

Siegfried Langer

Business Development Manager z/VSE and Linux on System z

IBM Germany Research & Development



ISY0456

Oracle and z Systems – A Great Combination

2015

IBM z Systems

Technical University

18-22 May | Dublin, Ireland



Discussion topics



- Value of IBM z Systems for Oracle
- Oracle & IBM working together
- Customer examples
- Oracle HA & DR considerations
- Consolidation made easy



World's leading businesses run on the mainframe



92
of the top 100
worldwide banks



10
out of 10 of the world's
largest insurers



23
of the top 25
US retailers



23
out of 25 of the world's
largest airlines

Optimized for high performance data and transaction serving



Performance

- 141 high-performance cores
 - delivering 40% more capacity
 - move data 2X as fast
- 320 Separate channels of dedicated I/O

Speed

- Simultaneous multi-threading for up to 30% more throughput

Memory

- 10TB memory to eliminate I/Os for up to 70% faster response time to introduce new in-memory workloads



The Value of IBM z Systems for Oracle



Reduce cost by simplifying and optimizing your business process infrastructure with highly utilized Server resources designed for the digital economy

z Systems provide unsurpassed scalability and virtualization capabilities that can help to support Hundreds to thousands of virtual servers

Reduce the IT equipment footprint and maximize its efficiency with the simplicity of a single server Solution

Use automated failover and rapid recovery for business-critical applications and data

IBM z Systems are recognized as the most available, scalable, and secure platform¹

IBM z Systems

The right choice for Oracle mission critical workloads



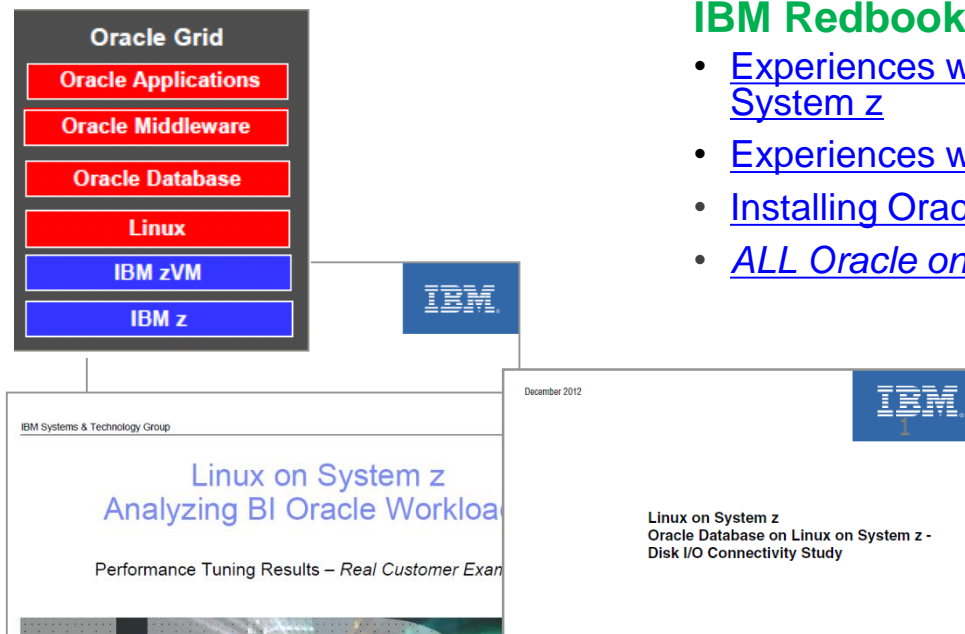
IBM z Systems® z13



IBM zEnterprise® BC12 (zBC12)

1 Forrester: Secure The Enterprise With Confidence Using A Mainframe Infrastructure, March 2013

Resources: IBM & Oracle working together




IBM Redbooks:

- [Experiences with Oracle Database 12c Release 1 on Linux on System z](#)
- [Experiences with Oracle 11gR2 on Linux on System z](#)
- [Installing Oracle 11gR2 RAC on Linux on System z](#)
- [ALL Oracle on System z Redbooks](#)

IBM Performance Papers:

- [Oracle Database on Linux on System z - Disk I/O Connectivity Study](#)
- [Analyzing BI Oracle Workloads](#)
- [Oracle Real Application Clusters on Linux on IBM System z: Set up and network performance tuning](#)



