

IPL3220

IT Optimization with Linux on z Systems

Siegfried Langer
Business Development Manager z/VSE and Linux on System z
IBM Germany Research & Development



2015

IBM Systems Technical University


IBM z Systems • IBM Power Systems • IBM Storage

October 5–9 | Hilton Orlando, Florida




Agenda

- **Infrastructure optimization**
 - *Cost and business value*
- **Linux on z Systems is Linux**
 - *What makes it different*
- **Application examples**
 - *Typical workloads*




Industry trend

– which industry product trend is depicted here?



Industry trend: commercial aircraft industry

Commercial Aircraft Cruise Speed



Fit for Purpose



All of these “tools” can very quickly move a person from one place to another.
But, which one is the right tool to move one person?
Fifty people? Five hundred people?

Platform Selection Is All about Using the Right Tool for the Right Job



IBM z Systems

- Up to 141 high performance cores (@ 5 GHz)
- Simultaneous multi-threading (SMT-2)
- Up to 10TB Memory
- Up to 320 16 GB/sec Channels
- Single instruction, multi data (SIMD)
- Specialty engines zIIPs, IFLs, ICFs



zBC 12




- Up to 13 high performance cores (@ 4.2 GHz)
- Up to 6 CPs (plus specialty engines)
- Up to 512GB Memory
- Specialty engines zIIPs, IFLs, ICFs
- Upgradable



