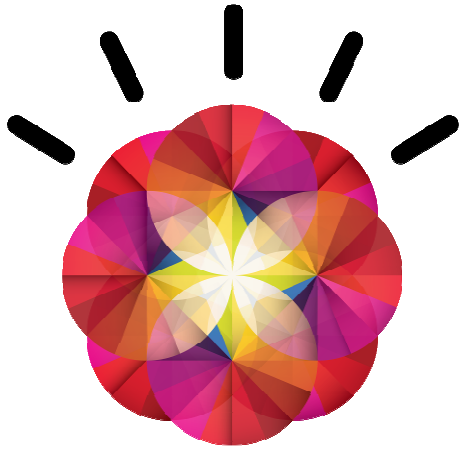

- IBM & Oracle relationship
- **Why Oracle on Linux on zEnterprise**
- Economics for consolidation
- Best practices for Oracle on Linux on System z
- Customer implementations

© 2013 IBM Corporation

Value of Oracle consolidation with System z

Do more with less

- Exploit the System z global virtualization with Oracle workload
- Consolidate servers, networks, 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

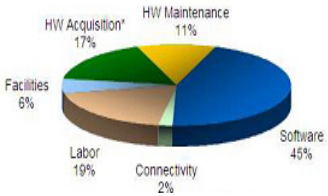

Reduce costs on a bigger scale

- **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 86 servers to a single IFL could potentially reduce licensing costs by as much as 97 percent.
- Minimize hardware needed for business continuance and disaster recovery

Manage growth and complexity

- 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 the Enterprise Linux Server means fewer cables, fewer components to impede growth

More flexibility, minimize lead time for new projects

- Consolidating Oracle and Linux environments to a single Enterprise Linux Server offers significant advantages in terms of flexibility
- Rapid provisioning reduces lead time for new IT projects, helping to increase business agility

© 2013 IBM Corporation

Why System z for Oracle

- High Availability Requirements
- Open Standards and Linux
- Disaster Recovery Requirements
- Scalability for growth
- Increased Performance Requirements
- Economics of Linux (IFL) Specialty Engines
- TCO and Total Cost of Acquisition advantage
- ‘Green’ Value from Mainframe
- zEnterprise servers can virtualize everything with up to 100% utilization rates
- zEnterprise can run Heterogeneous workloads
- System z has the highest security rating or classification for any commercial server

© 2013 IBM Corporation

Discussion Topics

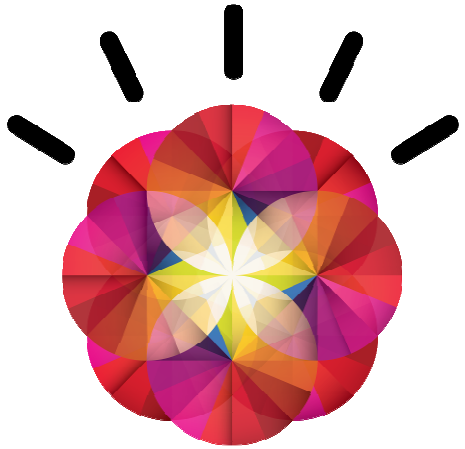

- IBM & Oracle relationship
- Why Oracle on Linux on zEnterprise
- **Economics for consolidation**
- Best practices for Oracle on Linux on System z
- Customer implementations

© 2013 IBM Corporation

Why High-End Servers?

Utilization on x86 systems

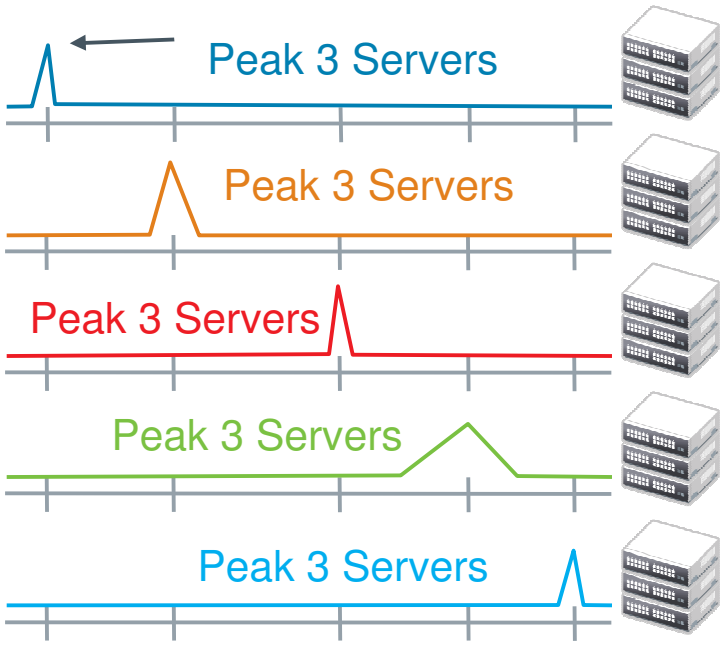

According to a study by Gartner, data centers that do not use virtualization have an average server CPU utilization rate of only 15%.

Mixed Utilization on IBM High End Servers

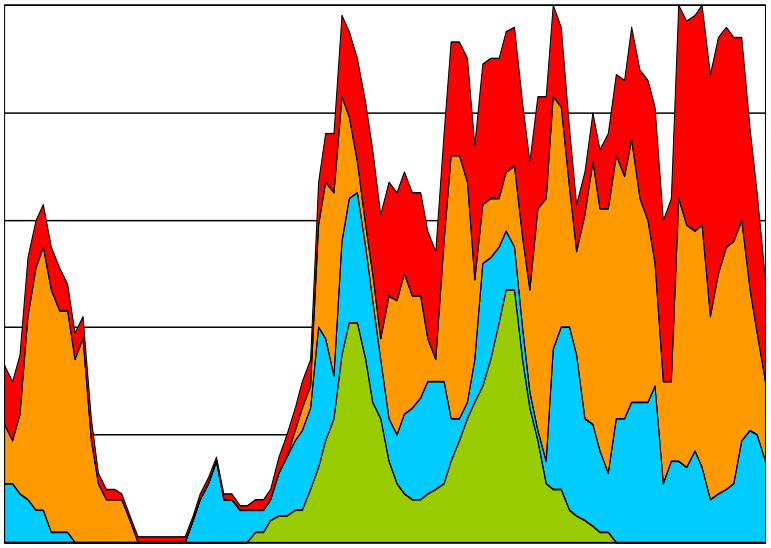

IBM High End Server: Up to 100% utilization

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

© 2013 IBM Corporation

Oracle Database Support for z/VM Server Virtualization

http://www.oracle.com/technetwork/database/enterprise-edition/db-virtualization-

VENDOR	Product	Virtualization Type	Processor Platform	Operating System(s)	Hard Partition Licensing	Basic Support		Live Migration	
						DB	RAC	DB	RAC
Oracle	Oracle Virtual Machine	Virtual Machine (PV, HVM, Hybrid)	X86, X64	Linux	Pinned Only	Yes	Yes [2,3]	Yes[1]	Test
		Virtual Machine (HVM, Hybrid)		Windows	Pinned only	Test	No	Test	No
HP	nPar	Physical Partition	IA, PA	HP-UX	Yes	Yes	Yes	N/A	N/A
	vPar	Virtual Partition	IA, PA	HP-UX	Yes	Yes	Yes	N/A	N/A
	Integrity Virtual Machines	Virtual Machine	IA	HP-UX	Capped only	Yes	No	N/A	N/A
IBM	Dynamic LPAR	Virtual Partition	Power	AIX	Yes	Yes	Yes	Yes	Test
	Micro-partitions	Shared Partition	Power	AIX	Capped only	Yes	Yes	Yes	Test
	VIO Server	Shared I/O	Power	N/A	N/A	Yes	Yes	N/A	N/A
	WPAR	OS Partition	Power	AIX	No	Yes	No	No	No
	zVM	Shared Partition	zSeries	Linux	No	Yes	Yes	N/A	N/A
SUN	Dynamic System Domains	Physical Partition	Sparc	Solaris	Yes	Yes	Yes[2]	N/A	N/A
	Logical Domains	Virtual Partition	Sparc	Solaris	No	Yes	Yes	N/A	N/A
	Local Solaris Containers	OS Partition	Sparc, X86, X64	Solaris	Capped only	Yes	Yes	N/A	N/A

NO VMWARE KVM XEN Support

© 2013 IBM Corporation

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

© 2013 IBM Corporation

Pricing based on PVU – Processor Value Units

PVU Table per Core (section 1 of 2 - RISC and System z)

Processor Technologies												
	Processor Brand			Processor Type								
				Cores per socket								
Processor Vendor	Processor Name	Server model numbers	Maximum number of sockets per server	One-Core (1)	Dual-Core (2)	Quad-Core (4)	Hexa-Core (6)	Octi-Core (8)	16-Core (16)	IFL Engine	Proc. Model Number	PVUs per Core
IBM	POWER7 [4]	770,780,795	> 4			■	■	■			All	120
		750,755,775 PS704, p460	4			■	■	■			All	100
		PS700-703, 710-740, p260, 7R2, p24L	2			■	■	■			All	70
	POWER6	550,560,570, 575,595	All		■						All	120
		520, JS12,JS22, JS23,JS43	All		■						All	80
	POWER5, POWER4	All	All		■						All	100
	POWER5 QCM	All	All			■					All	50
	z196, System z10 [1,5]	All	All							■	All	120
	z114, System z9 z990, S/390 [1,2,6]	All	All							■	All	100
	PowerPC 970	All	All		■						All	50
	PowerXCell™, Cell/B.E.™ 8i [3]	All	All	■							All	30
Any	Any single-core	All	All	■							All	100