IBM and Oracle business relationship



- **Oracle Software Stack is certified and supported on certified distributions of Linux (RHEL or SLES) running natively in LPARs or as a guest operating system in z/VM virtual machines deployed on the System z platform. (My Oracle Support reference Doc ID: 417770.1).**
- Products certified for the System z platform qualify for the same level of support as any other certified Oracle platform.
- There is a dedicated Oracle team @ Oracle specially trained to support customers running Oracle with Linux on System z servers.
- Oracle support policy for security patches for Linux on System z servers:
 - Security patches also known as “CPU patches” are now included in the quarterly PSU (Patch Set Updates) for all platforms.
- Products ported to Linux on System z servers will be supported according to the Oracle Lifetime Support Policy.
- IBM and Oracle Business Relationship:
 - The IBM and Oracle Web site hosted by IBM at: <http://www.ibm.com/solutions/oracle>
 - The IBM Partner Relationship Web site hosted by Oracle at: <http://solutions.oracle.com/partners/ibm>
 - Frequently asked questions from IBM and Oracle customers about Linux on IBM System z <http://www.ibm.com/support/techdocs>

Oracle on z Systems



“ System z is the most cost-effective platform for large Oracle workloads. Whether our customers need to consolidate or isolate processes, our Oracle services would be impossible without it. ”

- Lubo Cheytanov, founder and co-owner, L3C LLP

Read the full story ZSC03285USEN

<http://www-03.ibm.com/software/businesscasestudies/us/en/corp?synkey=W133353R73108L21>

Business need

L3C LLP needed to bring the robust reliability, security and affordability of the mainframe to its cloud customers, while also using the platform as a key differentiator for its managed services.

Solution

L3C deployed IBM® System z® servers running Linux to provide companies of any size—including small, mid-sized and very large enterprises—with scalable, cost-effective, high-performance cloud services.

Benefits

L3C can now provide Infrastructure-as-a-Service (IaaS) options, with differentiated qualities of service and price performance, to help expand its reach and reduce costs for customers.

Delivers extreme reliability and cost savings to cloud customers using IBM z Systems

Sparda Datenverarbeitung runs Oracle on z



“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 zEnterprise platform is perfect for consolidating Linux workloads because of the high I/O bandwidth, business continuity with capacity backup features.”

*“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

<http://www.youtube.com/watch?v=c7Z19IB5AmE>



A DBA's view: Sparda Datenverarbeitung

Presented at DOAG 2014 (German Oracle User Group) conference

Experience report

8 Years of Oracle Databases on Linux on System z

Liebhard Bidner – DBA
November 20, 2014



DOAG 2014 - Konferenz

Erfahrungsbericht

8 Jahre Oracle Datenbanken auf Linux on zSeries



Liebhard Bidner – DBA
20.11.2014


Sparda-Datenverarbeitung eG 20.11.2014 – LB/SY-DBS

Seite 1



What benefits do we get from zLinux under z/VM?

- Fast provisioning of Linux guests via z/VM cloning mechanism
- Fast and simple extension of zLinux system resources (CPU, memory)
- High performance and security
- Simple licensing model and savings of Oracle SW costs
- Relocation of complete zLinux guest-systems to the other datacenter with z/VM Live Guest Relocation feature
- Mirrored disk storage subsystem between two datacenters
 - Disaster Recovery with GDPS / XDR in z/VM and Linux swaps disk mirror in case of failure automatically via HyperSwap
- Oracle database versions plus RAC are fully supported and certified by Oracle.

Welche Vorteile bietet uns z/Linux unter z/VM ? 

- Schnelle Bereitstellung von z/Linux Gästen durch den z/VM Cloning – Mechanismus
- Schnelle und einfache Erweiterung von z/Linux Systemressourcen (CPU, Hauptspeicher)
- Hohe Performance und Sicherheit
- Einfaches Lizenzierungsmodell und Einsparung bei den Oracle Softwarekosten
- Verlagerung des gesamten z/Linux Gastsystems mit z/VM Feature „live guest relocation“ in das andere Rechenzentrum
- Gespiegeltes Plattensubsystem über zwei Rechenzentrumsstandorte
Stichwort „Disaster Recovery“ mit GDPS / XDR (Geographical Dispersed Parallel Sysplex und Cross Platform Disaster Recovery) im z/VM und z/Linux schwenkt der Plattenspiegel bei Ausfall einer Seite mittels Hyperswap automatisch auf die andere
- Oracle Database Versionen plus RAC auf System Z mit z/VM und z/Linux sind von Oracle voll unterstützt und zertifiziert.