IBM z Systems

z13




zBC 12



IBM z13: Advanced system design optimized for digital business

System I/O Bandwidth - 832 GB/Sec*




Meet the IBM LinuxONE Systems

*THE MOST TRUSTED, EFFICIENT AND HIGH PERFORMANCE
ENTERPRISE-GRADE LINUX PLATFORM*



Emperor™

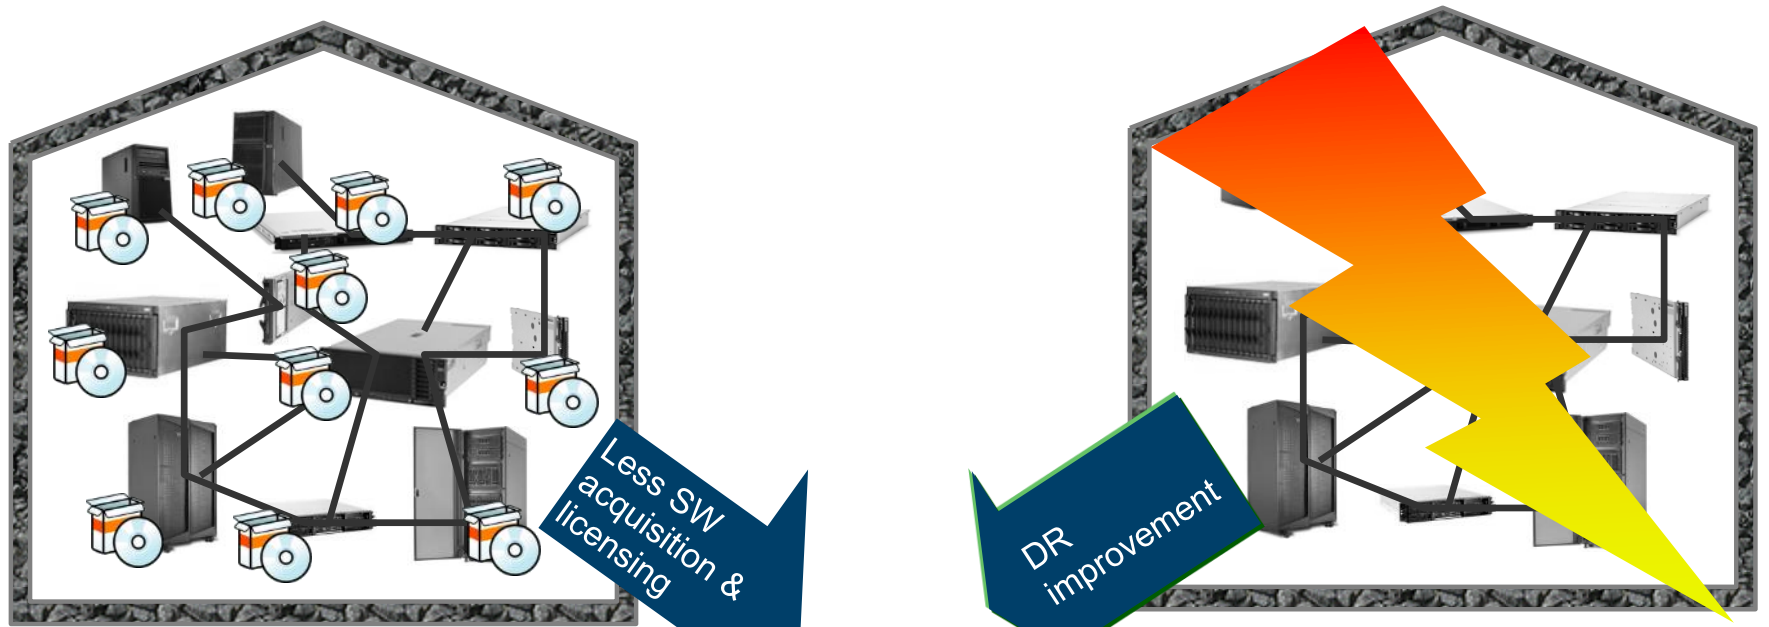


Rockhopper™



The Difference of Deployment on System z

Examples: Software Costs and Disaster Recovery



Less SW acquisition & licensing

DR improvement



Distributed software is often priced by the number of processor cores.

On System z, one IFL is equivalent to one core!



Coordinated near-continuous availability and DR solution for critical data on Linux



Complete Range of IT Cost Factors


Cost factors highlighted in yellow are the only ones that organizations routinely assess.

Full Range of IT Cost Factors			
Availability <ul style="list-style-type: none"> High availability Hours of operation 	Additional development and implementation <ul style="list-style-type: none"> Investment for one platform – reproduction for others 	Security <ul style="list-style-type: none"> Authentication and authorization User administration Data security Server and OS security RACF versus other solutions 	Controlling and accounting <ul style="list-style-type: none"> System analysis Cost
Backup/restore/site recovery <ul style="list-style-type: none"> Backup Disaster scenario Restore Effort for complete site recovery Storage area network (SAN) effort 	Operations <ul style="list-style-type: none"> Monitoring and operating Problem determination Server management tools Enterprise-wide integrated server management 	Operating concept <ul style="list-style-type: none"> Development of an operating procedure Feasibility of the developed operating procedure Automation 	Integration <ul style="list-style-type: none"> Integrated functionality versus functionality to be implemented (possibly with third-party tools) Balanced system Integration of and into standards
Infrastructure <ul style="list-style-type: none"> Space Power Network infrastructure Storage infrastructure Initial hardware costs Software costs Maintenance costs 	Resource utilization and performance <ul style="list-style-type: none"> Mixed workload/batch Resource sharing: shared nothing versus shared everything Parallel Sysplex versus other concepts Response time Performance management Peak handling and scalability Skills and resources <ul style="list-style-type: none"> Personnel education Availability of resources 	Deployment and support <ul style="list-style-type: none"> System programming <ul style="list-style-type: none"> Consistent OS and software levels Database effort Middleware <ul style="list-style-type: none"> Software maintenance Software distribution