Processor Technologies												
	Processor Brand			Processor Type								
				Cores per socket								
Processor Vendor	Processor Name	Server model numbers	Maximum number of sockets per server	One-Core (1)	Dual-Core (2)	Quad-Core (4)	Hexa-Core (6)	Octi-Core (8)	16-Core (16)	IFL Engine	Proc. Model Number	PVUs per Core
HP / Intel®	Itanium® 1,2	All	All		■	■					All	100
	PA-RISC	All	All		■						All	100
Sun / Fujitsu	SPARC64 VI, VII	All	All		■	■					All	100
	UltraSPARC IV	All	All		■						All	100
	SPARC T4	T4-4	4					■			All	100
		T4-1, T4-1B, T4-2	2					■			All	70
	SPARC T3	All	All					■	■		All	70
	UltraSPARC T2	All	All			■	■	■			All	50
	UltraSPARC T1	All	All			■	■	■			All	30
Any	Any single-core	All	All	■							All	100

* Requirements as of Publish Date: 17 May 2012

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

© 2013 IBM Corporation

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

Oracle Database

Database Products	Processor License	Software Update License & 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
Real Application Clusters One Node	10,000	2,200.00
Active Data Guard	10,000	2,200.00

Prices in USA (US Dollars)

Enterprise Edition 47,500 10,450.00 US$ annually

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

Intel Itanium 1.0

IBM System z (z10 and z114) 1.0

© 2013 IBM Corporation

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

© 2013 IBM Corporation

Oracle database consolidation example

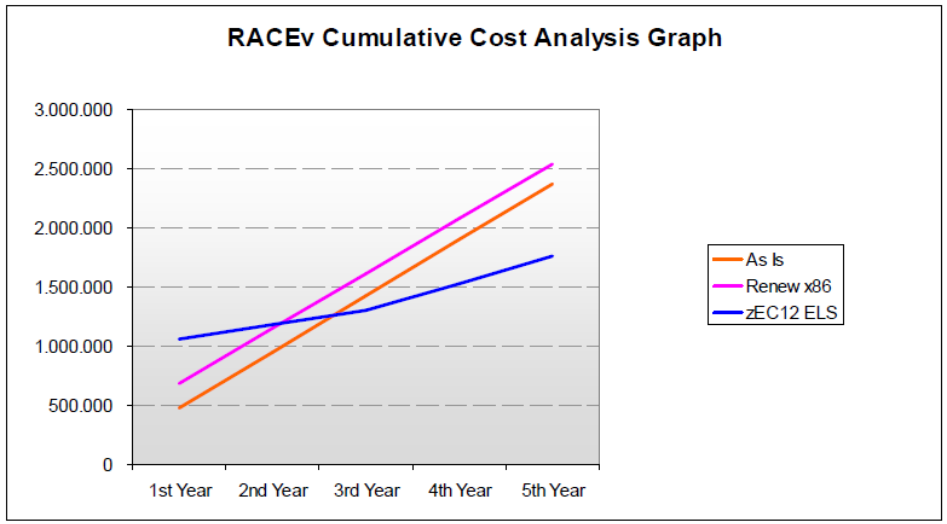

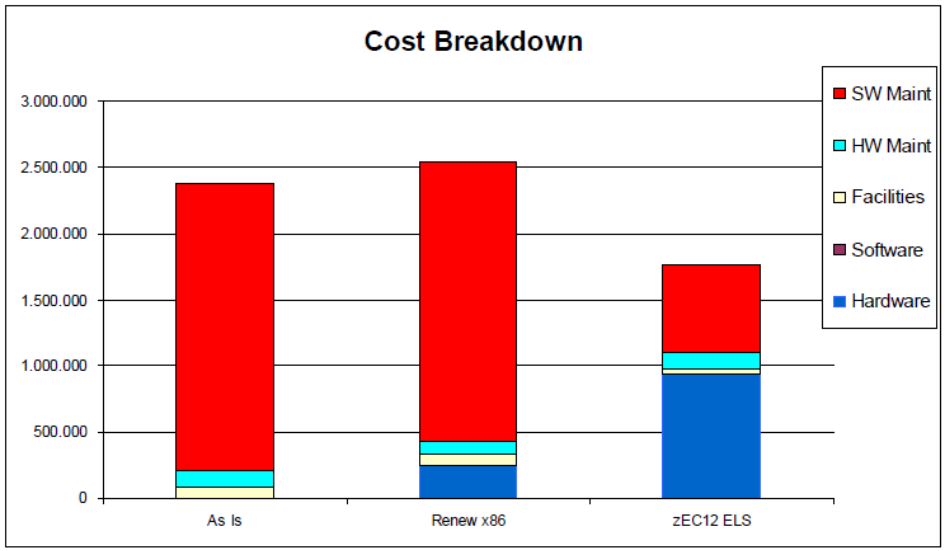

ProLiant BL260c G5 Xeon E5430
Quad Core 2.66GHz (1ch/4co)
- 50 physical Linux servers
  - 10 servers @ 25% utilization
  - 20 servers @ 15% utilization
  - 20 servers @ 10% utilization
- Oracle DB Standard Edition

Renewal: ProLiant BL280c G6 Xeon X5672 Quad Core 3.2GHz (1ch/4co)

Sizing information based on IBM internal RACEv tool and US prices. Actual results may vary.

© 2013 IBM Corporation

z/VM Virtualization: DCSS & NSS Technology Exploitation

- economic for maintenance and sharing

- **DCSS (Discontiguous Saved Segments)**
  - Can define an address range (MB boundary) to the system
  - **A single copy will exist and is shared among all users**
  - Virtual Machine "loads" the DCSS dynamically
  - Maps DCSS into its address space
  - Can be located outside virtual machine's defined storage
  - DAT architecture allows this to work with minimal CP involvement
  - Used to contain
  - Data (e.g., file system control blocks)
  - Code (e.g., CMS code libraries)

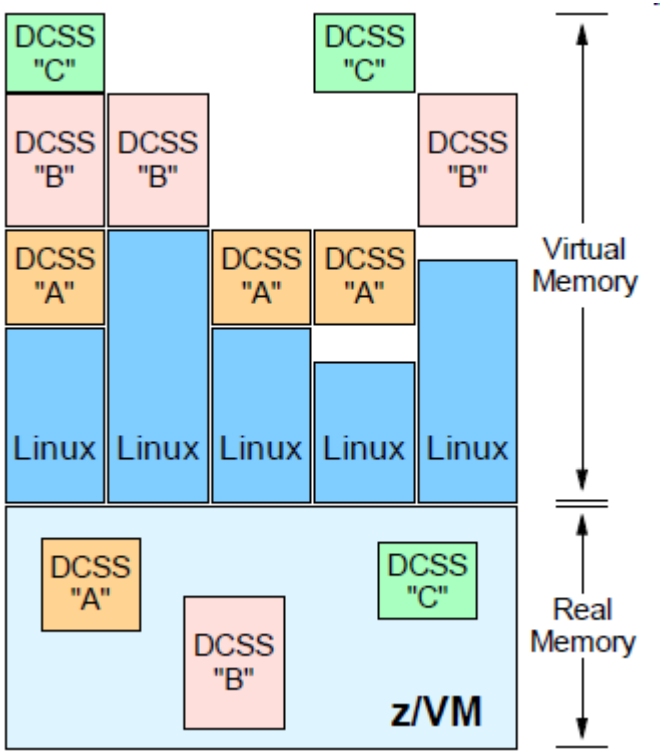

- **NSS (Named Saved Systems)**
  - Place kernel code in a segment
  - Able to IPL the NSS (boot the NSS)
  - 1 shared copy on system for N virtual machines instead of N copies

© 2013 IBM Corporation

Oracle HA with System z

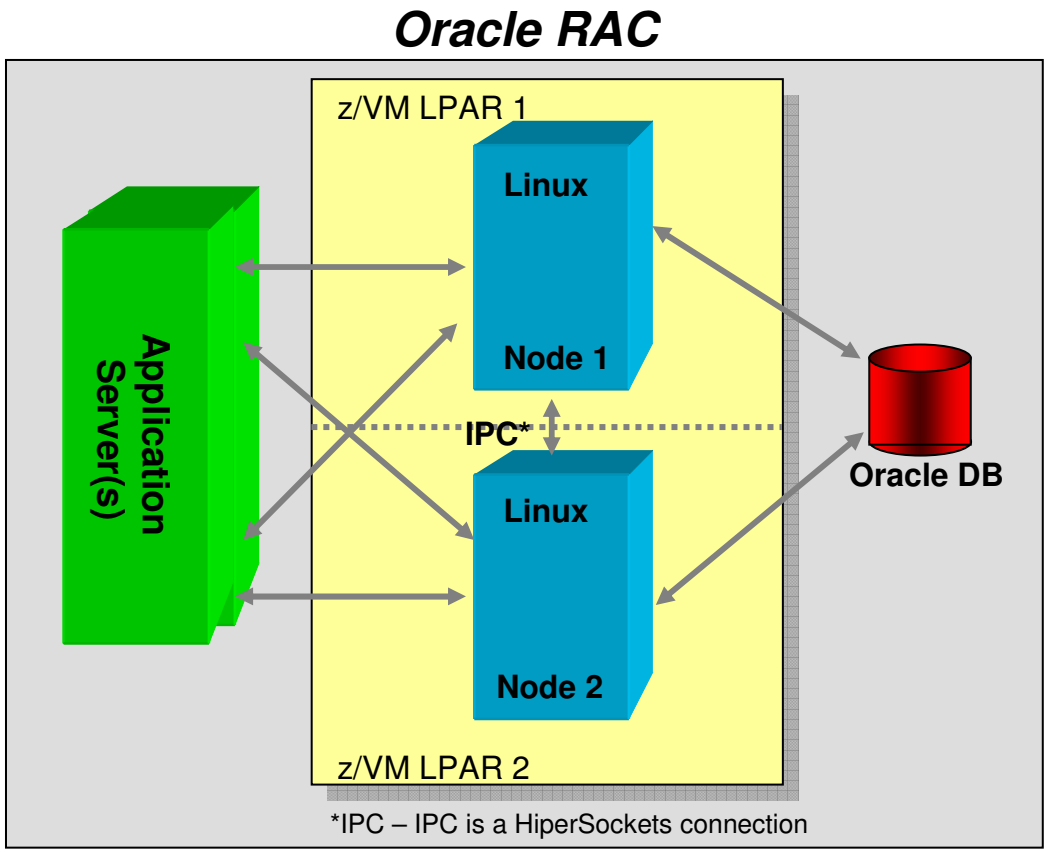

Server provided HA

Oracle

- **RAC**
- **Data Guard**
- **Flashback**
- **CRS**
- **Grid Control**

Operating system HA

Linux

- **Linux Clustering**

z/VM

- Mature Hypervisor
- Hardware assist
- z/VM SSI/LGR

Hardware provided HA

System z

- **Share CPUs in 2 LPARs**
- Spare CPUs avail
- N+1 power supplies
- Chip sparring in memory
- Concurrent maintenance

