Oracle consolidation on Linux on System z customer experience

Wilhelm Mild
IT Architect
IBM Germany

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?

IBM and Oracle Have a Long-Standing Relationship

Sustaining relationship of 150K + clients

- Oracle 25 years, PeopleSoft 23 years, JD Edwards 35 years, Siebel 13 years

Mutual executive commitment

- Dedicated, Executive-led Alliance teams, Regular Senior executive reviews

Vibrant technology relationship (Diamond Partner)

- Sustained investment in skills and resources including dedicated international competency centers

Market-leading services practice

- IBM GBS is Oracle's #1 SI partner (7,500 joint projects) with 5,000 people dedicated to Oracle

Unrivalled client support process

- Dedicated on-site resources and significant program investments ($77M on 1000+ assets)

Value of Consolidation with Linux on System z

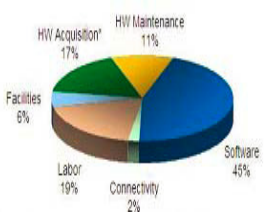

- Do more with less
  - Consolidate more servers, more networks, more applications, and more data with Linux running on z/VM
  - Achieve nearly 100% utilization of system resources nearly 100% of the time
  - Enjoy the highest levels of resource sharing, I/O bandwidth, and system availability
- Reduce 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
  - 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

Enterprise Linux Server and Oracle

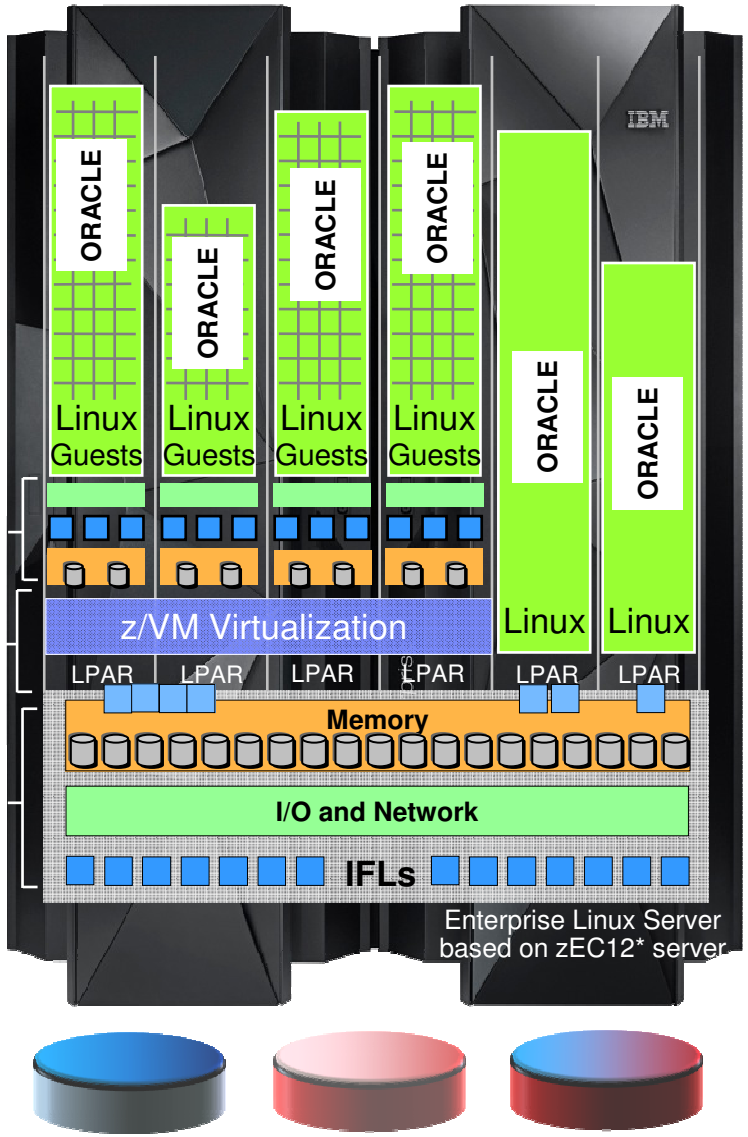

LPAR	Logical Partition = subset of hardware resources, virtualized as a separate computer; up to 60 LPARs can be configured
IFL	Integrated Facility for Linux = core; Enterprise Linux Server with zEC12* server: 5.5 GHz per core, up to 101 cores
Linux Guest	virtual Linux Guests running the workload such as database server, etc.; hundreds of virtual Linux Guests can be hosted on one Enterprise Linux Server (ELS)

* zEC12 = IBM zEnterprise EC12

Oracle 12c on Linux on System z

- Oracle on Linux on System z is the same Oracle as anywhere else - the code is ported to the new environment

- Oracle 12c (12.1) on SUSE x86-64 is the same as Oracle 12.1 on IBM Linux on System z.

- There is no difference between Linux distributions – it's the same Oracle image for any Linux on System z

- Supported is SUSE and Red Hat in exactly the same way

Oracle 12c pluggable concept

- Oracle Database 12c Release 1 introduces a new multi-tenant architecture that makes it easy to deploy and manage databases
  - –Oracle multitenant pluggable databases
  - –For consolidating multiple databases
  - –Automatic Data Optimization with Heat Map
- Virtualization based on container concept
  - –Common ASM
  - –Common SGA for pluggable databases
- A pluggable database may be moved from one container to another simply by unplugging it from one and plugging it into another
- Software patches and upgrades may be applied at the container level. Upgrading the one container database also upgrades all pluggable databases at once.

Why using High-End Servers ?

Utilization on x86 systems

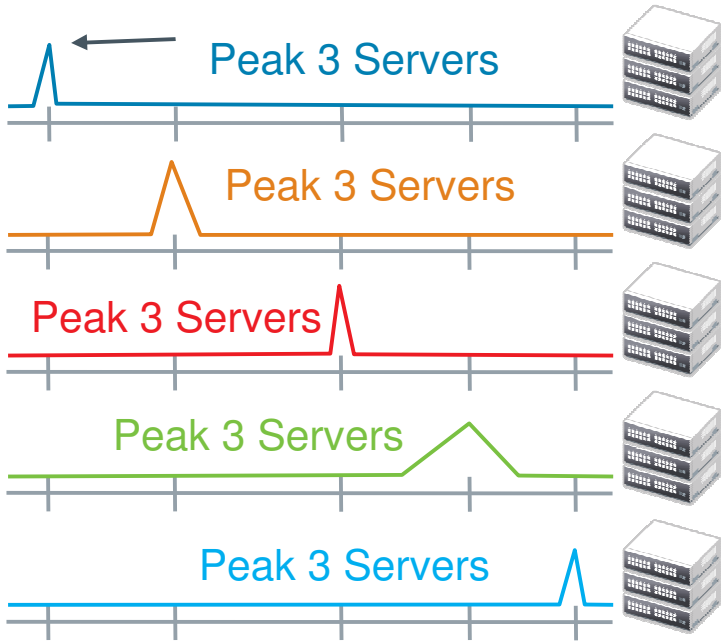

According to a study by Gartner, data centers that do not use virtualization have an average server CPU utilization rate of only 15%.

Mixed Utilization on IBM High End Servers

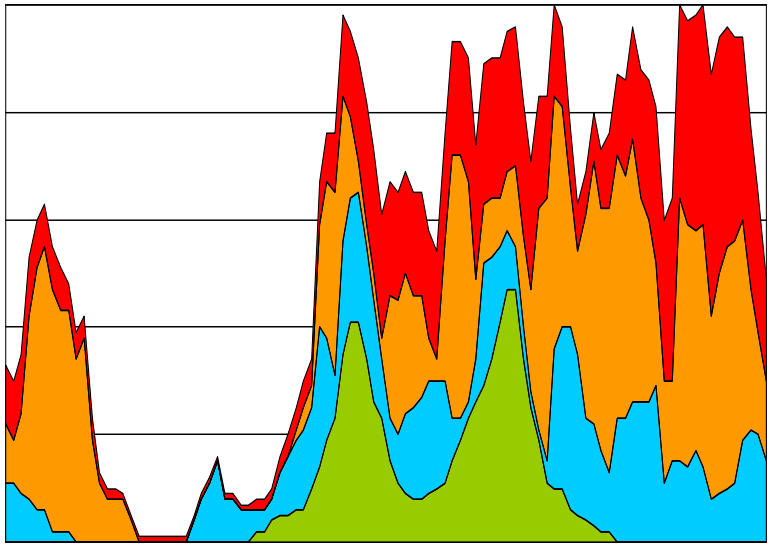

IBM High End Server: Up to 100% utilization

- 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 License	Software Update License & Support
Prices in USA (Dollar)		
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

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 Oracle on z projects

- **1 - Scope of the project delimitation – feasibility study**
  - Gather information on existing environment (servers, applications, network)
  - Fit fot 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

Oracle Solutions delivered …

WebLogic Server 12c Release 1 with Coherence Version 12.1.2 Q413

Siebel CRM Version 8.1.1.10 accessing Oracle DB 11g Release 2 10gR2 Q213

New Enterprise Manager Agent 12c Release 2 Version 12.1.0.2 Q212

Oracle DB 11g Release 2 Version 11.2.0.2.5 Q112

Oracle DB 12C Release 1 Version 12.1.0.1.0 Q114

Siebel Public Sector, Loyalty Management Version 8.2.2.3 accessing Oracle DB 11g Release 2 10gR2 Q213