across firewalls Application <ul style="list-style-type: none"> Technology upgrade System release change without interrupts 	Further availability aspects <ul style="list-style-type: none"> Planned outages Unplanned outages Automated takeover Uninterrupted takeover (especially for databases) Workload management across physical borders Business continuity Availability effects for other applications or projects End user service End user productivity Virtualization



Linux is Linux,

but are all Linux infrastructure solutions identical?



No

, while Linux is Linux, the underlying infrastructure (hardware and infrastructure software) directly affects the Linux workloads.



Enterprise grade Linux solution

While „Linux is Linux“, the underlying platform is providing differentiation of the Linux solutions.

An “enterprise grade Linux” solution, in our understanding, has defined characteristics:

- **IT simplicity, allowing to run up to hundreds of different workloads in parallel on one server**
- **Easy workload integration of new and existing data and applications**
- **Flexible server provisioning, simple to manage**
- **High productivity, based on efficient systems and life cycle management**
- **Highest resource utilization levels**
- **High levels of quality of service – security, availability, reliability**

“Enterprise-grade isn’t just about specific features, rather it is about delivering a strategy that enables a consistent architectural model with the support and service necessary for [the] ... complex environment that organizations find themselves in.” - Ben Kepes, contributor to Forbes

www.forbes.com/sites/benkepess/2013/12/18/what-does-enterprise-grade-really-mean



Why Linux on System z

- **Higher utilization**
 - ✓ Up to 100% (typically 80%+) → lower software license cost
- **Scalability**
 - ✓ Resources can be added transparently → incremental growth as needed
- **Faster processor**
 - ✓ Lower number of cores (IFLs) → lower software license cost
- **Capacity on Demand (CoD) and Capacity Backup (CBU)**
 - ✓ Flexibility to react to peaks or disaster → being prepared at lower cost
- **Highest reliability**
 - ✓ Better availability → higher quality of service
- **Integration and virtualization**
 - ✓ Two levels of virtualization (LPAR and z/VM) → security, flexibility, consolidation of conflicting workloads on single physical server
 - ✓ DR advantages → lower complexity and integration with z/OS or z/VSE systems
 - ✓ Easier systems management → lower management cost



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.*

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.**



Nationwide[®]
Insurance

“Our goal was server optimization and our approach was virtualization.”