- Guards against Linux failure, LPAR failure, z/VM failure, Oracle instance failure, LPAR maintenance
- Can be: Active/active, active/passive
- Not limited to two nodes

© 2013 IBM Corporation

Possible DR architecture, Oracle based (hot, warm)

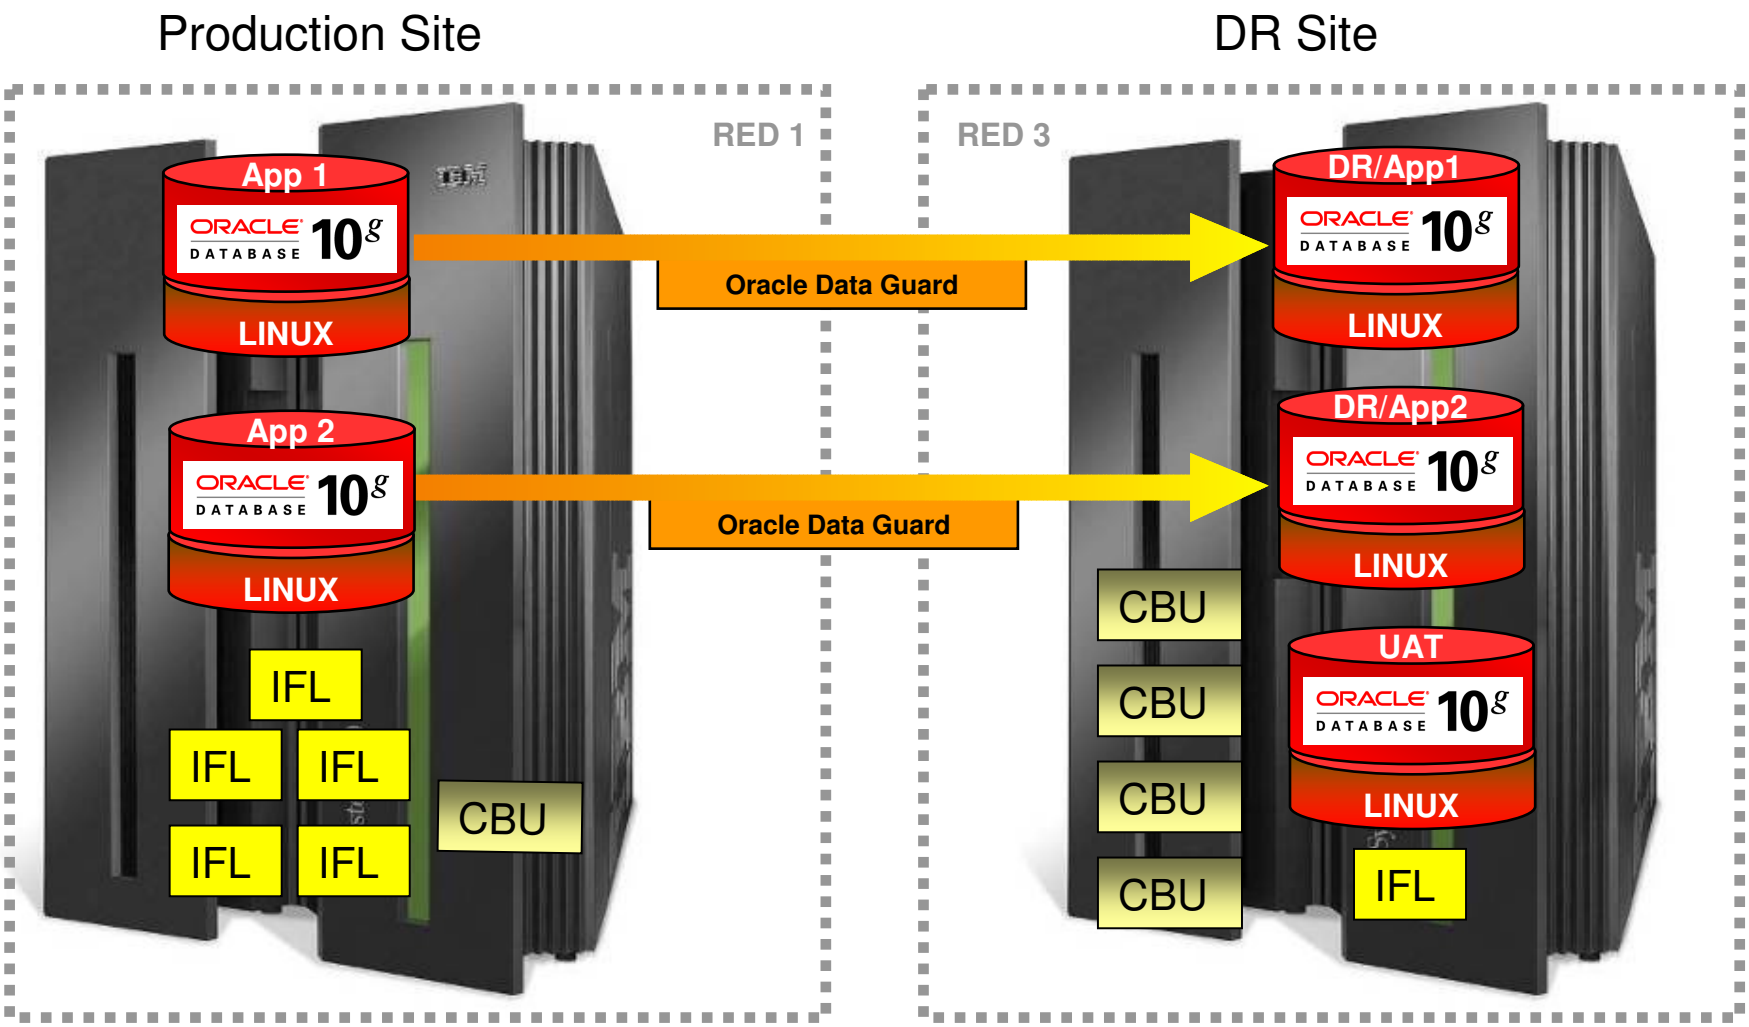

© 2013 IBM Corporation

Discussion Topics

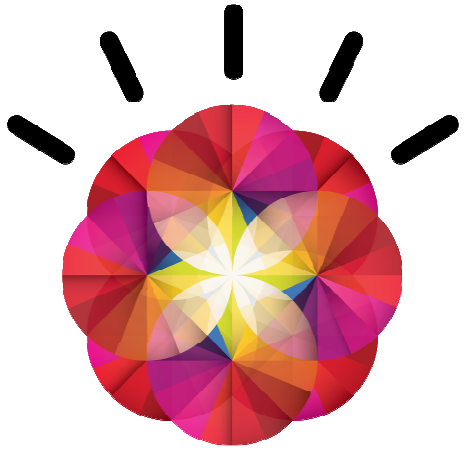

- IBM & Oracle relationship
- Why Oracle on Linux on zEnterprise
- Economics for consolidation
- **Best practices to start Oracle consolidation on Linux on System z**
  - Existing environment
  - Feasibility Study (Proof of Concept)
  - Sizing for System z
  - HA and DR
  - Oracle database migration services
- Customer implementations

© 2013 IBM Corporation

Discussion Topics

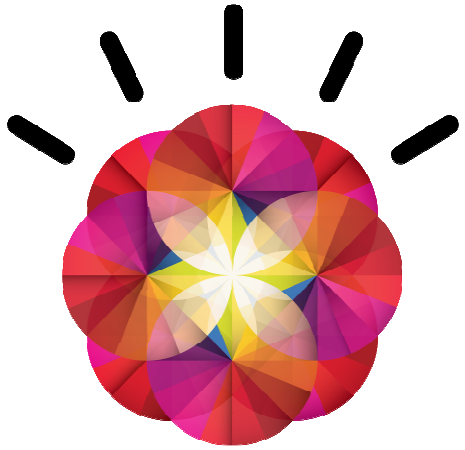

- IBM & Oracle relationship
- Why Oracle on Linux on zEnterprise
- Economics for consolidation
- **Best practices to start Oracle consolidation on Linux on System z**
  - **Existing environment**
  - Feasibility Study (Proof of Concept)
  - Sizing for System z
  - HA and DR
  - Oracle database migration services
- Customer implementations

© 2013 IBM Corporation

Evaluation steps for Consolidation projects

- **1 - Feasibility study**
  - Gather information on existing environment (servers, applications, network)
  - Prioritize the non-functional requirements (RAS, scalability, performance, management)
  - Fit for Purpose study
  - Cost & Value study
  - Assess the skills
- **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 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
- **4 - Pre-production tests**
  - Validation in the real environment
  - Environment health check before production (LPAR, z/VM, Linux, Middleware)
  - Skill transfer phase
- **5 – Launch to production**
  - Iterative put to production
  - Monitor the system to tune it accordingly
  - Reserve time for tuning

© 2013 IBM Corporation

Gather information about existing 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