New WebLogic Server 12c Release 1 Version 12.1.1x Q212

Oracle DB 11g Release 2 Version 11.2.0.3.1 Q112

WebLogic Server 11gR1 PS3 Q111

New Oracle DB 11g Release 2 Version 11.2.0.4 Q114

SOA WebCenter 11gR1 PS6 WebLogic Server 11gR1 11.1.1.6 Q213

Oracle DB 11g Release 2 Version 11.2.0.3.3 Q312

E-Business Suite Release 12.1x accessing Oracle DB 11g Release 2 Version 11.2.0.3 Q112

Oracle DB 11g Release 2 Version 11.2.0.3 PSU 3 Q311

New E-Business Suite Release 12.2 accessing Oracle DB 11g Release 2 Version 11.2.0.3 Q114

Oracle BIEE Version 11.1.1.7 with Oracle DB as DSR* 10g Release 2 11g Release 2 Q213

Oracle DB 11g Release 2 Version 11.2.0.3.4 Q412

WebLogic Server, SOA, WebCenter 11gR1 PS5 Q112

PeopleTools 8.51 accessing Oracle DB 10g Release 2 10.2.0.5 Q311

Oracle DB 11g Release 2 Version 11.2.0.3.6 Q213

Oracle DB 10g Release 2 Version 10.2.0.5.10 Q113

Oracle DB 11g Release 2 Version 11.2.0.3.2 Q212

Oracle DB 11g Release 2 Version 11.2.0.2 PSU 4 Q411

Oracle DB 11g Release 2 Version 11.2.0.3.7 Q313

Oracle DB 11g Release 2 Version 11.2.0.3.5 Q113

PeopleTools 8.51 accessing Oracle DB 11g Release 2 11.2.0.3 Q212

PeopleTools 8.52 accessing Oracle DB 10g Release 2 10.2.0.5 Q411

Oracle DB 11g Release 2 Version 11.2.0.3.8 Q413

Oracle DB 11g Release 2 Version 11.2.0.3 RHEL 6.2 Support Q113

PeopleTools 8.52 accessing Oracle DB 11g Release 2 11.2.0.3.1 Q212

PeopleTools 8.52 accessing Oracle DB 11g Release 2 11.2.0.2 Q411

New Oracle Billing and Revenue Management Version 7.5.0.0.0 accessing Oracle DB 11g Release 2 Version 11.2.0.3 Q113

Enterprise Manager Agent Release 12c Version 12.1.0.1 Q212

Hyperion EPM Version 11.1.2.3 accessing Oracle DB 10g Release 2 11g Release 2 Q213

New Hyperion EPM Version 11.1.2.2 accessing Oracle DB 10g Release 2 11g Release 2 Q212

Oracle DB 10g Release 2 Version 10.2.0.5 PSU 5 Q411

Oracle DB 11g Release 2 Version 11.2.0.3 Q411

PeopleSoft 9.2 PeopleTools 8.53 accessing Oracle DB 11g Release 2 11.2.0.3 10gR2 10.2.0.5 Q213

New PeopleSoft Campus Solutions V9 PT 8.52 accessing Oracle DB 11g Release 2 Q411

* Data Source or Repository

… for Linux on IBM zEnterprise

Oracle DB certifications on Linux on System z

- Oracle has been delivering database solutions on Linux on System z Servers since 2002
- On January 9th, 2014 delivered Oracle 12c Release 1 (12.1.0.1) Database for Linux on IBM System z Servers.

	SLES 10	SLES 11	RHEL 4	RHEL 5	RHEL 6
Oracle DB 12.1.0.1	NO	**YES**	NO	**YES**	**YES**
Oracle DB 11.2.0.3	**YES**	**YES**	**YES**	**YES**	**YES**
Oracle DB 11.2.0.2	**YES**	**YES**	**YES**	**YES**	NO
Oracle DB 10.2.0.5	**YES**	**YES**	**YES**	**YES**	NO

- **E-Business Suite on Oracle Database 11*g* Release 2** is supported as a mixed mode architecture (formerly " split tier architecture"). Database can run on Linux on z. EBS code runs on a different platform (AIX, Linux on x86, etc.)
- **Hyperion** Finance EPM is supported on Linux on System z in a split tier architecture.

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?

Workload is dependend on all layers

- Storage
  - type, ECKD, FCP
  - attachment perfomance
- Virtualization
  - Type, LPAR, z/VM, Oracle
- Oracle 12c database implementation
  - single
  - container
  - HA / clustered

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
  - 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

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 estimation
  - Memory required for Oracle PGA: As per DBA estimation
  - 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

Starting size = SGA + PGA + 0.5GB for Linux + ASM (if used)

- Memory over-commitment (relationship of virtual to real memory)
  - Limit/avoid memory over-commitment for critical production databases
  - Test/development guests can benefit from z/VM memory over-commitment capability

Linux Huge pages - recommendation

- With HugePages each OS pagetable mapping (virtual to physical) maps memory points to a 1MB page (as opposed to 4KB)
- Decreases page table overhead
- Pages are locked in memory and never swapped out, which provides RAM for shared memory structures such as SGA
- For Huge Pages you must use ASMM (Automatic Shared Memory Management)
- Huge Pages incompatible with AMM (Automatic Memory Management)
- Less operating System overhead, fewer cpu resources
- Recommended when SGA > 8GB

Linux Huge pages - recommendation

- If huge pages are configured, this amount of memory is no longer available for applications using 4K pages
  - Starting with Oracle 11g, the use of huge pages is done automatically
    - If the SGA can not be allocated as a whole in huge pages, the fall back is to allocate the whole SGA in 4KB pages, which can produce a heavy memory pressure.
  - Ensure to have enough huge pages defined that the full SGA from **all** Oracle databases in that system server fits into

- To verify usage of Hugepages
  - Monitor value of *HugePages_Free*
    - When starting Oracle the amount value of HugePages_Free must be lower (reduced by the SGA size)

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 #	# OEM Servers 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						

Processor	Feature	Utilization for Case 1 < Complementary Peaks Concurrent > 0%	40.0%	70.0%	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%