Sparda-Datenverarbeitung eG 20.11.2014 – LB/SY-DBS Seite 6



Conclusion – 8 years of Oracle under zLinux


Oracle is Oracle is Oracle ...
also with Linux on System z

During the last 8 years of production of our Oracle databases there were no failures or problem situations where we had to apply any Oracle patch specific to our database version on zLinux.

If required, it were always platform independent fixes for failure or problem situations.

Our Oracle databases under zLinux run absolutely stable, reliable, performing, and secured.

For us as Sparda Datenverarbeitung and our existing mainframe architecture this was and is the most effective and lowest cost platform for consolidation and virtualization.

FAZIT – 8 Jahre Oracle unter zLinux 

Oracle ist Oracle ist Oracle...

auch auf z/Linux.

In den letzten 8 Jahren Betrieb unserer Oracle Datenbanken gab es keine Fehler- oder Problemsituationen in denen wir einen Oracle Patch speziell für unsere Datenbankversion unter z/Linux einsetzen mussten.

Wenn, dann waren es immer plattformunabhängige Fehler- oder Problemsituationen die behoben wurden.

Unsere Oracle Datenbanken laufen unter z/Linux absolut stabil, zuverlässig, performant und abgesichert.

Für uns als Sparda Datenverarbeitung eG, mit unserer vorhandenen Mainframe Infrastruktur war und ist es die effektivste und kostengünstigste Plattform für eine Konsolidierung und Virtualisierung.

Sparda-Datenverarbeitung eG, 20.11.2014 – LB/SY-DBS Seite 11



Conclusion – 8 years of Oracle under zLinux


Oracle is Oracle is Oracle ...
also with Linux on System z

During the last 8 years of production of our Oracle databases there were no failures or problem situations where we had to apply any Oracle patch specific to our database version on zLinux.

If required, it were always platform independent fixes for failure or problem situations.

Our Oracle databases under zLinux run absolutely stable, reliable, performing, and secured.

For us as Sparda Datenverarbeitung and our existing mainframe architecture this was and is the most effective and lowest cost platform for consolidation and virtualization.

FAZIT – 8 Jahre Oracle unter zLinux 

Oracle ist Oracle ist Oracle...

auch auf z/Linux.

In den letzten 8 Jahren Betrieb unserer Oracle Datenbanken gab es keine Fehler- oder Problemsituationen in denen wir einen Oracle Patch speziell für unsere Datenbankversion unter z/Linux einsetzen mussten.

Wenn, dann waren es immer plattformunabhängige Fehler- oder Problemsituationen die behoben wurden.

Unsere Oracle Datenbanken laufen unter z/Linux absolut stabil, zuverlässig, performant und abgesichert.


Für uns als Sparda Datenverarbeitung eG, mit unserer vorhandenen Mainframe Infrastruktur war und ist es die effektivste und kostengünstigste Plattform für eine Konsolidierung und Virtualisierung.

Sparda-Datenverarbeitung eG, 20.11.2014 – LB/SY-DBS Seite 11

Oracle HA with System z



Oracle RAC



- 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

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 →

- Spare CPUs
- N+1 power supplies
- Chip sparring in memory
- Concurrent maintenance
- 50 years MTBF (system fail.)

About HA and DR ?

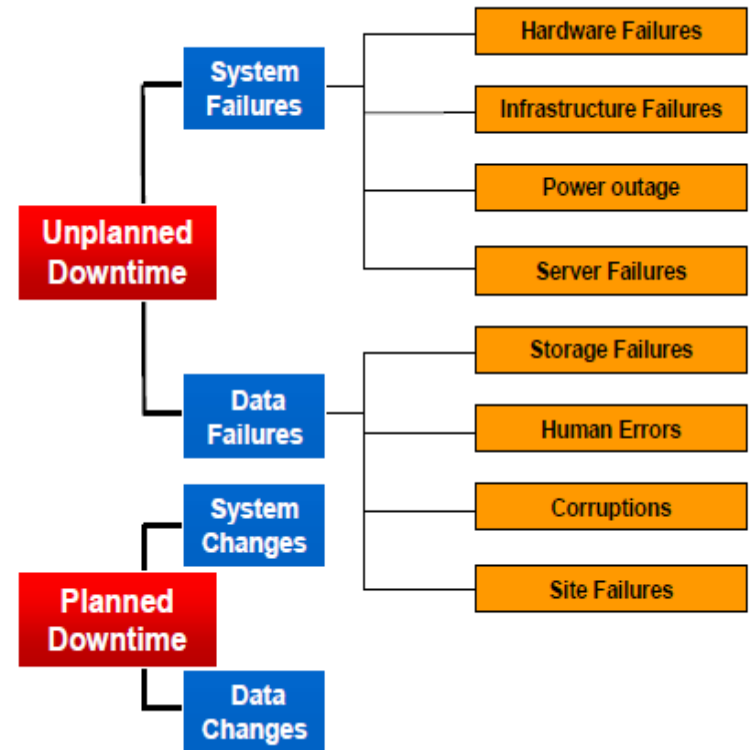


What is about High Availability ?