© 2013 IBM Corporation

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%

© 2013 IBM Corporation

Oracle Consolidation on System z study : methodology example

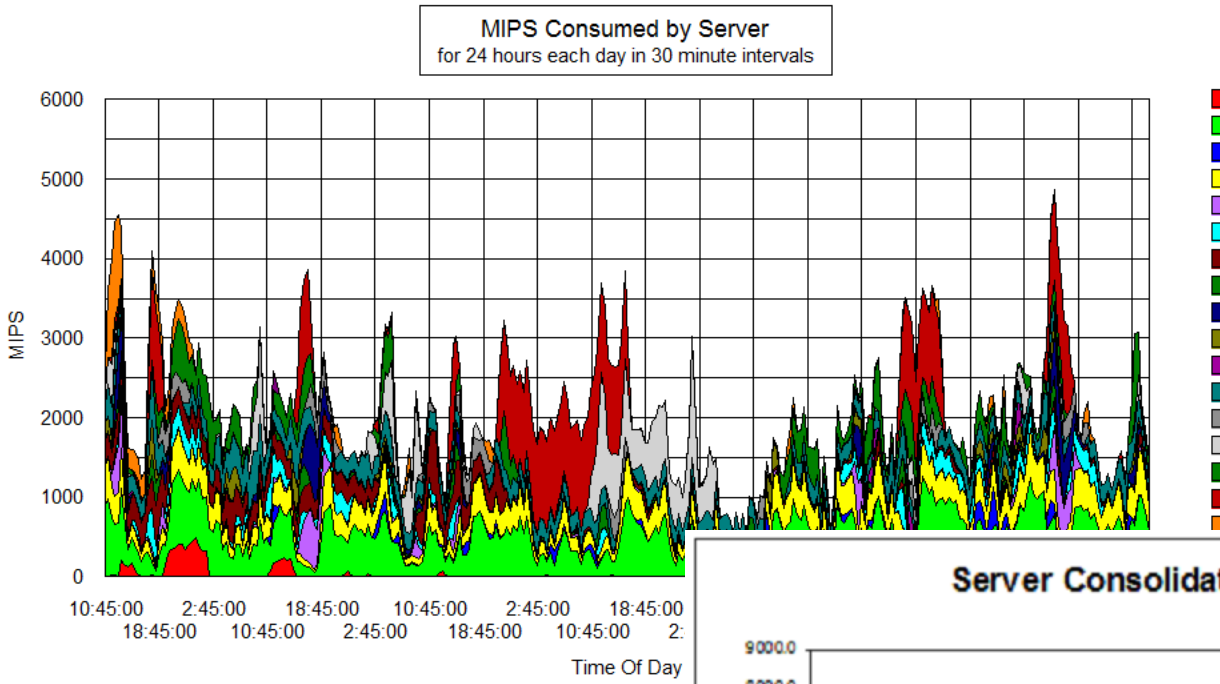

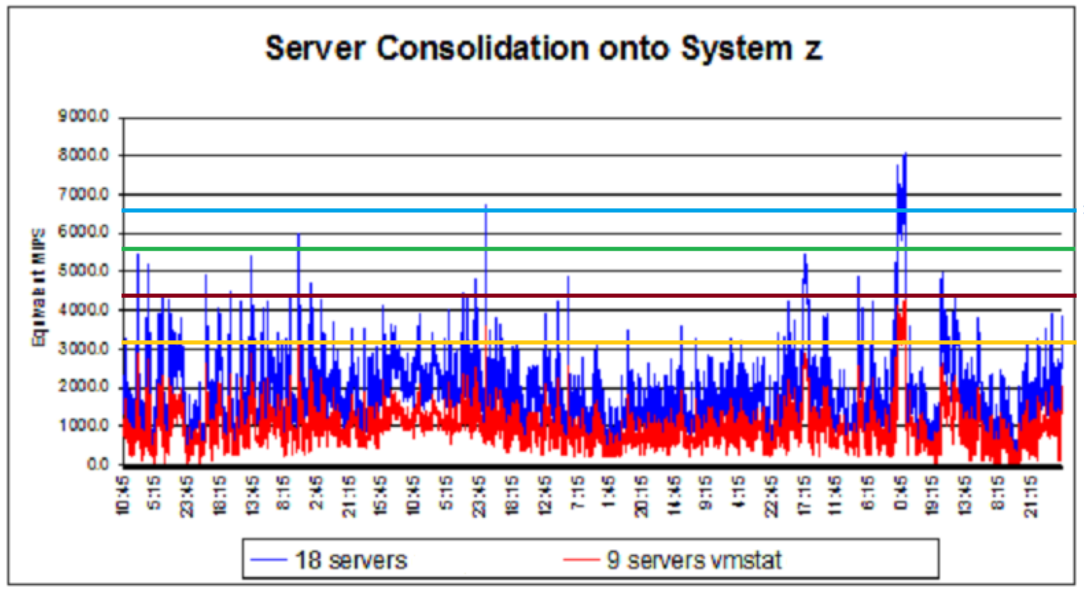

© 2013 IBM Corporation

Discussion Topics

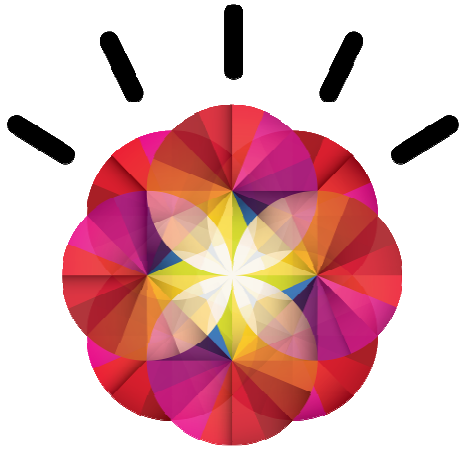

- IBM & Oracle relationship
- Why Oracle on Linux on zEnterprise
- Economics for consolidation
- **Best practices to start Oracle consolidation on Linux on System z**
  - **Existing environment**
  - **Sizing for System z**
  - Feasibility Study (Proof of Concept)
  - HA and DR
  - Oracle database migration services
- Customer implementations

© 2013 IBM Corporation

Oracle Database memory sizing for System z

- 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*
  SGA=System Global Area, PGA=Program Global Area
- 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 one guest
- System z memory = real memory for guests + memory for z/VM and expanded storage (2GB or 4 GB).

* Increase estimate when Oracle SGA is large and there are expected to be hundreds of dedicated server connections or use hugepages with Oracle 11gR2. A large overall virtual storage requirement may result in larger Page Tables in Linux which require extra guest storage. Consider hugepages but no AMM.

© 2013 IBM Corporation

Discussion Topics

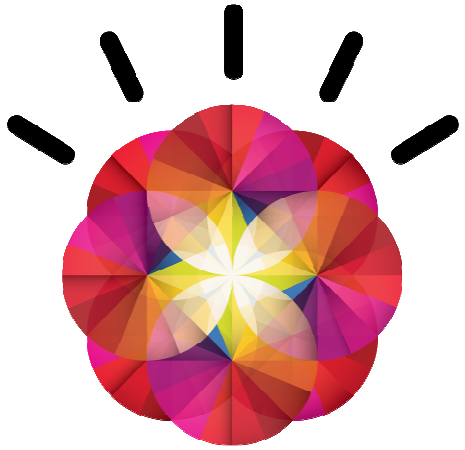

- IBM & Oracle relationship
- Why Oracle on Linux on zEnterprise
- Economics for consolidation
- **Best practices to start Oracle consolidation on Linux on System z**
  - **Existing environment**
  - **Sizing for System z**
  - **Feasibility Study (Proof of Concept)**
  - HA and DR
  - Oracle database migration services
- Customer implementations

© 2013 IBM Corporation

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
    - Storage Server and topology
    - 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!

© 2013 IBM Corporation

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!

© 2013 IBM Corporation

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;

© 2013 IBM Corporation

Best Practices: Use Oracle 11*g* OLTP improvements

Comparison: Oracle 10*g* versus 11*g* database

User scaling – transactional through-put

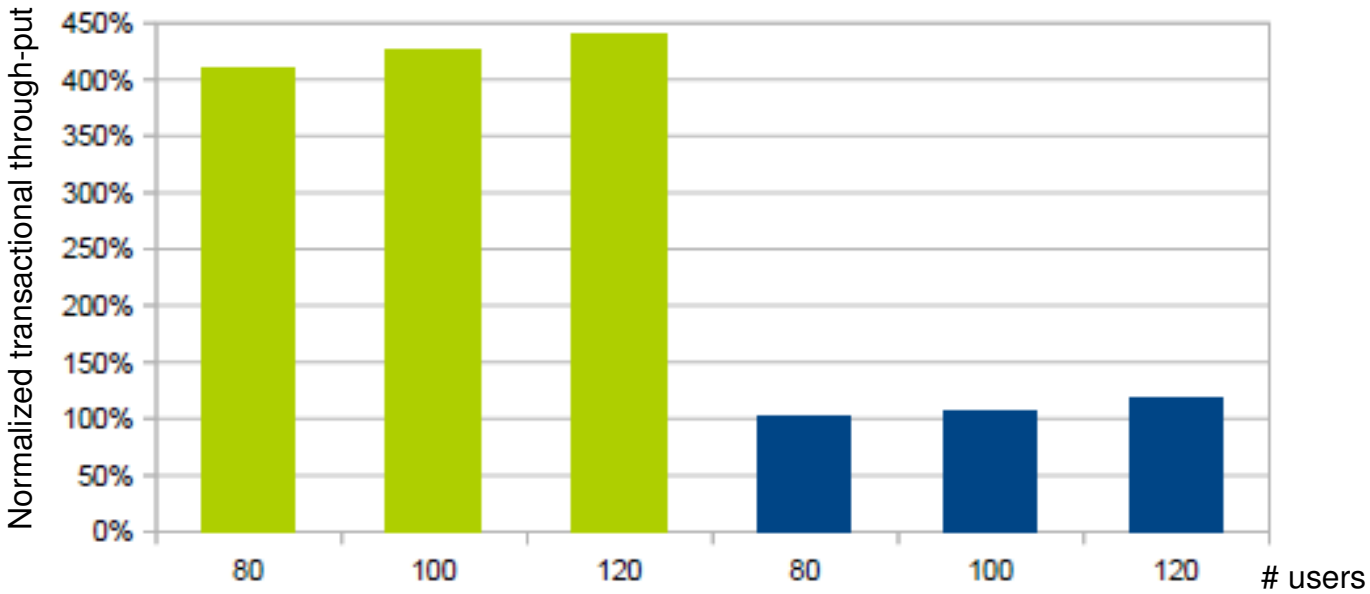

Recommendation: upgrade to 11*g*R2 if not already done

© 2013 IBM Corporation

Discussion Topics

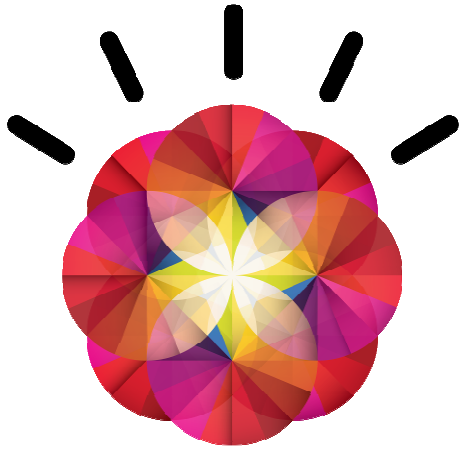

- IBM & Oracle relationship
- Why Oracle on Linux on zEnterprise
- Economics for consolidation
- **Best practices to start Oracle consolidation on Linux on System z**
  - **Existing environment**
  - **Sizing for System z**
  - **Feasibility Study (Proof of Concept)**
  - **Define Non-Functional Requirements: i.e. HA and DR**
  - Oracle database migration services
- Customer implementations