Summary of Servers to be Consolidated

Servers	Chips	Cores	Applications
36	35	172	36

Oracle Consolidation on System z study : methodology example

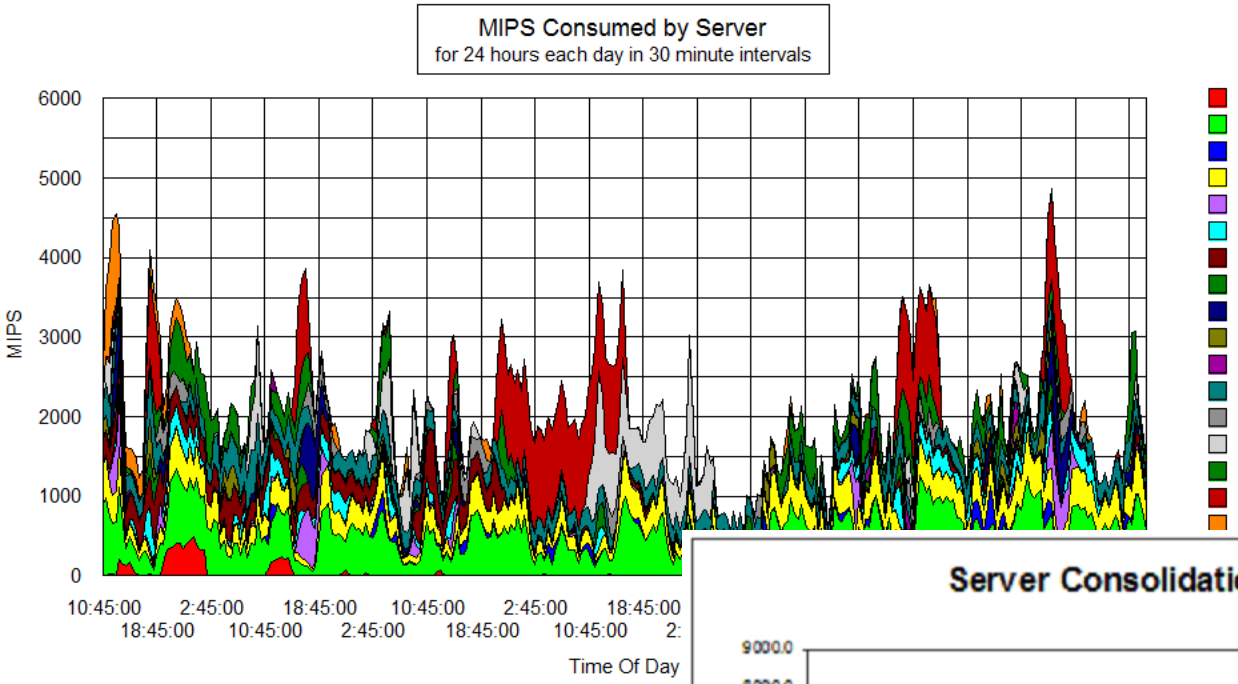

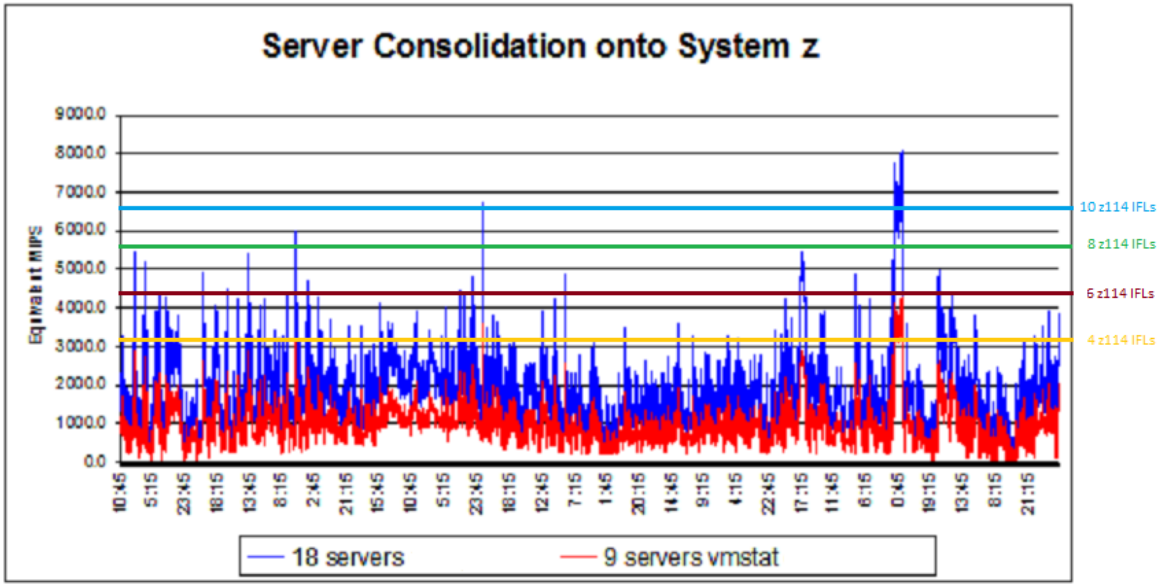

Oracle Consolidation on System z : provide accurate values otherwise…

Model	Sockets	Cores per Socket	Processor Speed
Intel(R) Xeon(R) CPU X5650	2	6	2.67GHz

- Before consolidation.
- Workload on Intel Xeon registered on July 2012

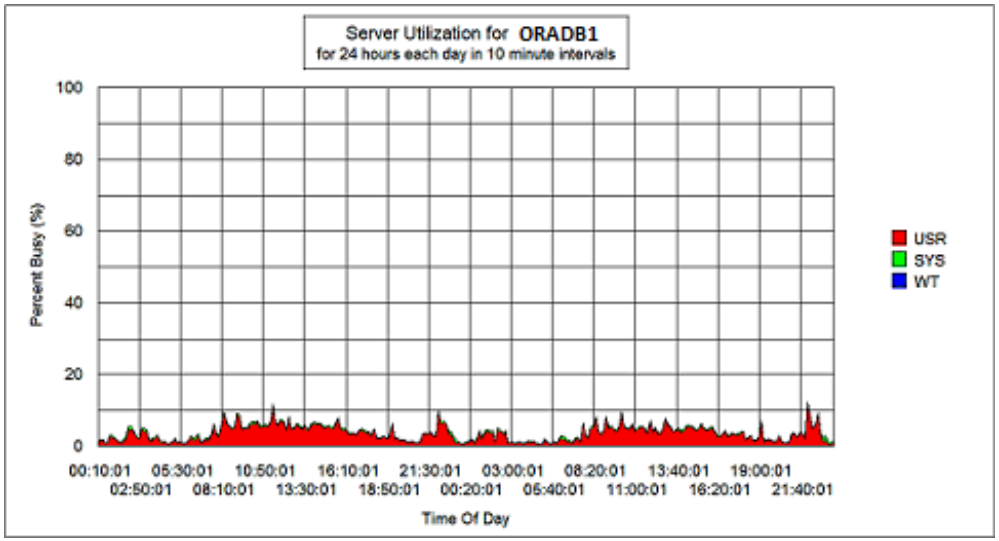

- After consolidation on Linux System z on Sept. 2013

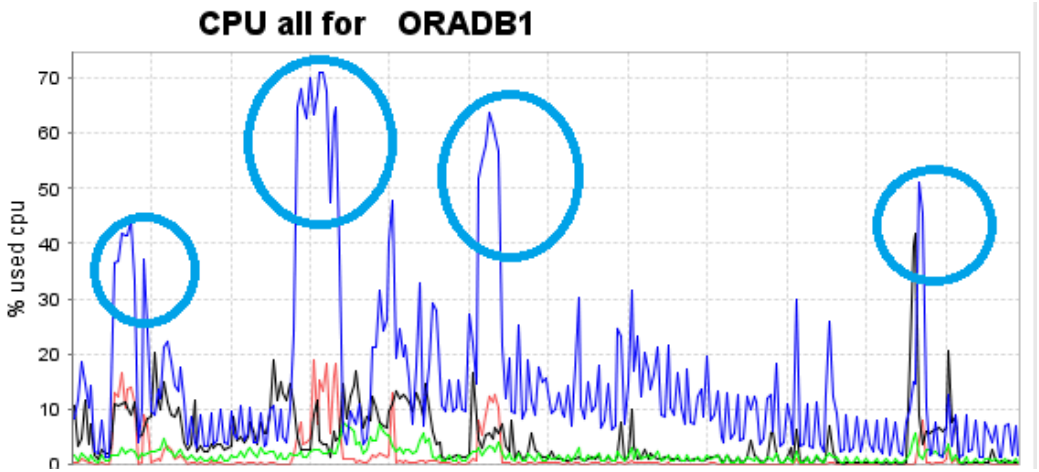

- => the workload is not the same

Need maximum I/O performance? Use FlashSystem with Linux on z!

Now certified to attach to Linux on System z, with or without an SVC, to meet your business objectives

Linux on System z can help achieve a smarter IT infrastructure that:

- Provides efficiency at scale on a single physical server
- Delivers industry-leading virtualization for effective deployment
- Enables flexible delivery of services through a private cloud
- Delivers real-time information and insight from data
- Provides unmatched security and reliability

And now you can leverage the "Economies of Scale" of Flash

- **Accelerate Application Performance**
- **Gain Greater System Utilization**
- **Lower Software & Hardware Cost**
- **Save Power / Cooling / Floor Space**
- **Drive Value Out of Big Data**

Linux on System z – Disk storage connectivity options

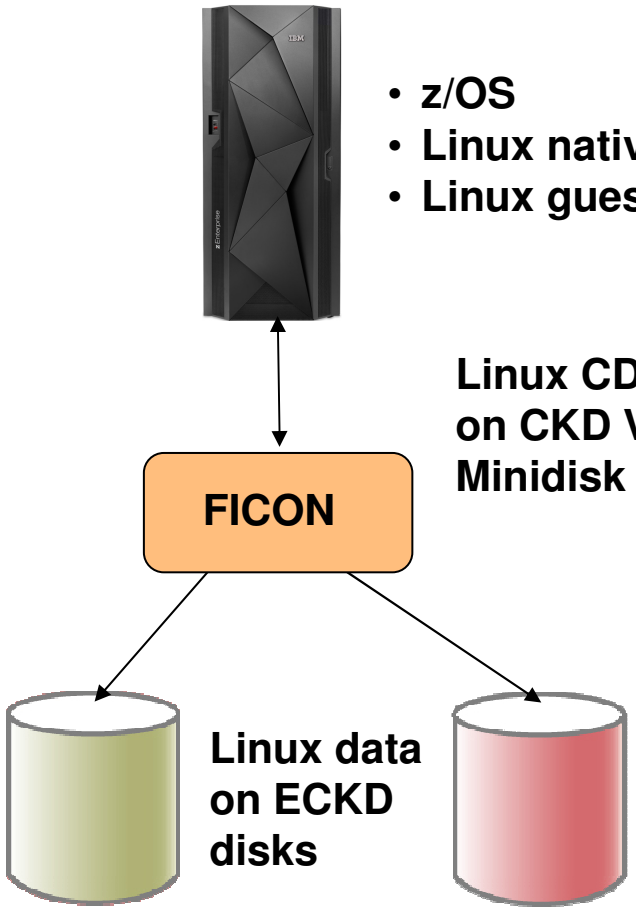

- z/OS
- Linux native
- Linux guest

Linux CDL on CKD VM Minidisk *

FICON

Linux data on ECKD disks

DS8000 only

- Linux native
- Linux guest

FCP

Linux on FBA VM Minidisk (EDEV)