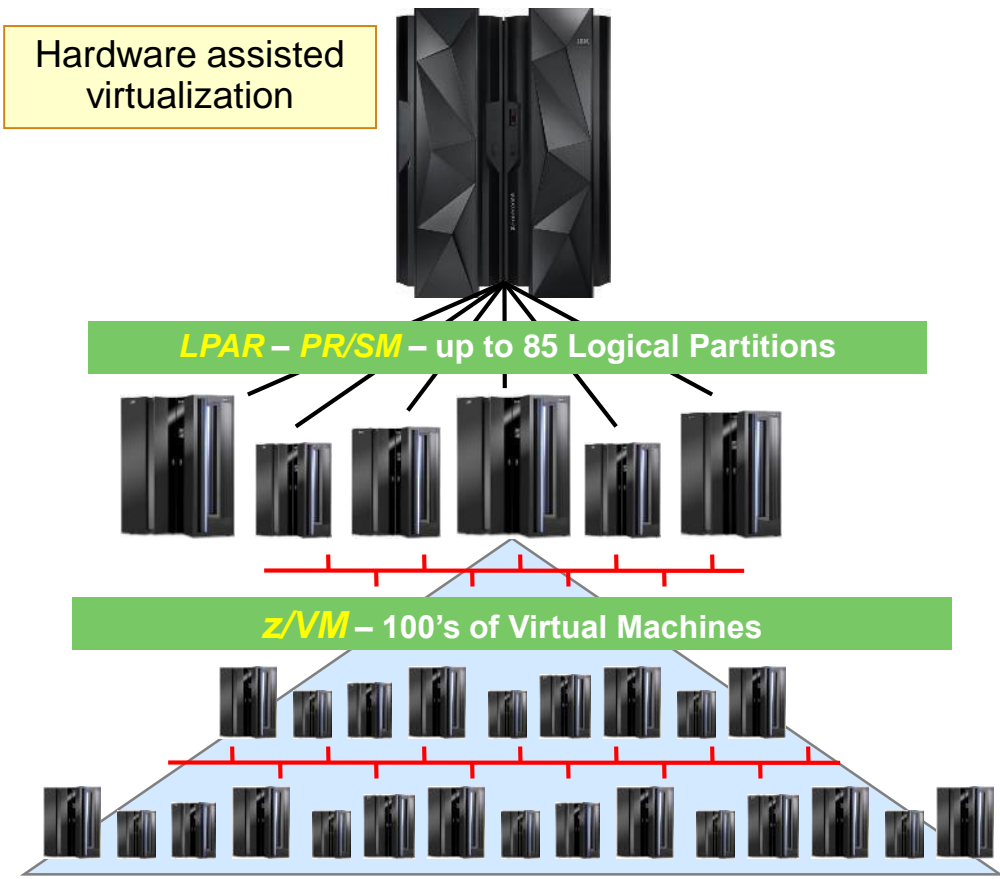
— Guru Vasudeva, Nationwide vice president and CTO

enterprisesystemsmedia.com/article/optimized-computing-nationwides-linux-on-system-z-virtualization-reduces-tc