© 2013 IBM Corporation

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

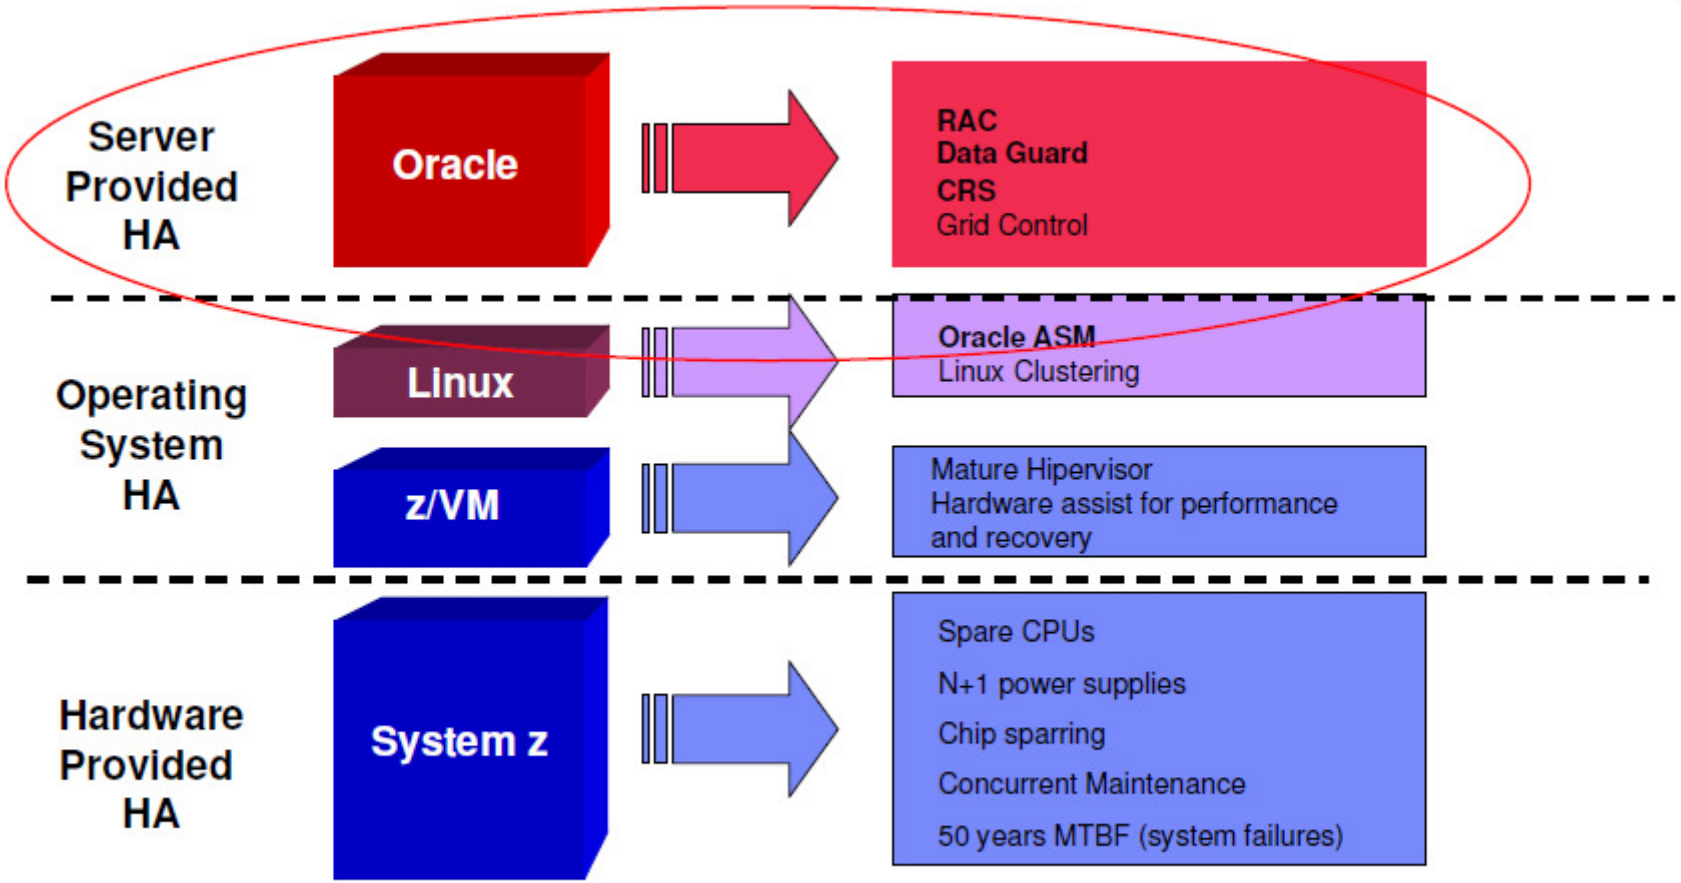

© 2013 IBM Corporation

Clustering (hot-standby or hot-hot)