Linux data on FBA disk

DS8000, XIV, Storwise, FlashSystem

Addressing I/O pressures together

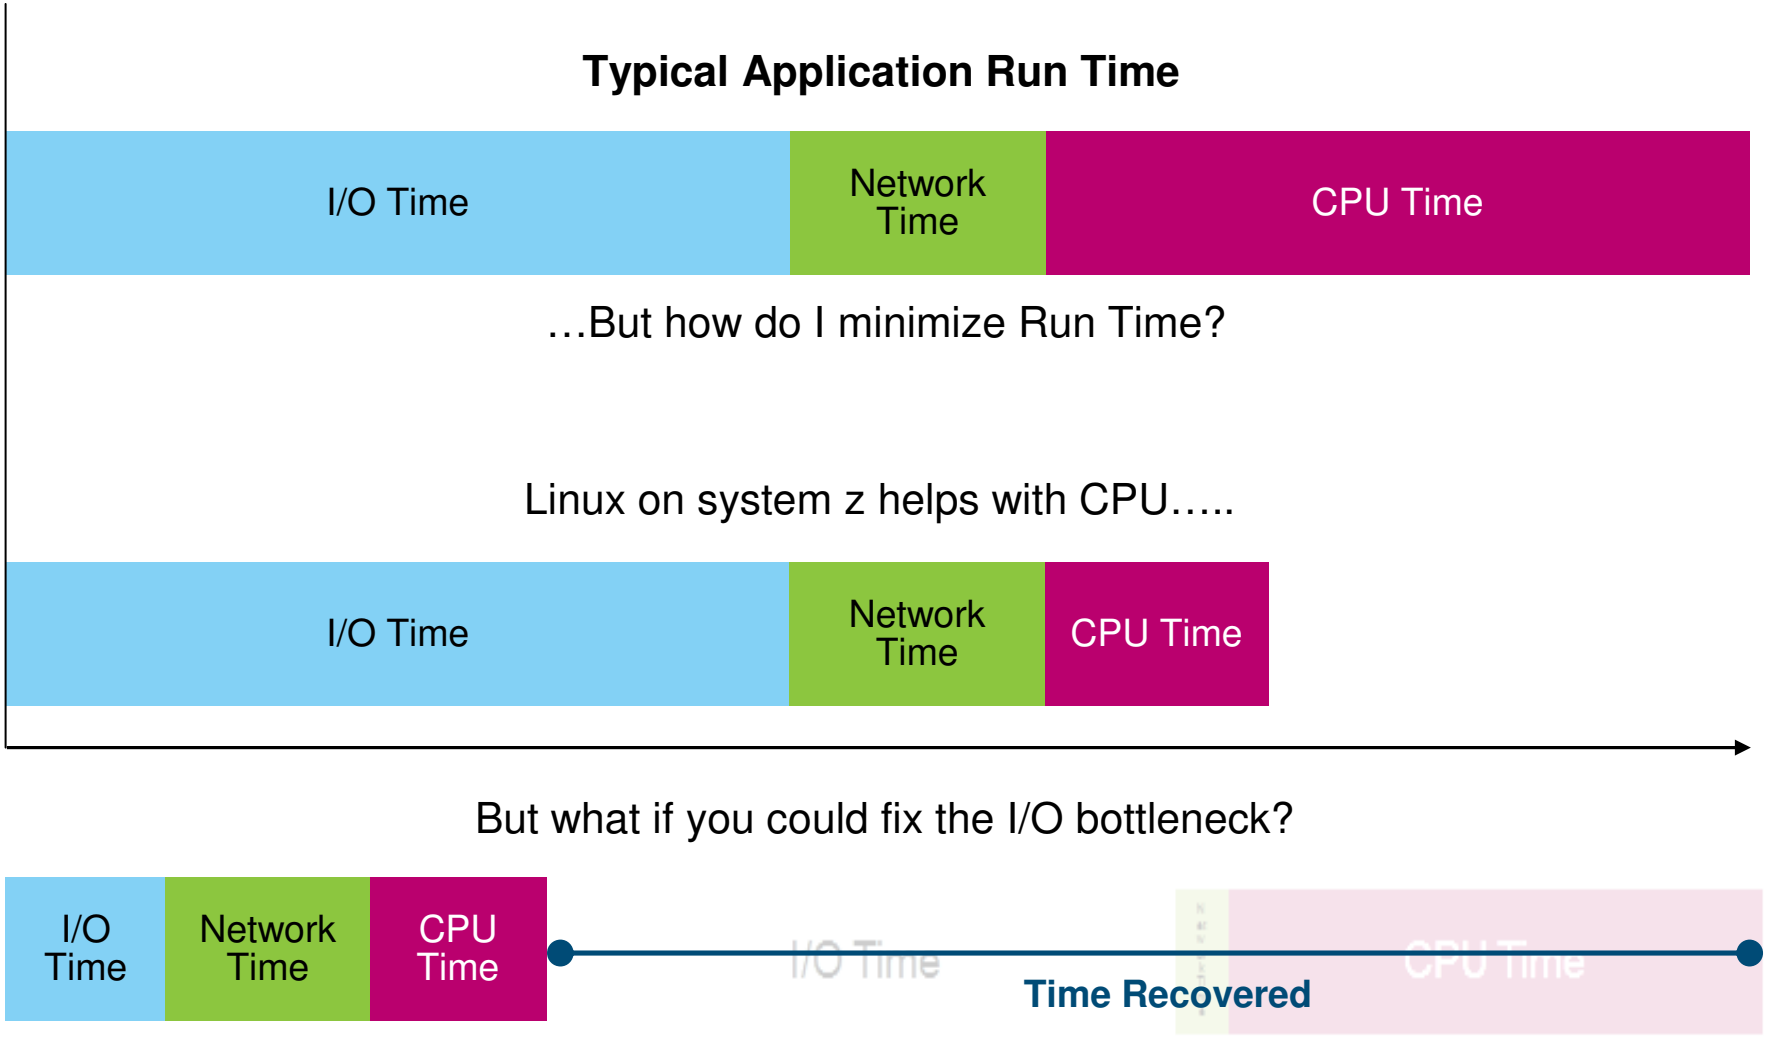

Easy Tiering => Performance made easy

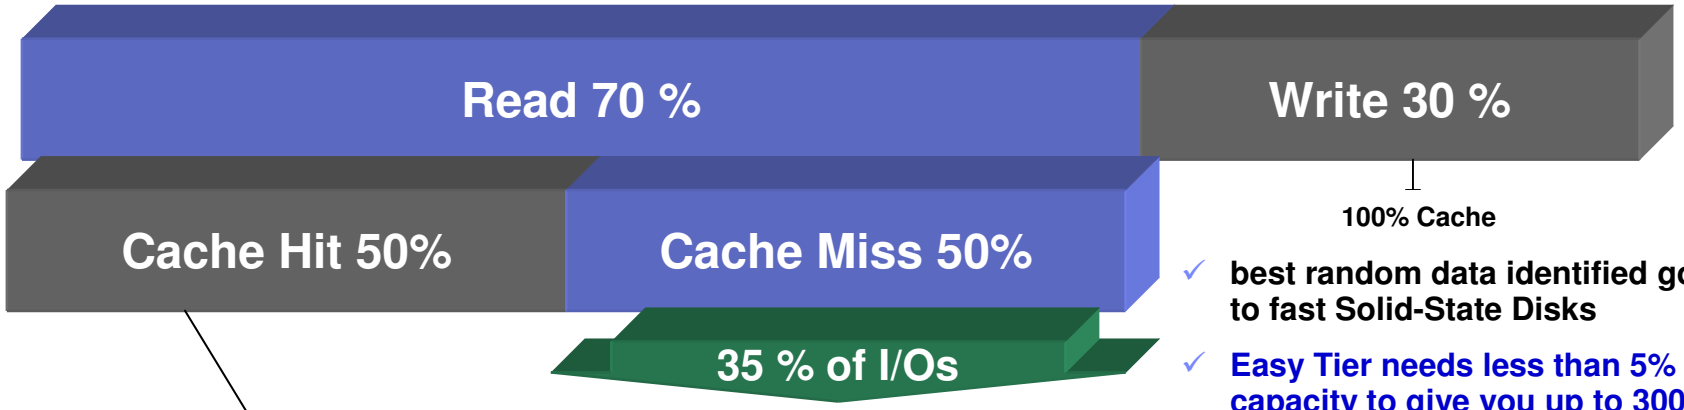

100% Cache

- ✓ best random data identified goes to fast Solid-State Disks
- ✓ Easy Tier needs less than 5% capacity to give you up to 300% performance

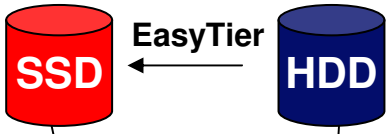

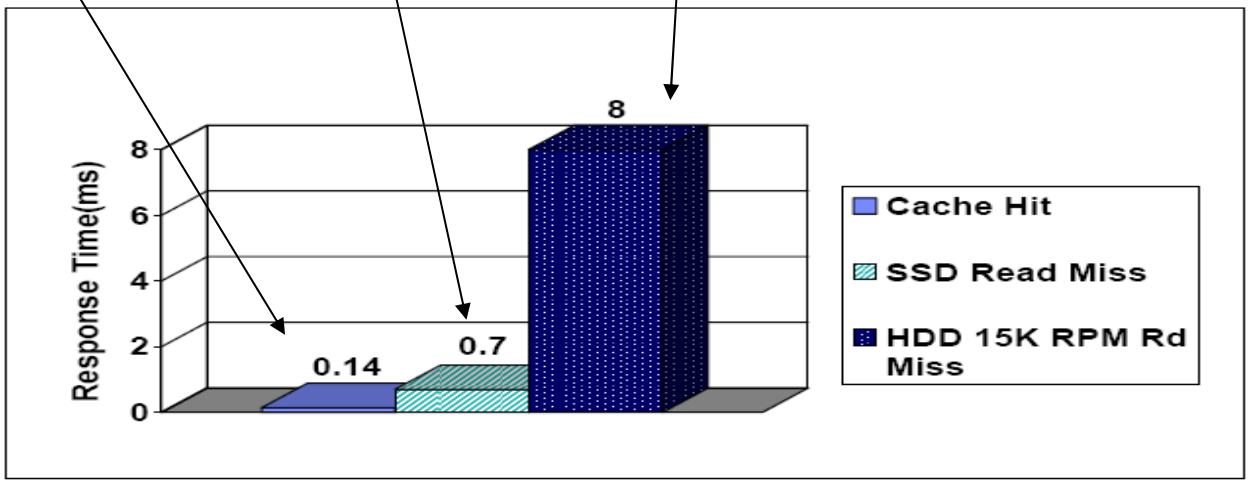

EasyTier for
- **DS8000**
- **Storwize V7000**
- **SVC**

Distributed Systems 4KB Read Response Time

DS8300 / DA FC3041

2010-05-11

IBM Flash in Storage

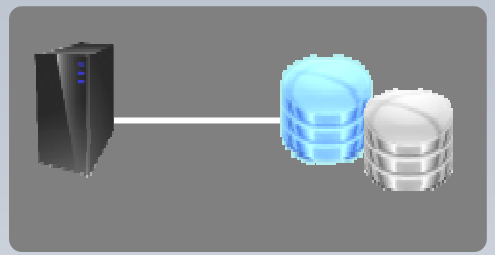

Flash in Storage

FlashSystem 820 & 840
FlashSystem Solutions

All Flash Array
Fibre Channel, InfiniBand, FCoE
4TB to 48TB per 3U
IBM MicroLatency™