- **Prevent** outages before they occur
- **Tolerate** outages so they are transparent to the business

What is about Disaster Recovery ?

- **Recover** quickly if an outage does occur
- *Last Customer Insurance for its business*
- *Last insurance for customer Data Integrity*



HA and DR are not opposed, they are complementary !

HA solution should always have a DR solution to cover HA solution failure.

High-Availability & Disaster Recovery




- **Keep it simple**
 - Processes are part of HA
 - If the HA/DR architecture is too complex, it can be too difficult to manage with normal administrators skills
 - Avoid the necessity of expert skills to manage a disaster situation
- **Refrain from wanting everything**
 - RTO=0 (no outage)
 - RPO=0 (no data loss)
 - Long distance DR without performance concerns
- **Find the good balance**
 - Business needs versus costs and complexity
- **Think globally, for all your IT**
 - Avoid to use a different HA or DR solution for each of your application
 - Create and update HA/DR scenarios and procedures
 - Do not improvise during a crisis



IBM GDPS appliance for Linux on z Systems



- The IBM GDPS appliance for Linux on z Systems provides high availability in case of system, application or network failure
- GDPS/Peer to Peer Remote Copy (GDPS/PPRC) multiplatform resiliency capability for customers who do not run the IBM z/OS operating system in their environment.
- This solution is intended to provide IBM z Systems clients who run IBM z/VM and their associated guests, for instance, Linux on z Systems, with similar high availability and disaster recovery benefits to those who run on z/OS.
- **The implementation of the new GDPS Appliance for Linux will offer business continuity for Linux-only environments.**



IBM FlashSystem & Linux on z Systems

Highest Reliability, Maximum Performance



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



IBM FlashSystem is certified ([see SSIC](#)) to attach to Linux on z to meet your business objectives

Performance of Linux on z with FlashSystem

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

- **21x** reduction in response times*
- **9x** improvement in IO wait times*
- **2x** improvement in CPU utilization*

New FlashSystem 900 and z Systems FiconExpress16s I/O cards can provide an even higher throughput

Why IBM FlashSystem for Linux on System z?

Extreme Performance

Cut IO Wait Time
80%+



3X
increase
IOPS



Latency
Under
100
Microseconds

Enterprise Reliability



Highest
Reliability
levels



Purposed-built, Enterprise
Architecture

Macro Efficiency

No application
Or architecture
Changes



Reduce floor space,
power & cooling



Benefits &
economics out
weigh disk



IBM MicroLatency™

Servers, Applications and
Databases are FASTER!



Go FROM 7
milliseconds to 700
microseconds



* IBM internal test results with IBM FlashSystem 820 and FiconExpress4s


Micro-latency effects storage run time



I/O Serviced by **Disk**

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


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



I/O Serviced by **IBM FlashSystem**

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

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





Virtualization management for z/VM and Linux virtual servers



- ✓ Simplify the OS administrative and management of virtualized servers all from a single dashboard
- ✓ Reduce the time it takes to perform complex virtualization management tasks
- ✓ Extend the reach of existing skills to deploy images that host Oracle
- ✓ Improve the quality and consistency of operations with a current and accurate view of your system using IBM Wave discovery
- ✓ Reduce risk of errors by delegating management scope to the appropriate teams, allows DBA's to efficiently manage Oracle deployments
- ✓ Accelerate virtualization steps like virtual server cloning and provisioning to make the transformation to cloud easier



Consolidation – made easy



Non-virtualized servers
(single and/or multiple DB instances)



Virtualized




Considerations:

- Migration time & effort
- Flexibility
- Separation of applications
- Workload management (SLAs)
- Accounting

Consolidation – made easy



Non-virtualized servers
(single and/or multiple DB instances)