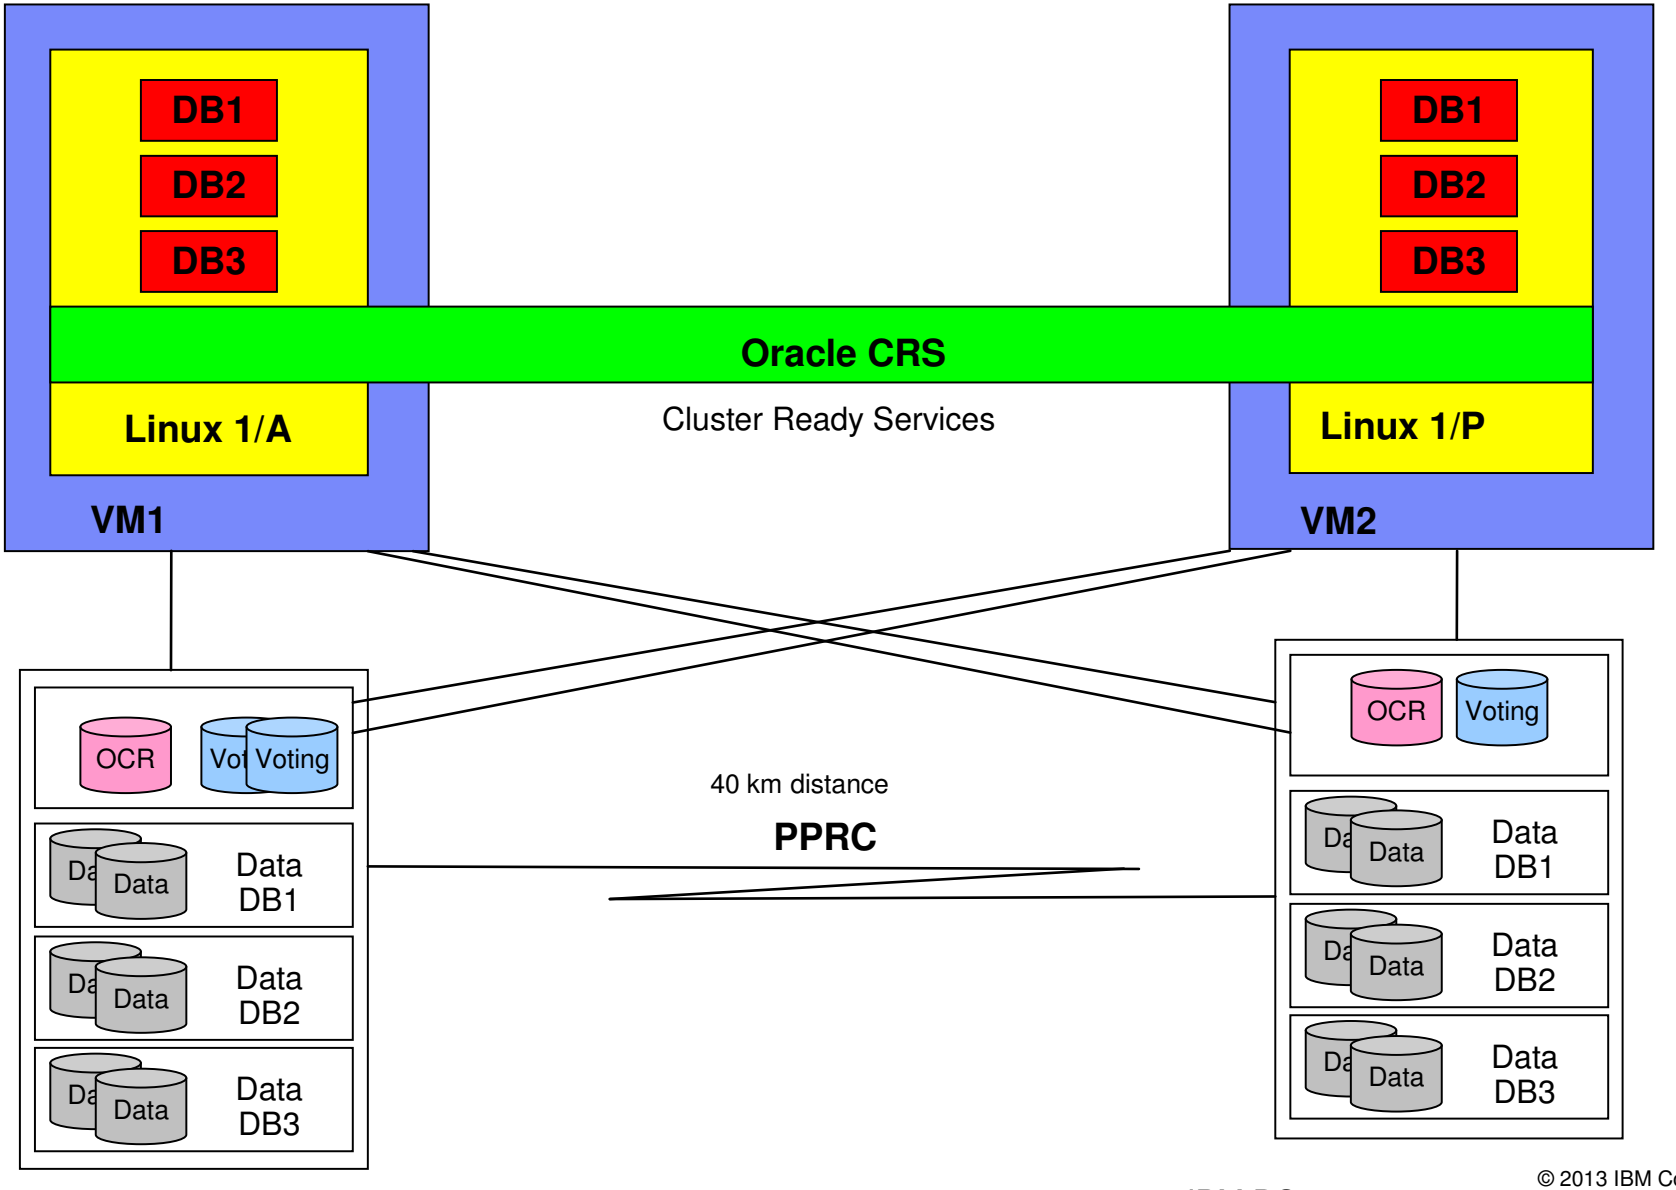

© 2013 IBM Corporation

Disaster Recovery Database – Data Guard

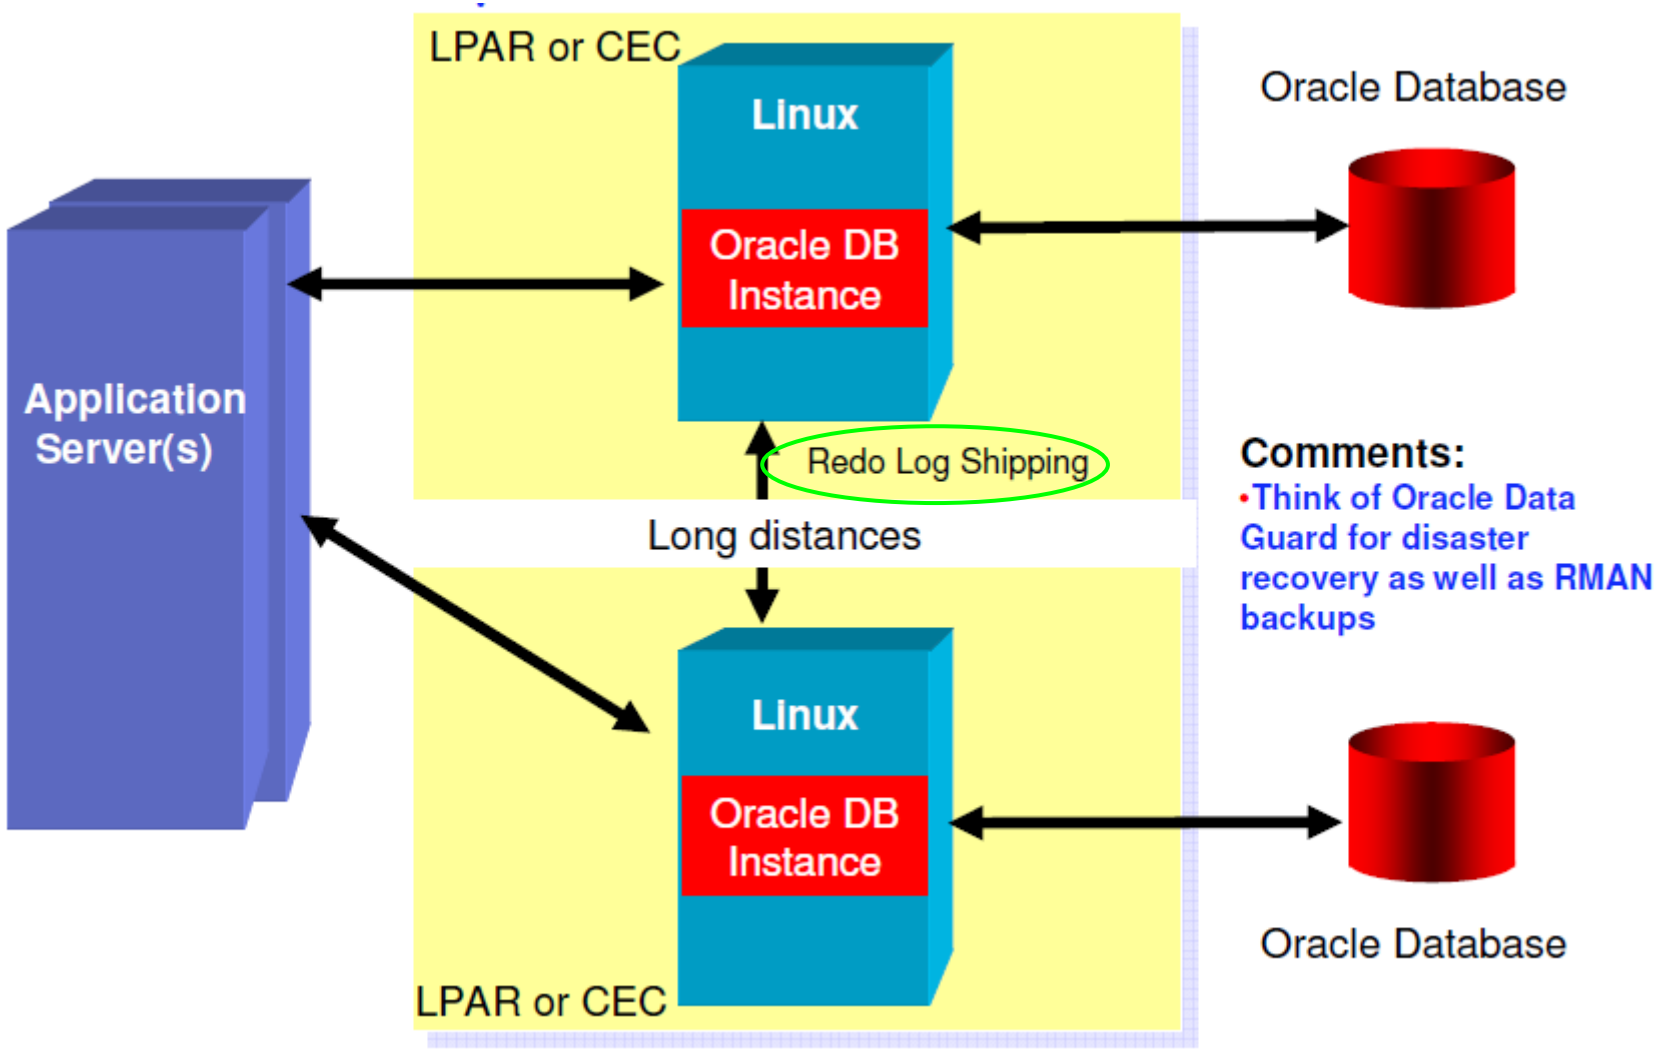

© 2013 IBM Corporation

Disaster Recovery Database – Data Guard

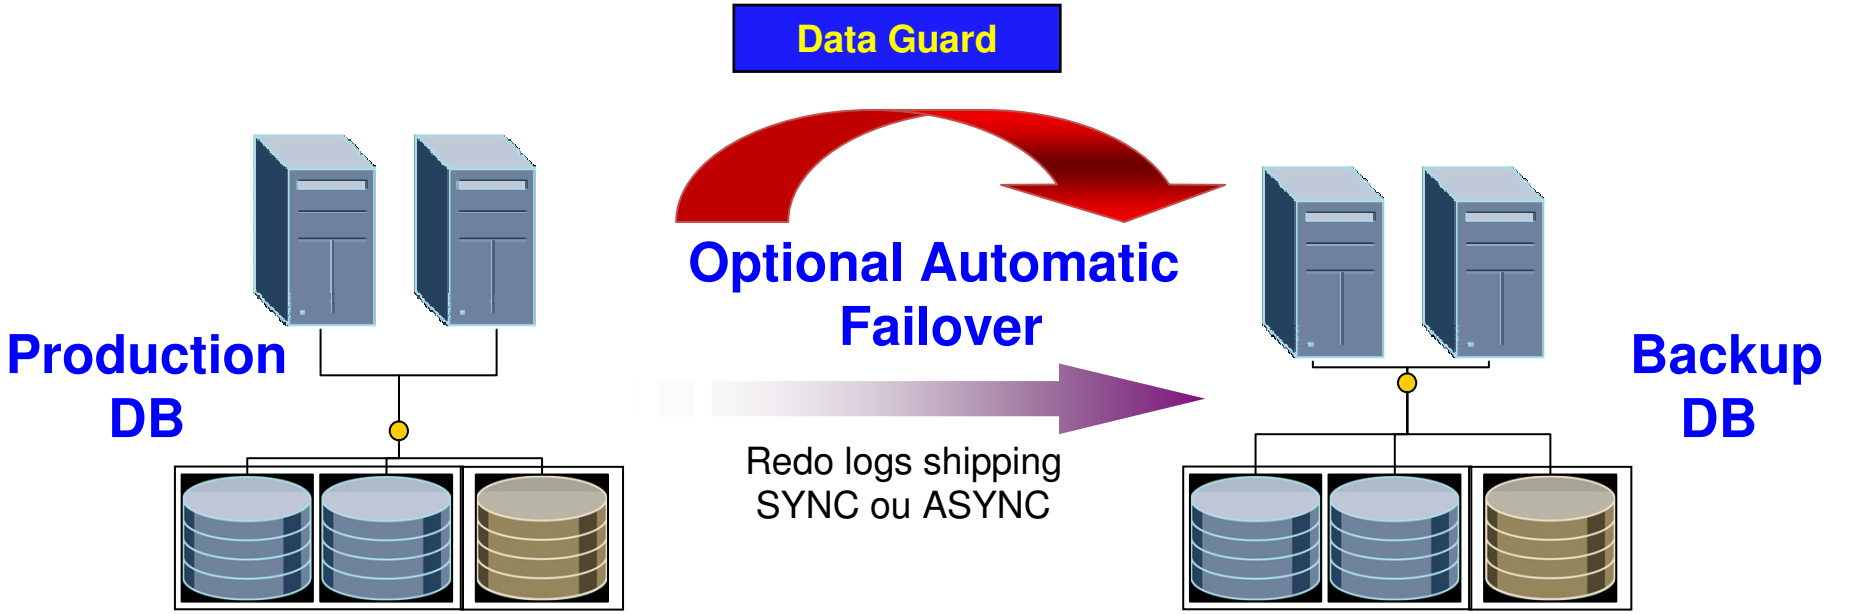

Oracle Data Guard is a database software infrastructure that automates the creation and maintenance of a duplicate, or standby copy, of the production (or primary) database. If the primary database becomes unavailable (disasters, maintenance), the standby database can be activated and can take over the data serving needs of the enterprise.

Oracle Data Guard Configuration

Oracle Dataguard maintains standby databases as consistent copies of the production database.

These standby databases can be located at remote disaster recovery sites thousand of miles away from the production data center, or they can be located at the same city, campus, or even in the same building.

Oracle Data Guard can switch any standby database to the primary role, minimizing the downtime associated to the outage and preventing any data loss.

© 2013 IBM Corporation

Discussion Topics

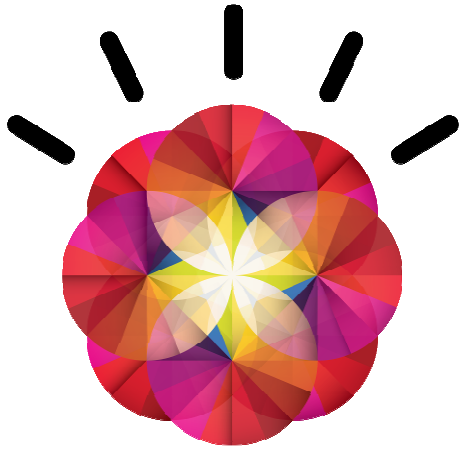

- IBM & Oracle relationship
- Why Oracle on Linux on zEnterprise
- Economics for consolidation
- Best practices to start Oracle consolidation on Linux on System z
  - Existing environment
  - Sizing for System z
  - Feasibility Study (Proof of Concept)
  - HA and DR
  - **Oracle database migration services**
- Customer implementations

© 2013 IBM Corporation

Migration Factory

- **Server Workload Migration**
  - Migration Estimating
    - Migration Factory in a Box (MiaB) – a suite of pre-defined Migration Factory services based on a set of pre-determined assumptions, constraints and scope criteria that define the 'Box'
    - Oracle and SAP Makeover – based on a combination of historical & customer-provided data we develop a high-level technical and financial roadmap for migrating Oracle databases, Oracle E-Business Suite and SAP for common scenarios
  - Migration Analysis & Planning
    - Rapid Assessment – high-level analysis of low to medium complexity workloads
    - Full Assessment – detailed analysis of medium to high complexity workloads
  - Workload Migration
    - Custom Code Porting/Migration
    - Database Migration, Conversion, Upgrade
    - ERP Application Migration, Upgrade

- Migration Portal www-03.ibm.com/systems/migratetoibm
- Migration Factory www-03.ibm.com/systems/migratetoibm/factory
- Migration Factory Expert Videos www-03.ibm.com/systems/migratetoibm/videos.html
- developerWorks www-128.ibm.com/developerworks/ondemand/migrate

© 2013 IBM Corporation

Migration Factory: XenoBridge

XenoBridge is an IBM Migration Factory database migration tool that automates the process of database migration

- Provides advanced parallel processing techniques to migrate databases in a minimum of downtime
- Utilizes vendor-supplied utilities to ensure data integrity is maintained during the migration process

By wrapping vendor-supplied utilities with an intelligent migration engine, the throughput obtained by XenoBridge is typically a factor of 100 times that of traditional methodologies

Data is transferred across a network – meaning that the location of the source and target is transparent

XenoBridge is full featured

- Not only move the rows in a database rapidly
- Also for homogenous migrations – it can build the target database based on the object definitions contained in the source database … for example …
  - Oracle 8 → Oracle 10
  - Oracle → DB2 / Sybase
  - Informix → DB2

Email: migr8te@us.ibm.com

© 2013 IBM Corporation