DS8870

Enterprise All Flash or Hybrid
Fibre Channel, FICON
Easy Tier

XIV

Cloud-optimized scale-out
Fibre Channel, iSCSI
Up to 12TB Flash cache

Storwize

All Flash or Hybrid
Software defined storage

Business Critical Data Economics on *DS8870 Flash*
Optimized - delivers immediate ROI…

Today

- ✓ ***Superior Data Economics***
- ✓ ***Leading Performance*** *Across Varying & Dynamic Workloads*
- ✓ ***Proven Security & High Availability*** *Resiliency Architecture & Multi-site Disaster Recovery*

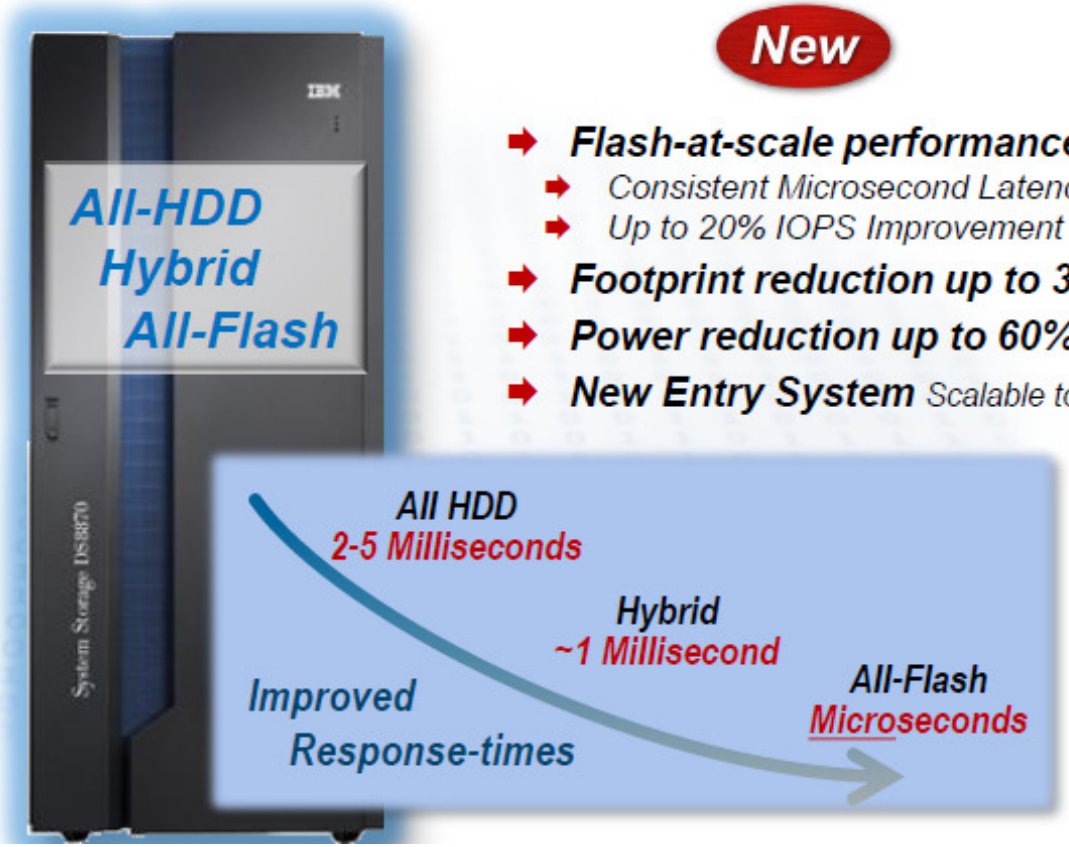

New

- ➡ ***Flash-at-scale performance***
  - ➡ *Consistent Microsecond Latency*
  - ➡ *Up to 20% IOPS Improvement*
- ➡ ***Footprint reduction up to 30%*** *vs. HDD*
- ➡ ***Power reduction up to 60%*** *vs. HDD*
- ➡ ***New Entry System*** *Scalable to 1,056 drives*

Microlatency effects storage run time

I/O Serviced by Disk

1. Issue I/O request ~ 100 µs
2. Wait for I/O to be serviced ~ 5,000 µs
3. Process I/O ~ 100 µs

- Time to process 1 I/O request = 200 µs + 5,000 µs = 5,200 µs
- CPU Utilization = Wait time / Processing time = 200 / 5,200 = ~4%

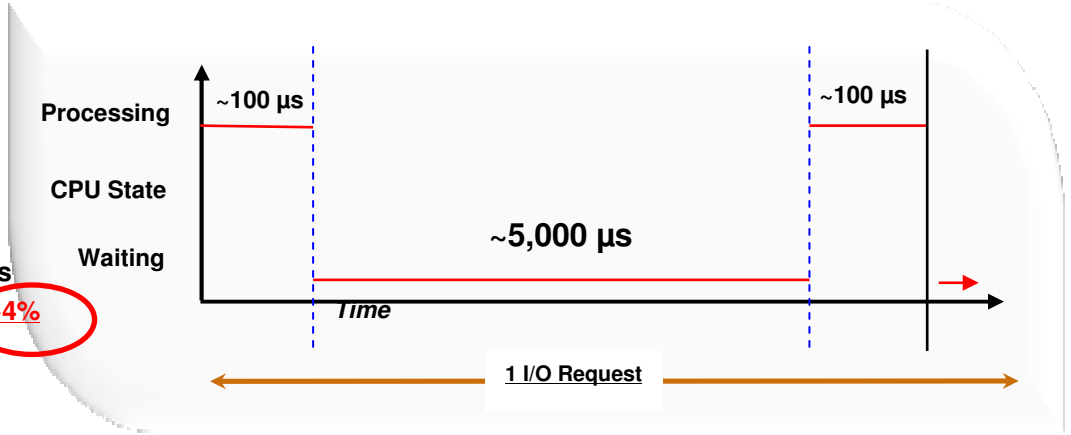

I/O Serviced by IBM FlashSystem

1. Issue I/O request ~ 100 µs
2. Wait for I/O to be serviced ~ 200 µs
3. Process I/O ~ 100 µs

- Time to process 1 I/O request = 200 µs + 200 µs = 400 µs
- CPU Utilization = Wait time / Processing time = 200 / 400 = 50%

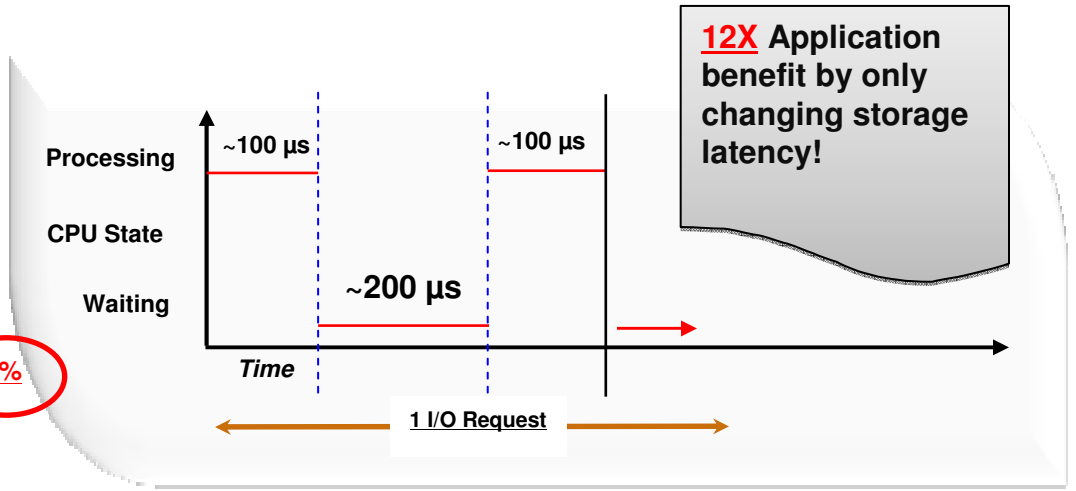

FlashSystem 820 Oracle performance results with Linux on System z

Performance of Linux on System z with FlashSystem and Oracle:

I/O bound Oracle databases can benefit from IBM FlashSystem over spinning disks.