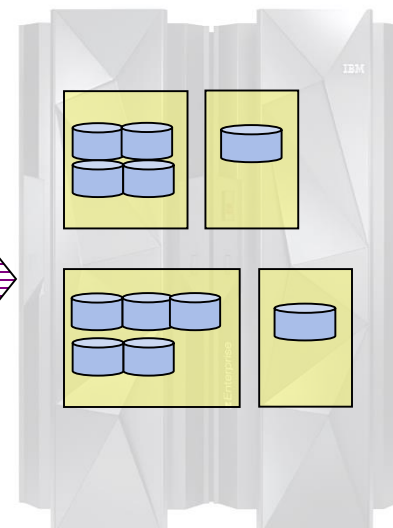


Consolidate

Virtualized



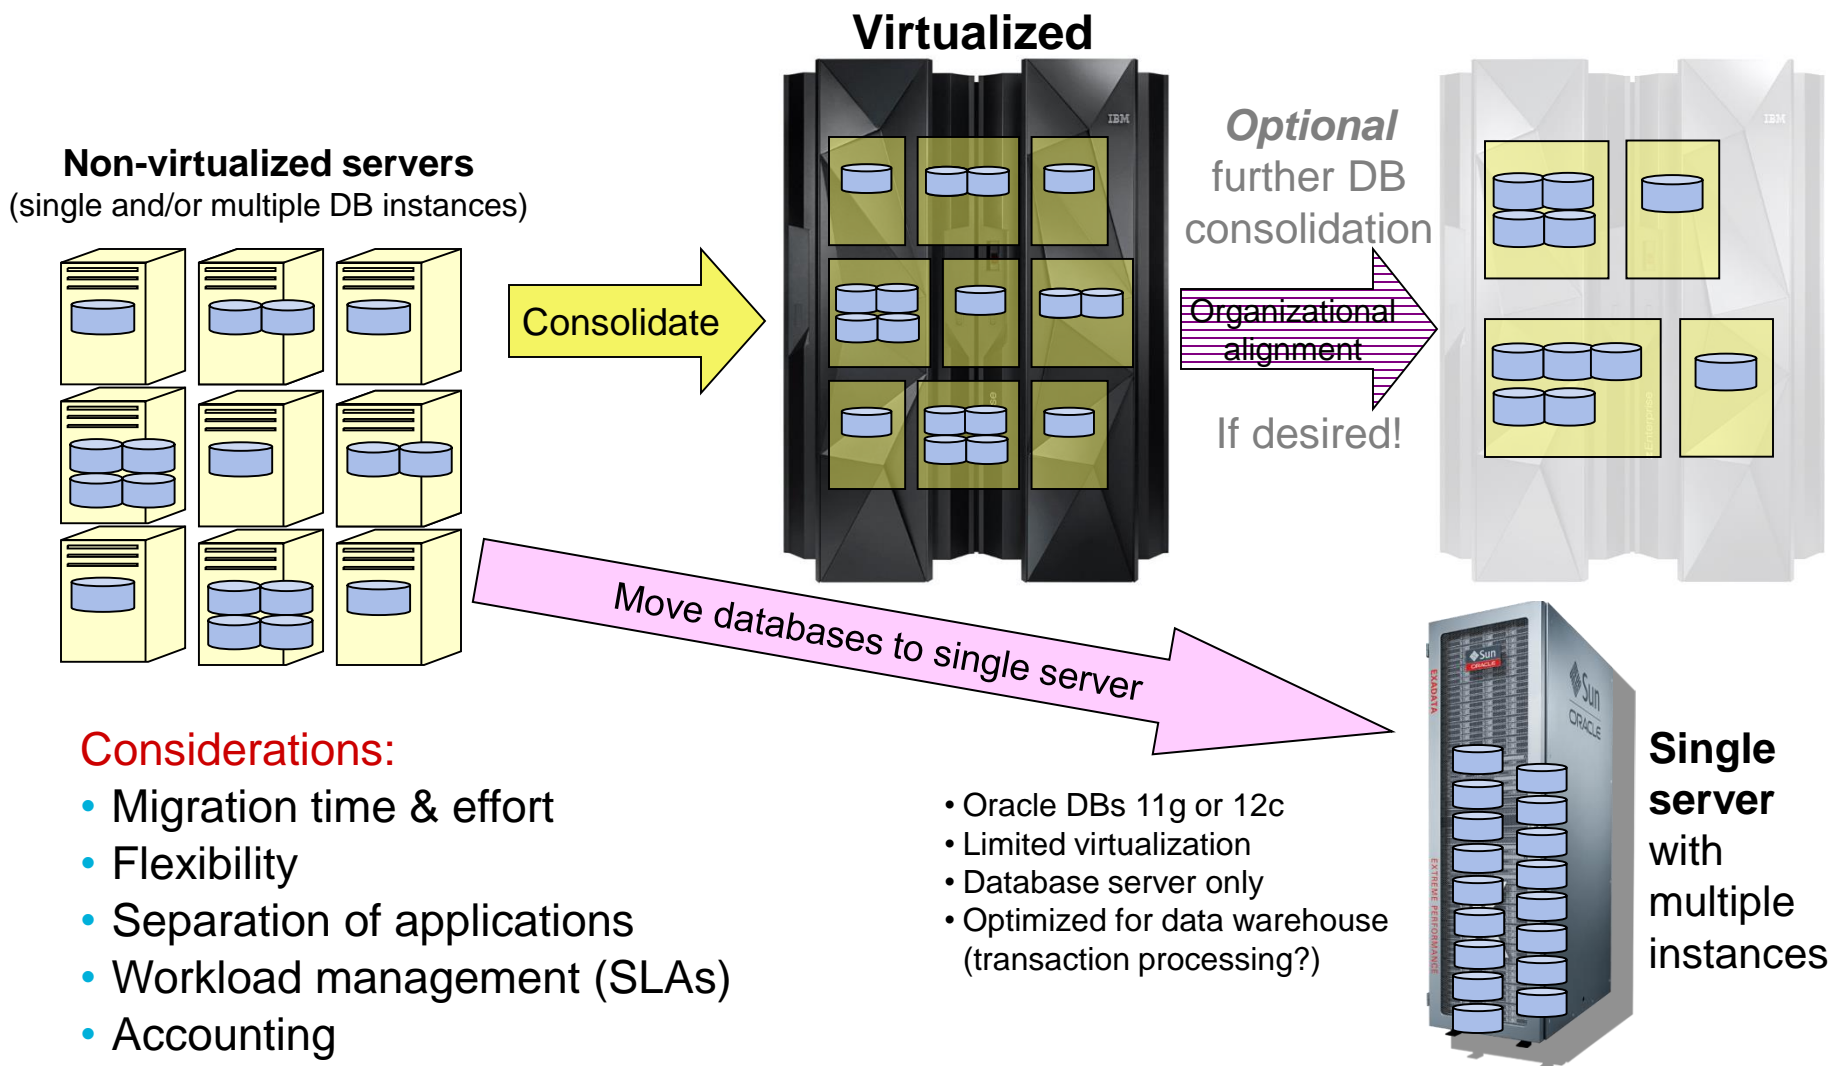
Organizational alignment



Considerations:

- Migration time & effort
- Flexibility
- Separation of applications
- Workload management (SLAs)
- Accounting

Consolidation – made easy



z Systems – Extreme Virtualisation



Build-in and Shared Everything Architecture


Hardware assisted virtualization



LPAR – PR/SM – up to 85 Logical Partitions




z/VM – 100's of Virtual Machines



Production and test & development on same box (separated by LPAR)

Production



Test & Development

Description of the Test




- There are 8 transaction types in the Brokerage high volume trading test, with the following transaction mix
 - CE (Customer Emulation)
 - Broker Volume (BV) – 20
 - Customer Position (CP) – 60
 - Market Watch (MW) – 180
 - Trade Lookup (TL) – 5000
 - Trade Order (TO) – 2000
 - Trade Update (TU) - 5000
 - MEE (Market Exchange Emulation)
 - Trade Result (TR) matches Trade Order (almost 1:1 in count)
 - Market Feed (MF) – 1 transaction per 10 Trade Result transactions
- CE transactions are driven through concurrent connections (threads) at a fixed interval rate of 70 ms
- MEE transactions are automatically triggered by the CE Trade Order transactions