Discussion Topics

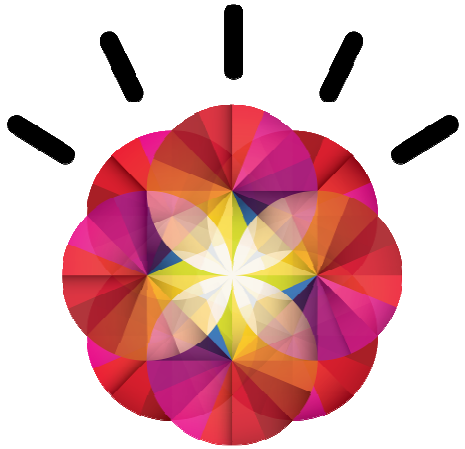

- IBM & Oracle relationship
- Why Oracle on Linux on zEnterprise
- Economics for consolidation
- Best practices for Oracle on Linux on System z
- **Customer implementations**

© 2013 IBM Corporation

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

- System z customer consolidation example:
  - Oil and 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 – z10 BC, 3 IFLs, 24 GB
    - SLES10, Oracle 10g EE
    - Databases: 7 production, 400 GB – 3 TB, 7 virtual servers/database

© 2013 IBM Corporation

Start of a 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
  - Plan is to go towards a « Cloud » environment

© 2013 IBM Corporation

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

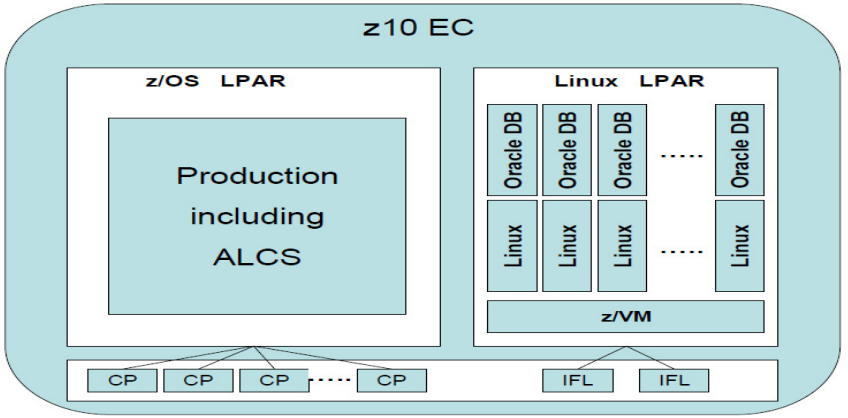

- Results
  - The Poc was done at their site with the help of BOE people (on site) and Mop people (remotely)
  - After this first step, a workshop was done at their site for HA/DR with Oracle DB on System z environment
  - Customer now in production

© 2013 IBM Corporation

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

© 2013 IBM Corporation

Optimized Computing: Nationwide's Linux on System z Virtualization Reduces TCO and Time to Deploy

Nationwide was already experiencing serious technology pain points from the continuous growth of its business. Among these were:

- *Too many distributed physical servers with low utilization*
- *A lengthy provisioning process that delayed the implementation of new applications for headquarters and agencies, and for new customers for Nationwide's human resources outsourcing business*
- *Limitations in data center power and floor space*
- *High Total Cost of Ownership (TCO)*
- *Difficulty allocating processing power for a dynamic environment.*

> *"Our goal was server optimization and our approach was virtualization."*
>
> — Guru Vasudeva, Nationwide vice president and CTO

TCO results that Nationwide has experienced:

- Monthly Web hosting fees have gone down by 50 percent.
- Hardware and operating system support needs have decremented by 50 percent.
- CPU utilization is up an average of 70 percent, with the elimination of many physical servers with below average utilization.
- Middleware licensing costs for WebSphere, Oracle, and UDB have dramatically fallen.
- There has been an 80 percent reduction in data center floor space needs, and power consumption is down.
- The net of the effort is a $15 million savings for Nationwide IT over the past three years.

www.mainframezone.com/it-management/optimized-computing-nationwides-linux-on-system-z-virtualization-reduces-tco-and-time-to-deploy

© 2013 IBM Corporation

Dundee City Council

Dundee City Council delivers value through new technologies

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

Overview

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

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

Dundee is Scotland's fourth largest city, home to 145,000 people. A former industrial centre, Dundee has transformed itself into a UK centre for life sciences and digital media. As a result, the city has been named one of the world's top seven intelligent communities for three of the past four years (see www.intelligentcommunity.org).

Dundee City Council employs around 10,000 people, and provides a wide range of municipal services for citizens, many of which rely on IT support. The council runs numerous applications to support both internal processes and public-facing systems, such as its Web portal (www.dundeecity.gov.uk), which provides information and online services.