- **21x** reduction in response times
- **272%** improvement in CPU utilization
- **957%** improvement in IO wait times

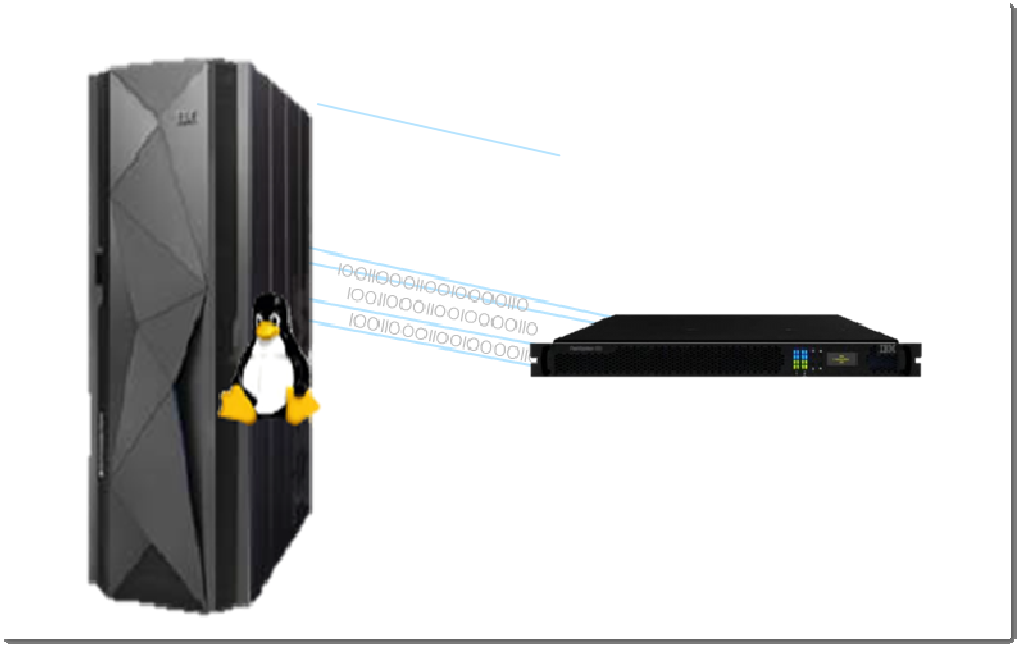

*System z FiconExpress **8s** I/O cards can provide an additional 10% throughput running with FCP*

Now certified to attach to System z for Linux on System z, with or without an SVC

Source: IBM Actual Results and Estimates

Summary disk connections for Oracle

- Storage server:
  - Storage pool striping with many ranks
  - Use disks from both internal servers
- Infrastructure:
  - Use the highest possible link speed of the FICON/FCP channel, but ensure it is supported from all elements in the path
- FICON Disks
  - Use an appropriate amount of HyperPAV aliases (e.g. 10 per LCU), the larger the more are needed
  - FICON devices are easy to administer, they save CPU due to the use of the SAPs
- FCP Disks
  - Use multipath policy multibus and rr_min_io_rq/rr_min_io in the area of 100
  - FCP devices are driven by the CPU/IFLs, but this allows higher throughput values
- Further tuning options: FCP queue depth (default 32)
  - Indicator : iostat/sar reports larger values for avgqu-sz (for example > 10)

Evaluate your Storage

- Test your storage subsystem, (without Oracle database installed !) with disk utility **Orion** provided by Oracle to help decide various storage configurations for your Oracle Databases.
  http://www.oracle.com/technetwork/topics/index-089595.html

- You can also use the Oracle I/O calibrate routine from the script provided from the Oracle Database PL/SQL Packages and Types Reference 12c Guide:
  http://docs.oracle.com/cd/E16655_01/appdev.121/e17602/d_resmgr.htm#ARPLS67598

- Oracle's /IO calibrate routine does not harm the database or the underlying data files. I/O calibrate does require a database to be created.

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 - Verify 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
    - Partitioning
    - 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!

Effective Virtualization with Linux on z and z/VM shared memory

Linux Shared Memory Exploitation for many Virtual machines

z/VM Discontiguous Saved Segments (DCSS)

- **DCSS support is Data-in-Memory technology**
  - Share a single, real memory location among multiple virtual machines
  - Can reduce real memory utilization
- **Use Cases:**
  - **As fast Swap device**
  - **For sharing read only data**
  - **For sharing code (e.g. program executables/libraries)**
- **The large DCSS allows the installation of a full middleware stack in the DCSS (e.g. WebSphere, Databases, etc)**
- **The DCSS becomes a consistent unit of one software level**
- **NSS – Named Saved System – for a bootable Linux image**

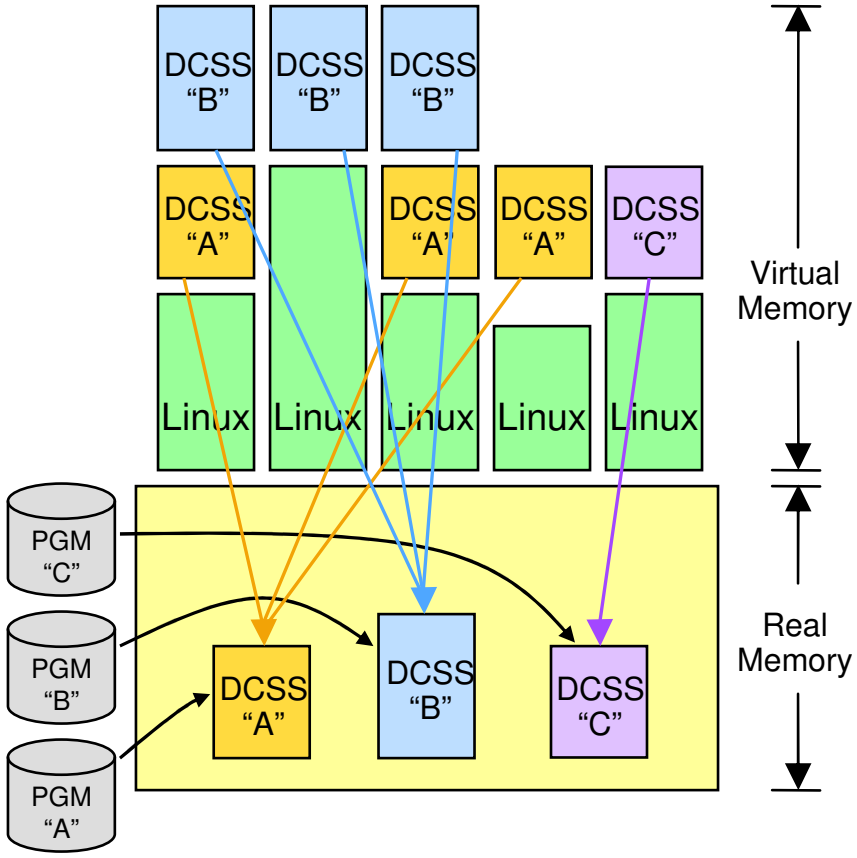

http://public.dhe.ibm.com/software/dw/linux390/perf/ZSW03186USEN.PDF

43

Effective Virtualization with Linux and z/VM SRM

- **Real memory constraint corrected by z/VM**
  Virtual Machine Resource Manager (SRM)
  - Linux images signaled to reduce virtual memory consumption
  - Demand on real memory and z/VM paging subsystem reduced
- **z/VM Virtual Disks in Storage (VDISK)**
  - Simulate a disk device using real memory
  - Use VDISKs for Linux swap devices instead of real disk volumes
  - Reduces demand on I/O subsystem
- **Linux guest: shared program executables**
  - Execute-in-place (xip2) file system
  - Access to file system is at memory speeds; executables are invoked directly out of the file system (no data movement required)
- **Data-in-Memory technology**
  - Share a single real memory location among multiple VMs
  - Reduce real memory utilization