z Systems – Extreme Virtualisation

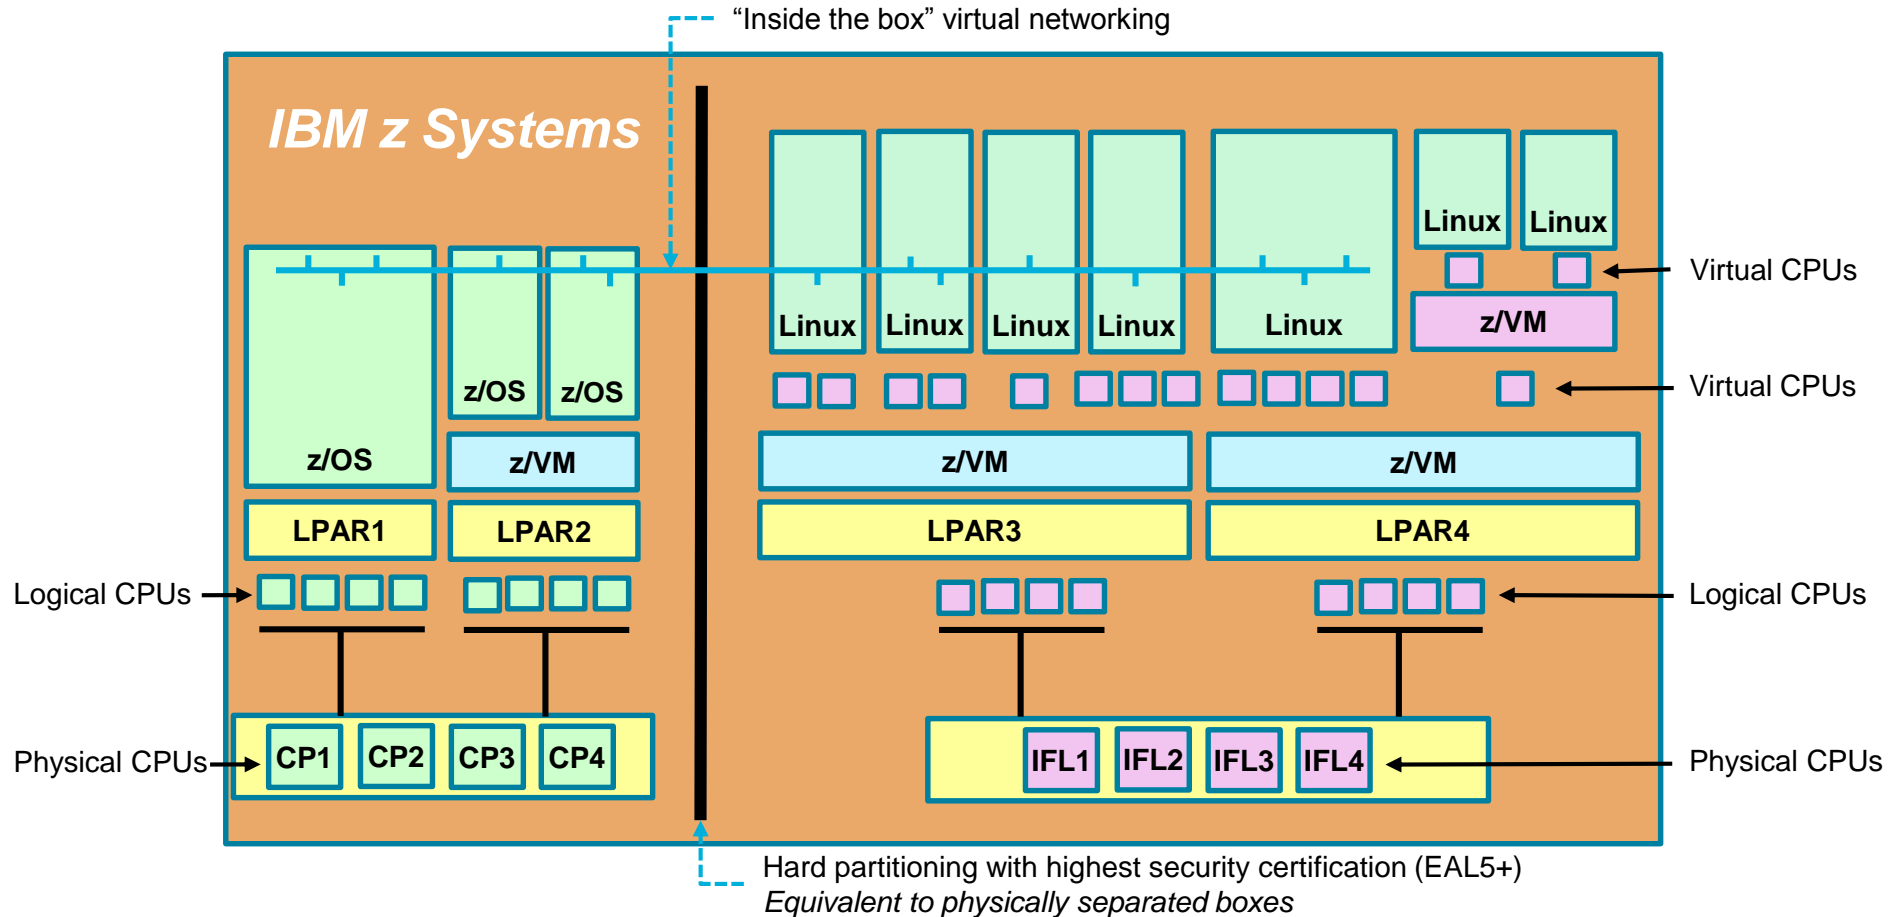
Build-in and Shared Everything Architecture



- ### z Systems
- Provisioning of virtual servers in seconds
 - High granularity of resource sharing (<1%)
 - Upgrade of physical resources without taking the system down
 - Scalability of up to 1000's of virtual servers
 - More with less: more virtual servers per core, sharing of physical resources
 - Extensive life-cycle management
 - HW-supported isolation, highly secure (EAL5+ or EAL4+ certified)