Linux on System z

For several years, the council has run all its core IT systems (mostly Oracle databases and applications) on SUSE Linux Enterprise Server, running on IBM System z servers.

"Running Linux on the System z platform is a cost-efficient approach, especially for software like Oracle, which is licensed on a per-processor basis," explains Tim Simpson, IT Support Manager at Dundee City Council. "We can run 60 virtual machines on just four System z processors – whereas an equivalent x86-based architecture might require several processors for each server! So the savings can be considerable."

Leasing leading-edge hardware

The council's existing servers – a pair of z9 Business Class machines – were leased from IBM, and the existing lease was due to expire.

"The best thing about our leasing strategy is that it allows us to continually upgrade to the latest, fastest IBM hardware, while maintaining our costs at a steady level," says Simpson. "When our latest lease was coming up for renewal, we realised it was a good opportunity to rethink our storage architecture too."

At the time, the council's storage infrastructure was based on a mixture of SGI and IBM storage arrays, virtualised using a solution from LSI.

For several years, the council has run all its core IT systems (mostly Oracle databases and applications) on SUSE Linux Enterprise Server, running on IBM System z servers.

"Running Linux on the System z platform is a cost-efficient approach, especially for software like Oracle, which is licensed on a per-processor basis," *explains Tim Simpson, IT Support Manager at Dundee City Council.*

The z10 BC machines each contain two IFL processors, and run approximately 60 virtual Linux servers in total.

"We can run 60 virtual machines on just four System z processors – whereas an equivalent x86-based architecture might require several processors for each server!
So the savings can be considerable."

Dundee City Council
www.dundeecity.gov.uk

© 2013 IBM Corporation

Sparda Datenverarbeitung eG chooses IBM zEnterprise

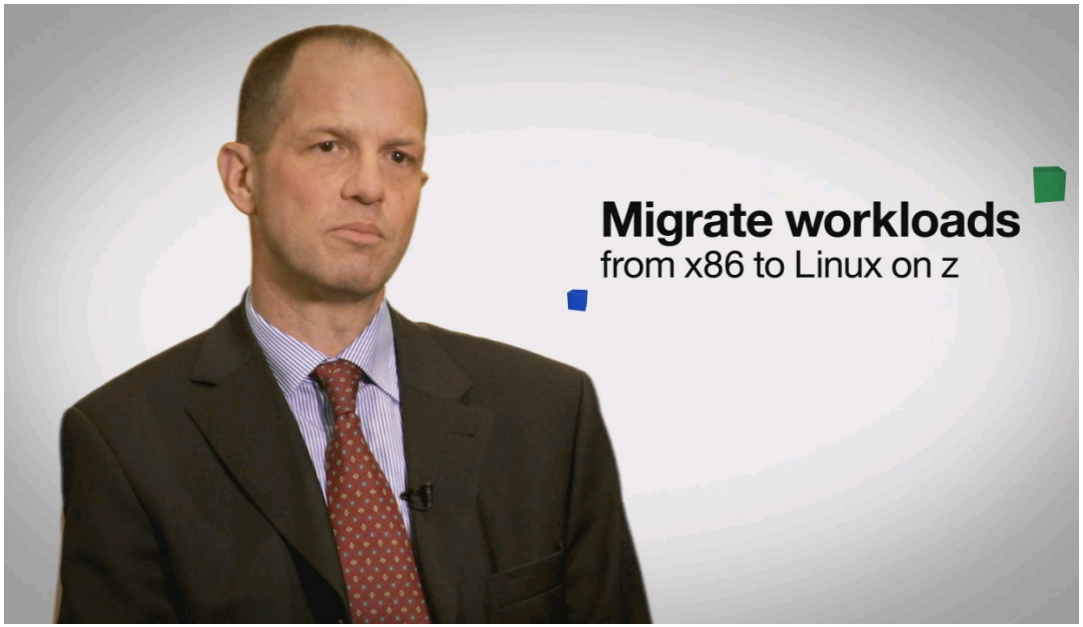

"Oracle has been consolidated on this platform we are using right now only Oracle on the z196 platform,"

Bernd Bohne, Sparda-Datenverarbeitung e.G., Manager, Central Systems

"Over the years, the mainframe transformed from traditional workloads, quite simple, to a universal platform for new workloads as well.

And we see a lot of new applications that are coming to this platform.

Especially for Linux, it's perfect.

The z/Enterprise platform is perfect for consolidating Linux workloads because of the high I/O bandwidth, business continuity with capacity backup features."

Watch and listen to

- Bernd Bohne, Sparda-Datenverarbeitung e.G., Manager, Central Systems
- Marie Wieck, IBM, General Manager, Application Integration Middleware
- Steve Mills, IBM, Senior Vice President & Group Executive, Software & Systems

ibm.com/systems/z/resources/sparda_bank_video.html

© 2013 IBM Corporation

Reference Customers
Linux on System z with Oracle

Bank of New Zealand

© 2013 IBM Corporation

Light Oracle on Linux on z environment for Demo's and POC

Light Oracle on zLinux offerings for System z are predefined remote system environments for Linux & Oracle on System z, with supporting middleware, that are "automatically" provisioned and made available for small PoC's or Functional Benchmarks.

The objective is to provide small, low cost, low touch, short term, repeatable system environments to IBM FTSS teams and Business Partners to support STG sales & growth opportunities, such as new workloads, new clients, or new ISVs.

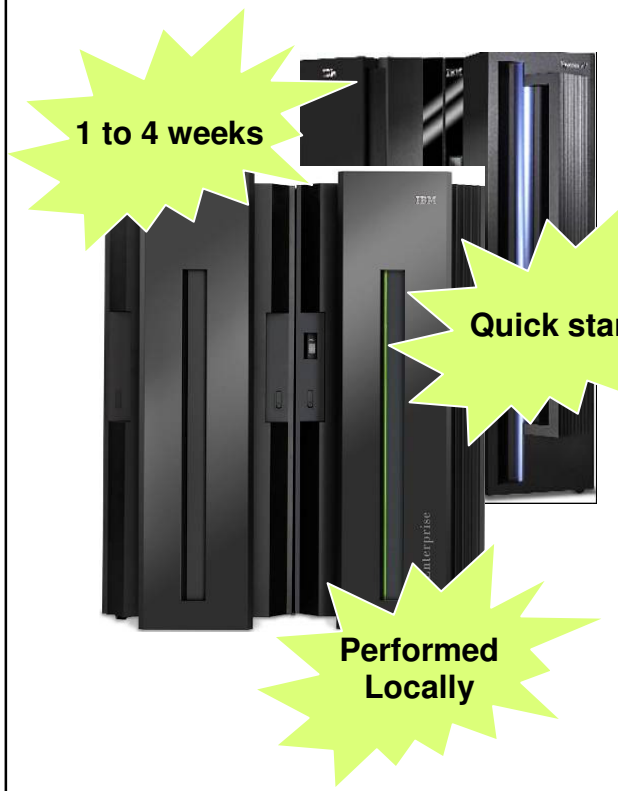

zLight Benchmark Offering:

- 2 pre-selected environments (fast-start, last certified &updated levels) are available now (up to 4 weeks).
  - Linux on System z; Linux on System z guest hosted on a z/VM LPAR; up to 4 CP shared, 4GB memory, 14GB storage; software code (WebSphere, Oracle, DB2…) downloadable from a NFS. SUSE (SLES 11 SP2) and Redhat (RHAT 6.0) versions of Linux are available.
  - Oracle solution on System z; Oracle 11g (Single Instance on SLES 11 SP2 or REDHAT 5.7) on one Linux guest + Oracle RAC 11g on 2 Linux guests (cluster, on SLES 11 SP2 or REDHAT 5.7), 4GB and 4CP per guest, up to 300GB storage.
- Support :
  - First level support provided by local IBM FTSS or BP
  - Second level support provided by the STG Mainframe Benchmark Centers in Montpellier
  - « On-demand » specific support (to be defined)
- POC or functionnal tests performed remotely by the customer/ISV and IBM team
- Infrastructure & IBM products 2nd level support assumed by STG Benchmark Center skilled people for no / small costs

KEY ADVANTAGES

- **Pre-generated and Up-to-date environments**
  - Help to answer to short duration RFP
  - Practice, education to improve knowledge and to move faster than competition
- **Prove customer solution on STG platforms**
  - Validate technical alternatives and produce inputs for future innovation plan for Customer
  - Support small business opportunities as well as strategic ones.
- **Validation of customer applications and larger benchmark needs** (unit/acceptance tests, build DB volumes and needed growth and workloads, perform small performance tests, …)
- **Easy access to IBM Benchmark Centers infrastructure & support** for customers & local team with limited travel capabilities.

© 2013 IBM Corporation

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)

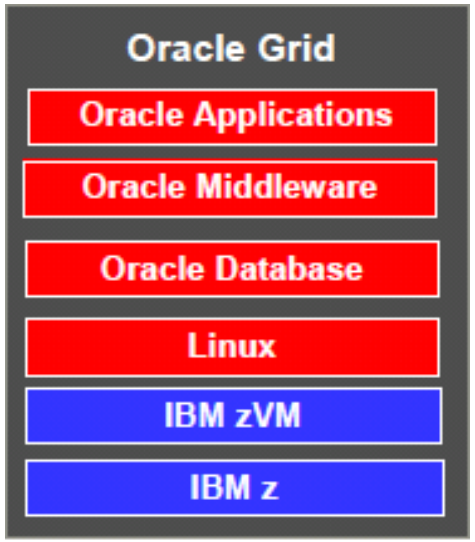

http://www.redbooks.ibm.com/abstracts/sg247191.html

© 2013 IBM Corporation

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

© 2013 IBM Corporation

धन्यवाद

Hindi

多謝

Traditional Chinese

ขอบคุณ

Thai

Спасибо

Russian

Gracias

Spanish

Danke

German

شكرا

Arabic

Merci

French

Obrigado

Brazilian Portuguese

Bedankt

Nederlands

多谢

Simplified Chinese

Thank You

English

நன்றி

Tamil

ありがとうございました

Japanese

감사합니다

© 2013 IBM Corporation