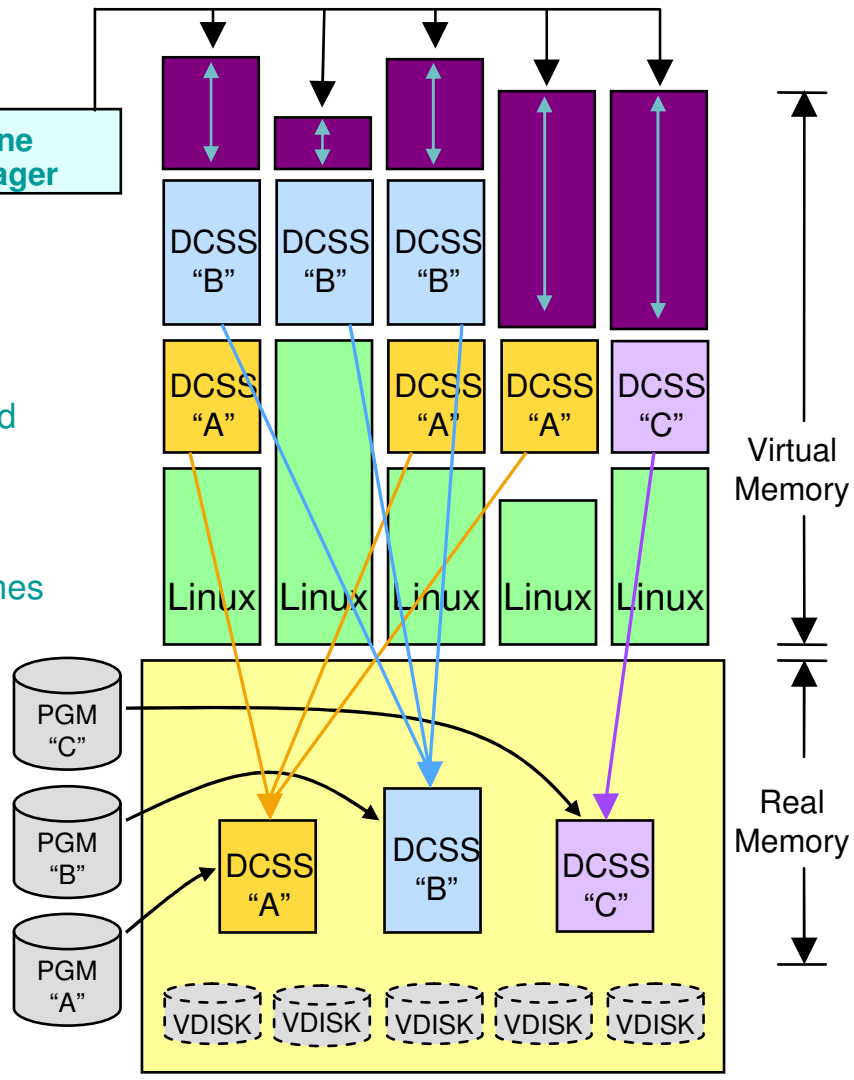

http://public.dhe.ibm.com/software/dw/linux390/perf/ZSW03186USEN.PDF

inactive virtual memory

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 (Before 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!)

Performance measurement before and after migration

To avoid the complaints such as “it was running faster before the migration to this new system” we advise to take real performance measurements on the source system before the migration, and on the target system after the migration

- at the Operating System level
  - SAR
  - NMON
  - VMSTAT
  - IOSTAT
- at the Oracle Database level
  - AWR reports
- at the application level
  - Duration of some batches
  - Response time for some complex user transactions…

These performance measurements provide the baseline prior to any migration/operation.

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 develop new Oracle servers
  - They wanted to test their own infrastructure (« background task »)

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

- Results and return of experience
  - No planning, no dedicated resources =>it took a long time to take a decision (several months)
  - Customer in production now
  - 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

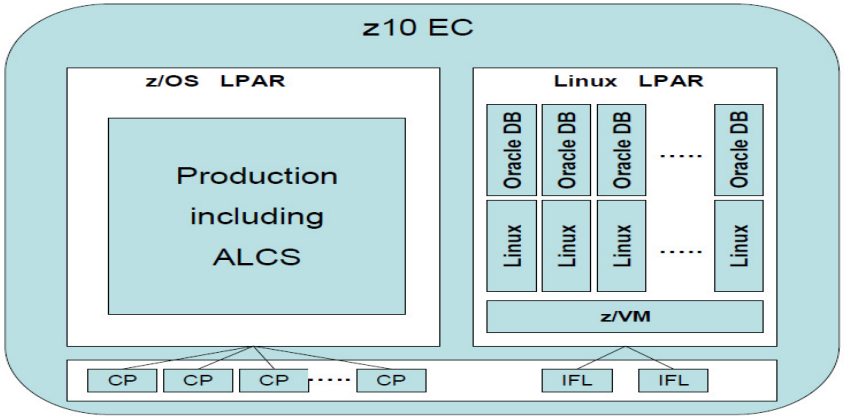

- Results
  - The Poc was done at their site with the help of BOE people (on site) and Mop people (remotely)
  - After this first step, a workshop was done at their site for HA/DR with Oracle DB on System z environment
  - Customer now in production

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 close the deal
  - Customer is in production now

Some Reference Customers
Linux on System z with Oracle

Bank of New Zealand

Business Connection

Nationwide®
Insurance

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?

• 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

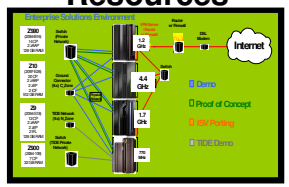

• Technical Expertise from IBM Developmentm, Lab Boeblingen

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

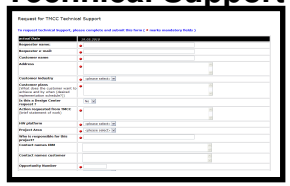

The IBM Oracle Centers

MISSION

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

STRENGH

Cross platform team with strong knowledge on Oracle products and a wide network within IBM and Oracle ecosystem

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…)

ibmoracl@us.ibm.com

IBM Oracle Competence Centers

IBM Oracle Competence Centers

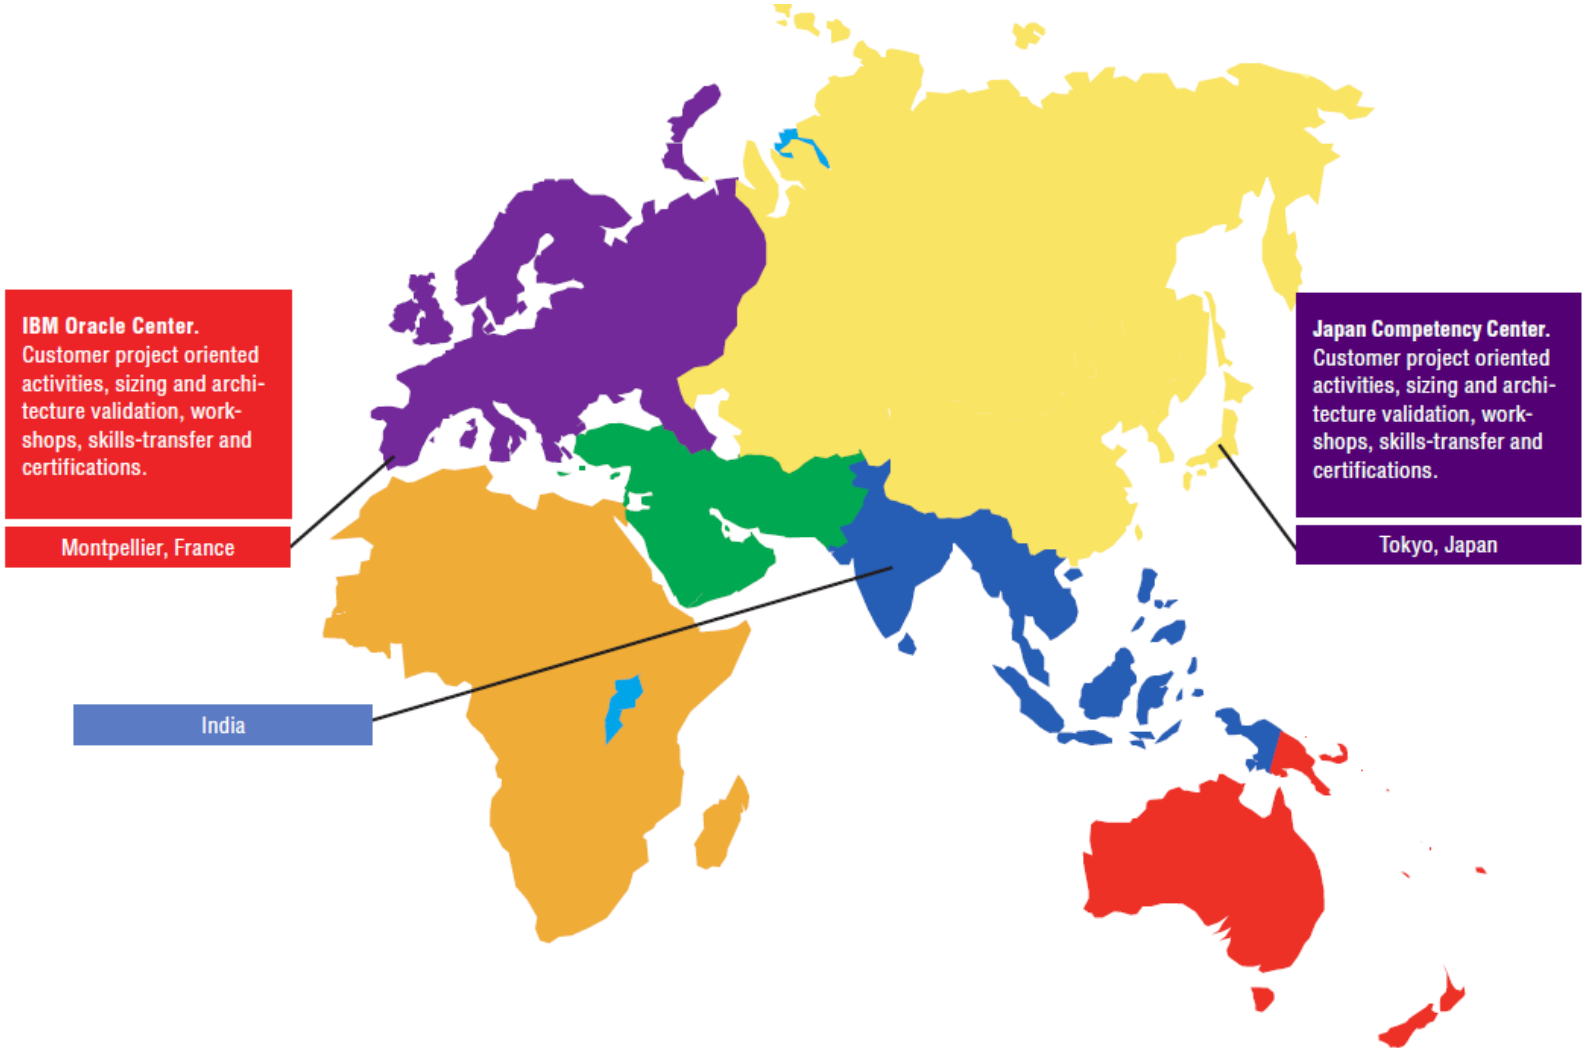

Deploy Oracle Software to the “Best Fit” Technology

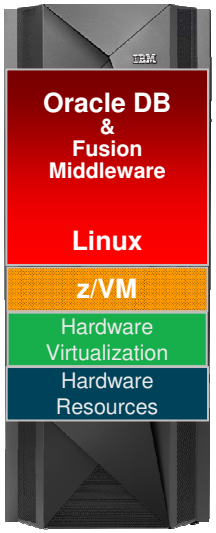

Oracle software deployments (incl. consolidations) with the Enterprise Linux Server provides an excellent price performance.