- ### Distributed platforms
- Limited virtual server scalability per core
 - Scaling requires additional physical servers
 - Operational complexity increases with growth of virtual server images
 - Single level of virtualization (no VM on VM)

IBM z Systems Virtualization Options



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

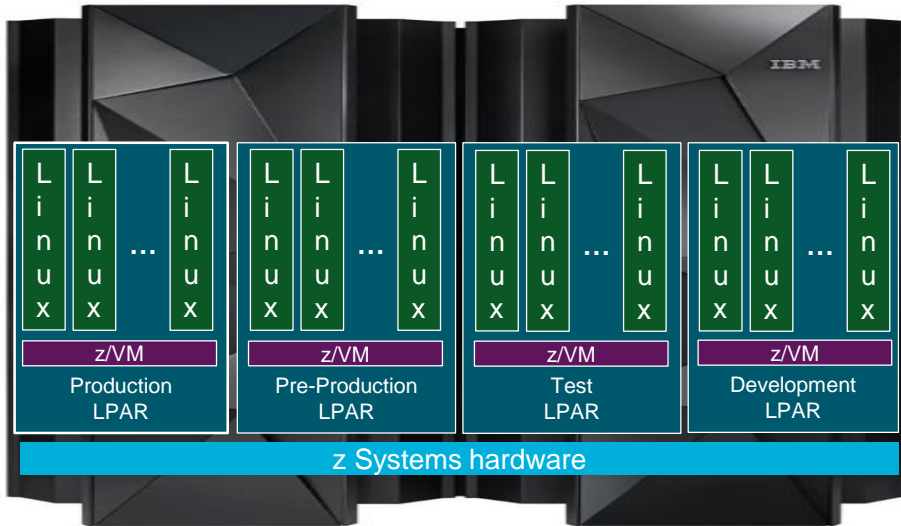


Test & Development




Comparing System Environments

Example



- LPAR allows to run potentially conflicting workloads in same box
- z/VM hypervisor with high level of resource management

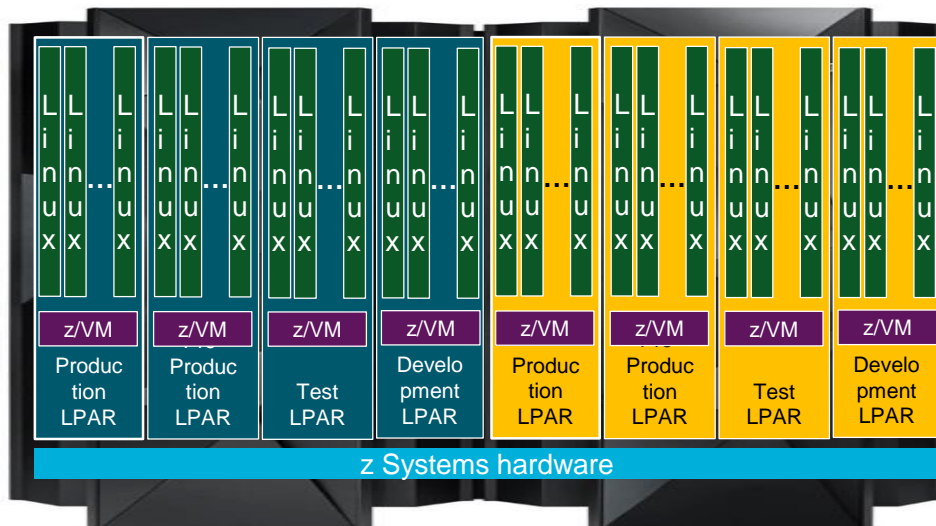


- **Multiple systems required**
 - Total number of cores increases
 - SW license cost increases
 - External network required
 - No/limited sharing of resources



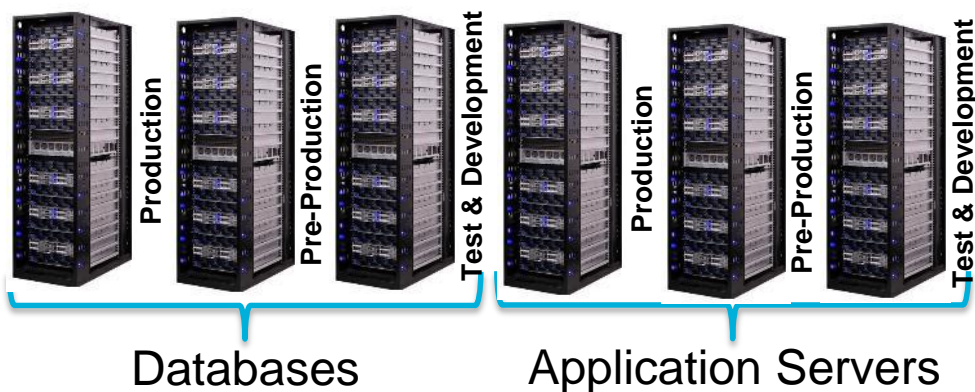
Comparing System Environments (Example)