Calculating TPS for the Test



- Since CE threads drive transactions at a minimum interval of 70 ms (can be longer when response time exceeds the injection interval):
 - Max # Transactions per second from each CE thread = $1000 \text{ ms} / 70 \text{ ms} = 14.286$
 - Total # of CE Transactions per 60 second interval (measurement granularity) across 16 threads = $14.286 \times 60 \times 16 = 13,715$
 - Estimated # Trade Order transactions in 60 second interval, using the transaction mix in previous slide = $2,000 / 12,260 \times 13,715 = 2,237$
- MEE thread transaction execution is triggered by the Trade Result transactions at the following rate:
 - Estimated # Trade Result transactions = # Trade Order Transactions = 2,237
 - Estimated # Market Feed transactions = $1/10 \times \# \text{ Trade Result transactions} = 1/10 \times 2,237 = 224$
- Total # CE + MEE Transactions in 60 second interval = $13,715 + 2,237 + 224 = 16,176$
- Max TPS assuming 100% of CE transactions execute within 70 ms = $16,176 / 60 = 270$ TPS
- Assuming SLA, where 90% of transactions need to complete within measurement interval of 60 seconds, Min TPS = $0.9 \times 270 = 243$ TPS

Conceptual Diagram for the Test

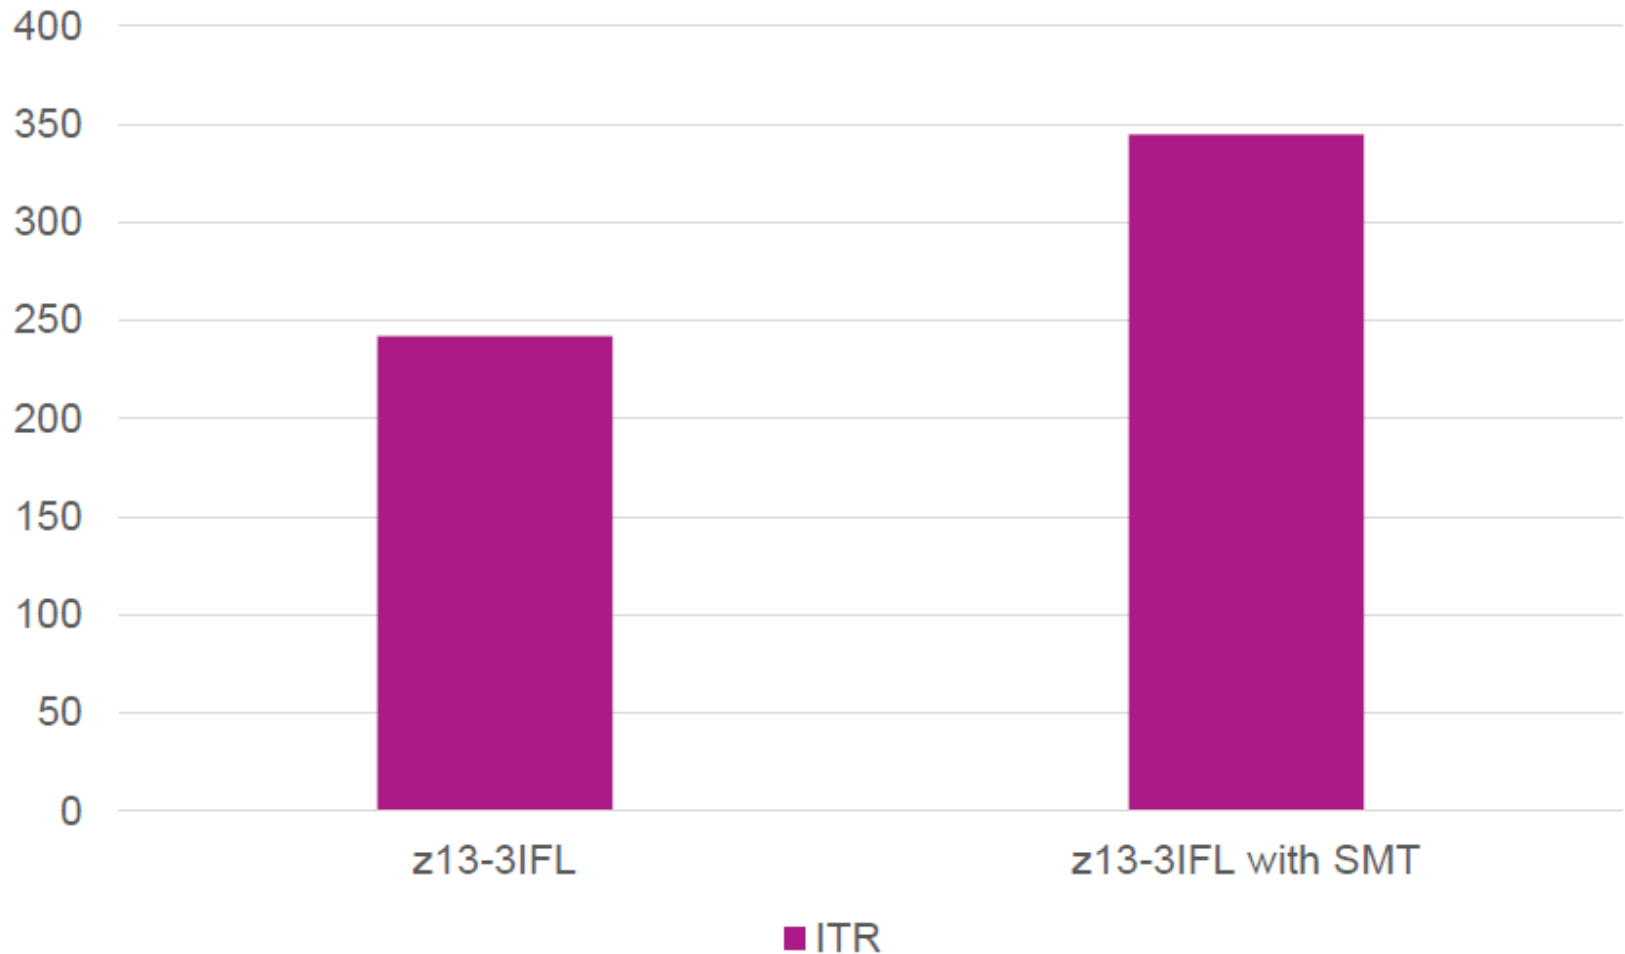


X Transactions types not used in the Test

SMT measurement for brokerage workload



42.5% SMT Benefit






Outstanding Scalability


Potential for economic growth and flexible configuration

- Highest levels of resource sharing – including the over-commitment, cooperative memory management, I/O bandwidth
- In-memory emulated storage achieves data transfers on memory-speed
- Very fast internal I/O connections, no external networking
- Dynamically add processors, memory, I/O adapters, devices and network cards ... no disruption



- Unused resources for peak utilization are provided to other virtual servers during off-peak hours ... automatically

Scale horizontally



Linux on z Systems

means an enterprise grade Linux solution




While „Linux is Linux“, z Systems server and virtualization technologies provide an enhanced Linux solution

Having an enterprise grade Linux solution brings:

- IT simplicity to run hundreds of workloads on one server
- Workload integration inside a single server
- Flexible server provisioning and growth inside the server
- High productivity through efficient life cycle management
- High utilization of shared resources