- From an Oracle licensing perspective 1 IFL = 1core
- Less operational efforts
- High levels of security and availability

Business Connexion – South Africa	**Sparda Datenverarbeitung eG** – Germany
▪ ICT services to the financial sector, government, … and more	▪ IT provider for approximately 4.2 million customers
▪ Approximately 50 virtual Linux servers; flexible environment for hosted services; high performance for Oracle databases	▪ Runs a number of very large Oracle databases, where the virtual Linux server requires 30 GB memory and ~350 GB storage
▪ Enabled competitive pricing for client services	▪ Experienced >99% availability, which proves the Linux reputation

The Met Office forecasts a bright outlook for Linux on zEnterprise

Saving software licensing and hardware lifecycle costs by consolidating applications and systems

The need

The Met Office uses post-processing systems to tailor its weather forecasts for specific clients' needs. Running these systems on a distributed Linux infrastructure was becoming complex and expensive.

The solution

Following a comprehensive evaluation and benchmarking process, the Met Office decided to migrate suitable candidates from its distributed Linux landscape onto a pair of IBM zEnterprise 196 servers.

- **Oracle licensing costs cut by a factor of 12**
- I/O-intensive workloads perform considerably better on zEnterprise than on commodity servers
- Fewer physical servers means a more manageable Linux landscape and lower hardware lifecycle costs

> *"By consolidating distributed commodity servers you can save a great deal of money. When we looked at all of the parameters, it just made sense to move the workload to the mainframe."*
>
> — Martyn Catlow, portfolio lead for centralised IT infrastructure, the Met Office

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

- Linux on System z is Oracle's platform for the mainframe
- Oracle database 12cR12 available on Linux on System z (since 1Q2014)

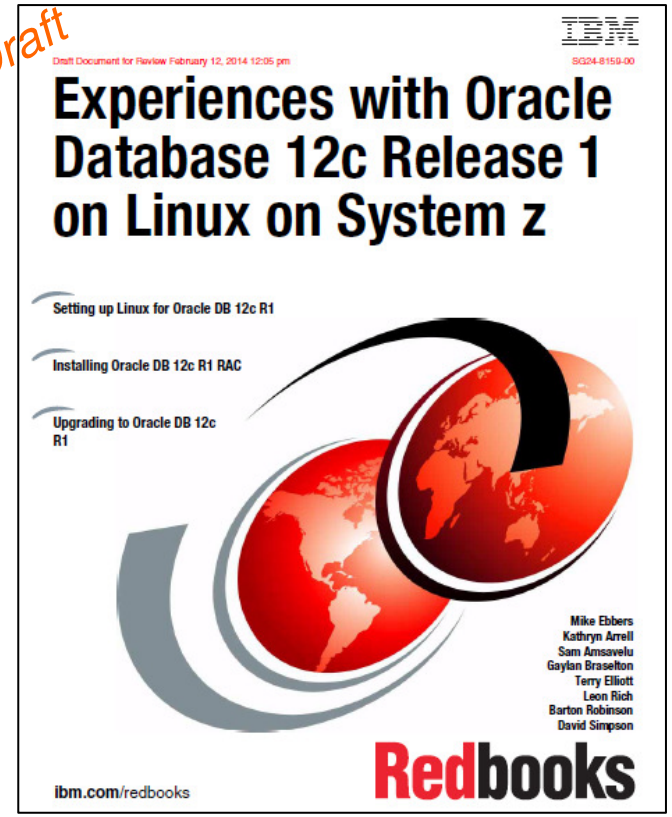

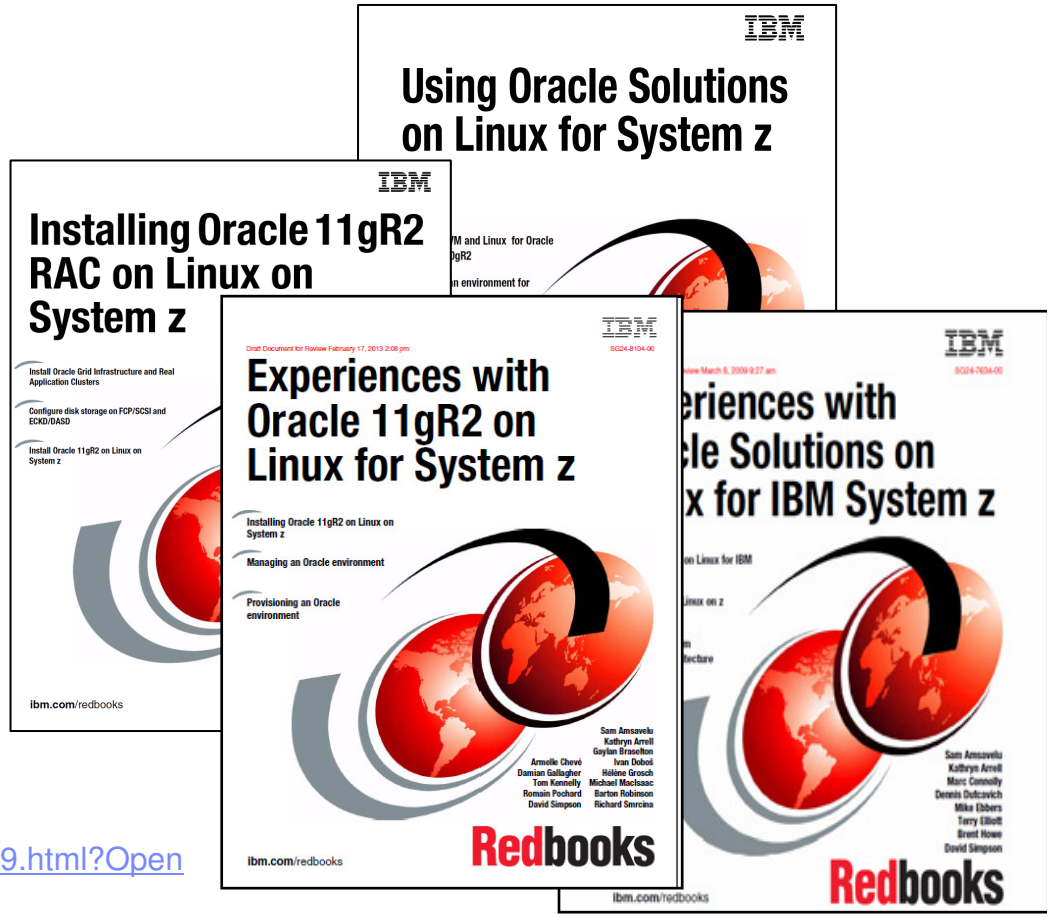

http://www.redbooks.ibm.com/redpieces/abstracts/sg248159.html?Open

Resources

- **RedBooks**
  - Experiences with Oracle 11gR2 on Linux for System z
    http://www.redbooks.ibm.com/redpieces/pdfs/sg248104.pdf
  - 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
  - Database Tuning for Linux on System z
    http://www.ibm.com/developerworks/linux/linux390/perf/tuning_database.html

Questions?

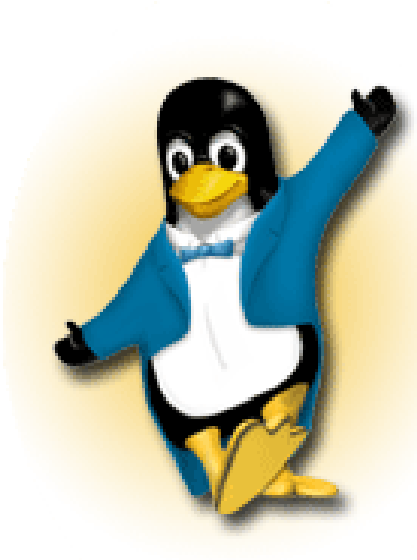

Wilhelm Mild

IBM Executive IT Architect

IBM Deutschland Research & Development GmbH
Schönaicher Strasse 220
71032 Böblingen, Germany

Office: +49 (0)7031-16-3796
wilhelm.mild@de.ibm.com

THE Open GROUP
Master
Certified IT Architect

धन्यवाद
Hindi

多謝
Traditional Chinese

ขอบคุณ
Thai

Спасибо
Russian

Gracias
Spanish

Thank You
English

شكراً
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

Trademarks

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

* 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 SpeedStep, 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.
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.

This information 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) ("SEs"). 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 SE 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.