Adding Application Servers



- **Adding LPARs**

- Adding IFLs and memory as required
- Internal interconnection (no cables)
 - HiperSockets (LPAR to LPAR)
 - VSWITCH & VLAN in z/VM



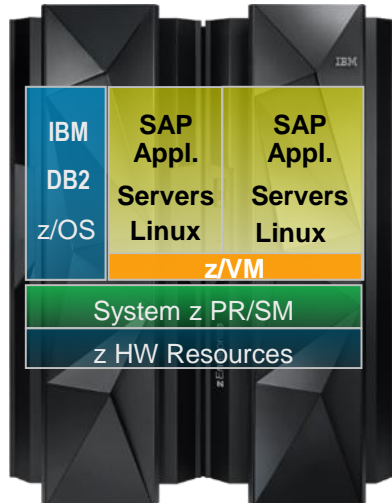
- **Multiple systems required**

- Management complexity increases
- Big external network required
- No/limited sharing of resources

Additional High Availability and Disaster Recovery servers not shown here!



SAP Application Server Deployment and Consolidation on System z



Business Continuity

- DB on z/OS
- Data Sharing in Parallel Sysplex®
- HA with Tivoli System Automation

Server Consolidation

- Internal near memory-speed communication
- Scale-up and scale-out capabilities
- Fabulous performance throughout

Embasa - Brazil

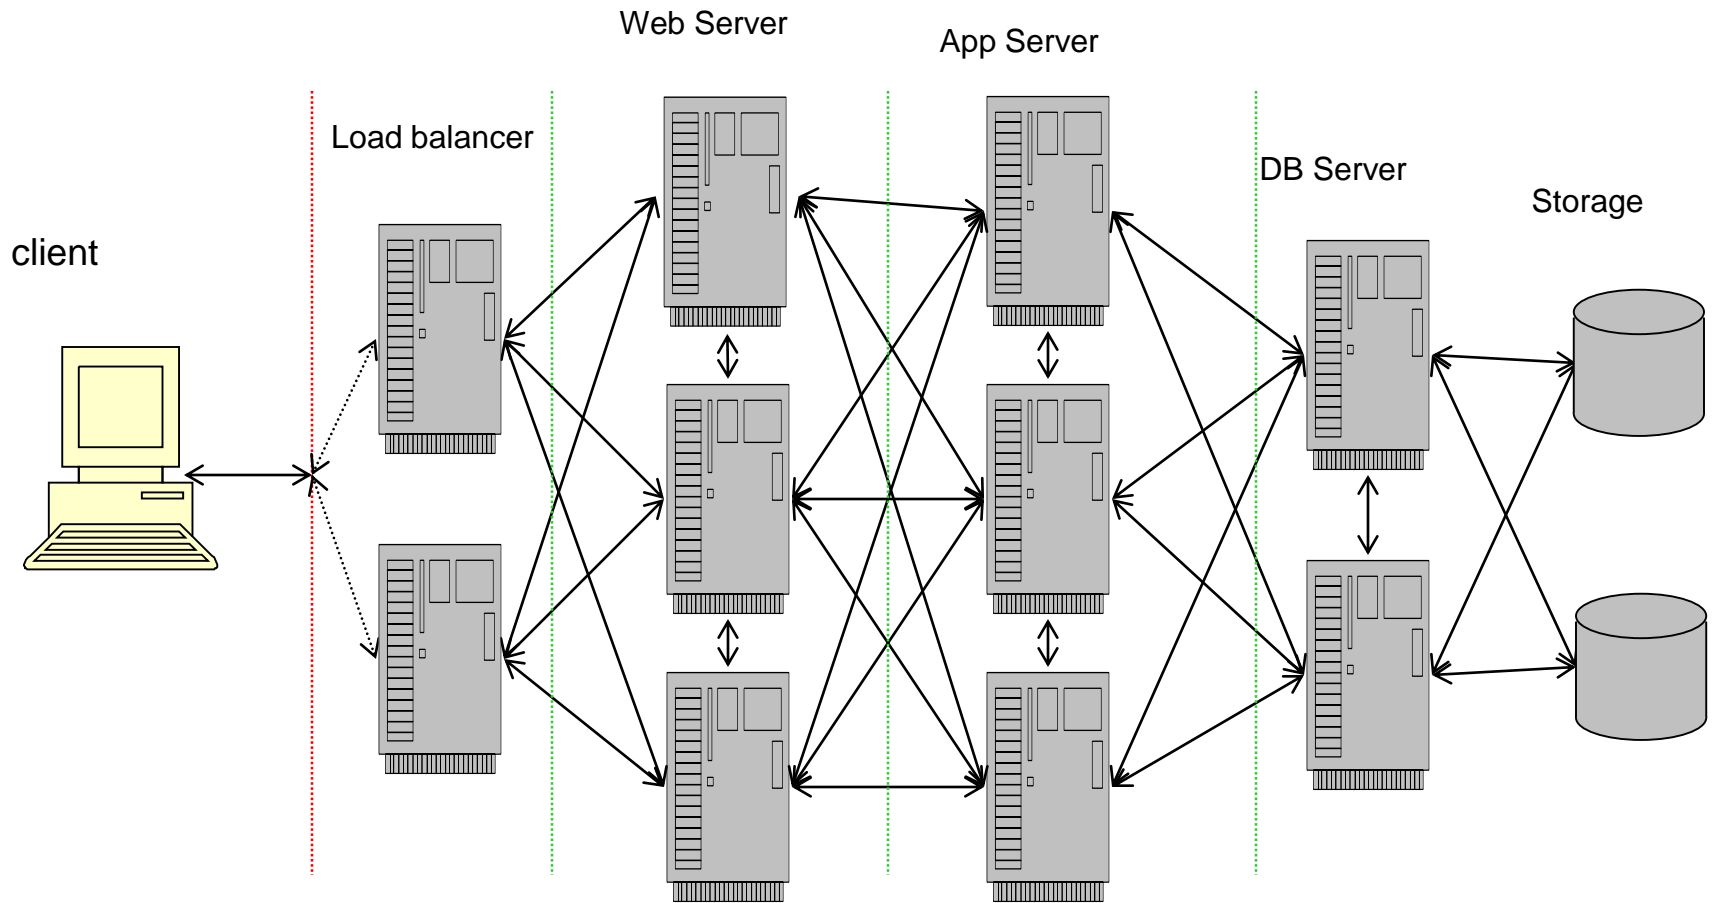
- **Manages one of the largest water treatment services**
- **Needed a high-performance, cost-effective way to introduce SAP software while continuing with the tried and trusted database solution**
- **Commercially attractive “Solution Edition” gave confidence to go ahead**