- Highest levels of security and quality of service – including business continuity

Linux on z Systems provides security, availability, and scalability to deploy (consolidate) all kinds of workloads



IBM and Oracle workshop for Running Oracle DB R12c on IBM z13



What: IBM and Oracle workshop for Running Oracle DB R12c on latest Mainframe technology

When: 29 – 30 June 2015

Where: IBM Client Center at IBM Germany Research & Development, Schoenaicher Str. 220, D-71032 Boeblingen, Germany

General Information

Linux on z Systems takes advantage of the qualities of service in the latest IBM® z13 server and in z/VM®, making it a robust industrial strength Linux. This provides an excellent platform for hosting Oracle® solutions that run in your enterprise, especially for Oracle DB R12c.

Join us for a two days workshop for Running Oracle DB on Linux on z Systems. Meet with experts from IBM and from Oracle to discuss the new and exciting functions of Oracle DB R12c and learn how the combination with the IBM z13 platform can help you optimize your Oracle infrastructure and gain advantages for your data center.

Enrollment and Costs/Fees

This workshop is offered at **no-fee** to qualified customers and Business Partners. Travel, hotel and living expenses will be your responsibility. Please click on the link below to enroll:

www.ibm.com/events/IBMOracleWorkshop



Questions



Siegfried Langer
Business Development Manager
z/VSE & Linux on System z

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

Phone: +49 7031 - 16 4228

Siegfried.Langer@de.ibm.com



Continue growing your IBM skills



ibm.com/training provides a comprehensive portfolio of skills and career accelerators that are designed to meet all your training needs.

- **Training in cities local to you** - *where and when you need it, and in the format you want*
 - Use [IBM Training Search](#) to locate public training classes near to you with our five Global Training Providers
 - Private training is also available with our Global Training Providers
- Demanding a high standard of quality – **view the paths to success**
 - Browse [Training Paths](#) and [Certifications](#) to find the course that is right for you
- If you can't find the **training that is right for you** with our Global Training Providers, we can help.
 - Contact IBM Training at dpmc@us.ibm.com



Global Skills Initiative