Endress+Hauser – Germany

- **Specialist in measurement technology; 89 companies across 42 countries**
- **Detailed cost-benefit analysis compared Linux on System z to Power/x86 servers. z/OS, z/VM and a total of 80 IFLs**
- **Simple and intuitive user management tools make it possible for just 1.5 FTEs to administer the entire Linux landscape**



HA with Independent Tiered Execution Streams



Insurance Company Consolidated 292 Servers to a z10

Before

3560E-24TD	8
3560E-24TD	17
3560E-12D	6
50 Ft UTP Cable	584
10GB Eth Fiber Cable	60

Over 600 Cables

Backbone Network

The diagram only shows **30** of **292** servers

After

Backbone


Just 10 Cables!

WebSphere Application Server (WAS) on Linux on z Systems

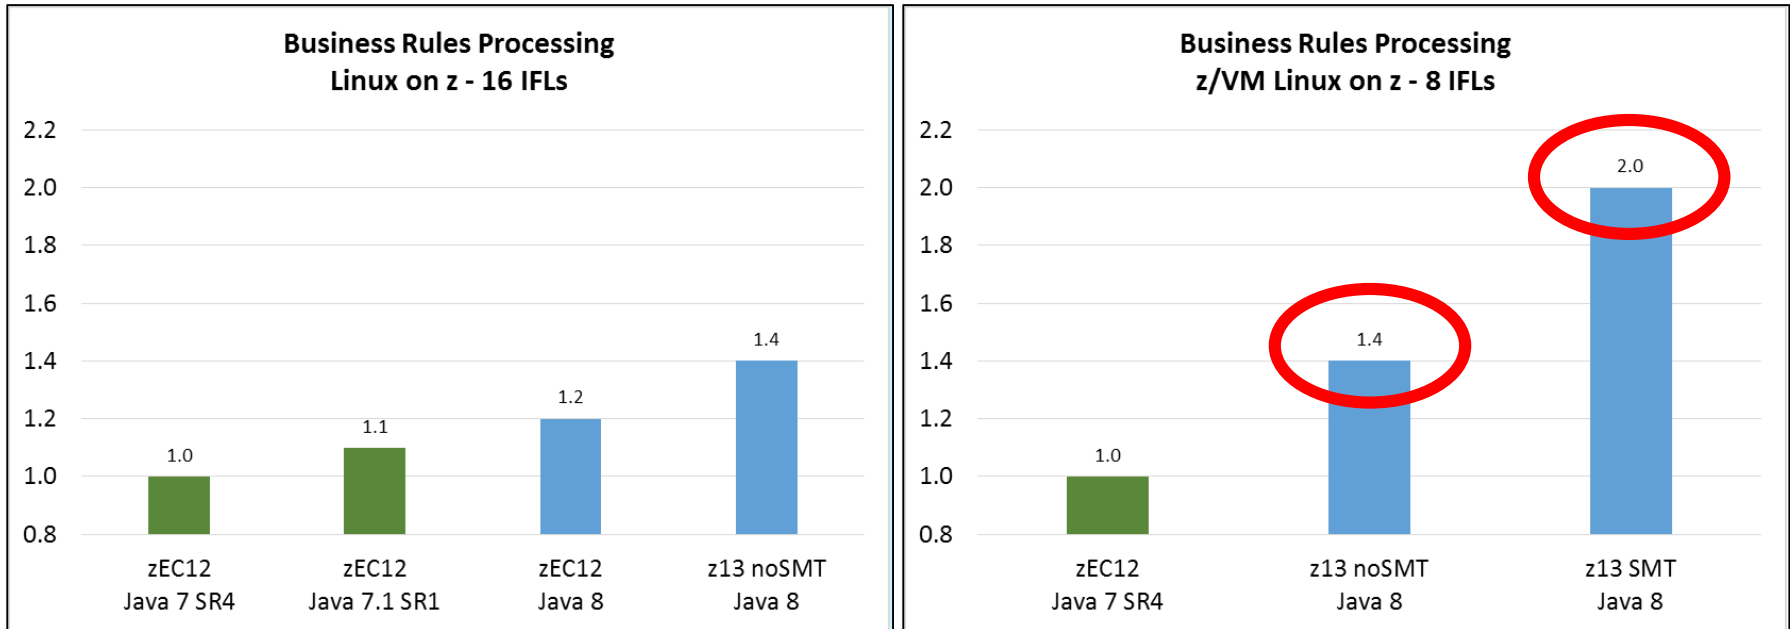


4x aggregate HW, SDK and WAS improvement:

WAS 6.1 (Java 5) on z9 to WAS 8.5 (Java 7) on zEC12



IBM Business Rules Processing with IBM Java 8 and z13

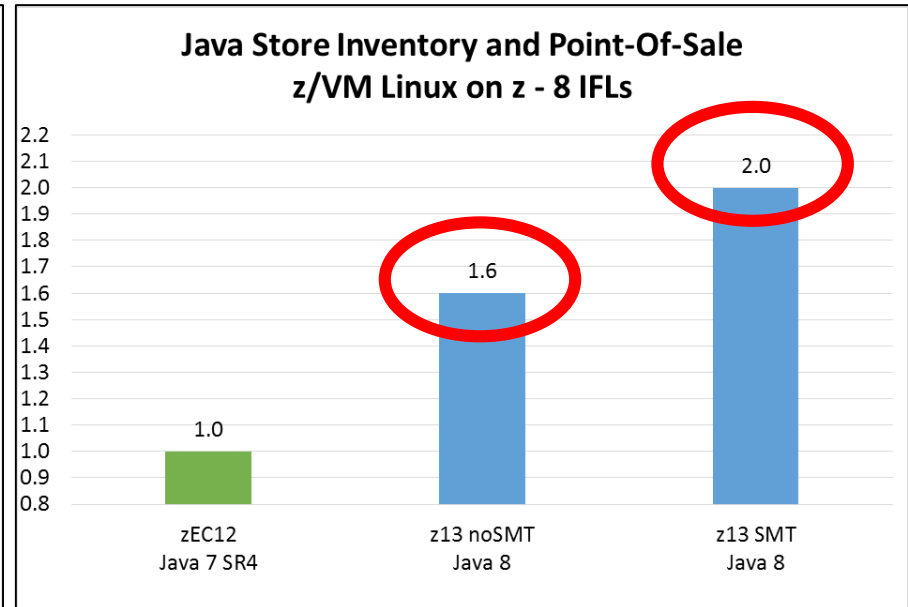



In this specific example, SMT gives you **around 43%** throughput improvement

(Controlled measurement environment, results may vary)



Java Store, Inventory and Point-of-Sale App with IBM Java 8 and z13




In this example, SMT gives you ***around 25%*** throughput improvement

(Controlled measurement environment, results may vary)




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%.

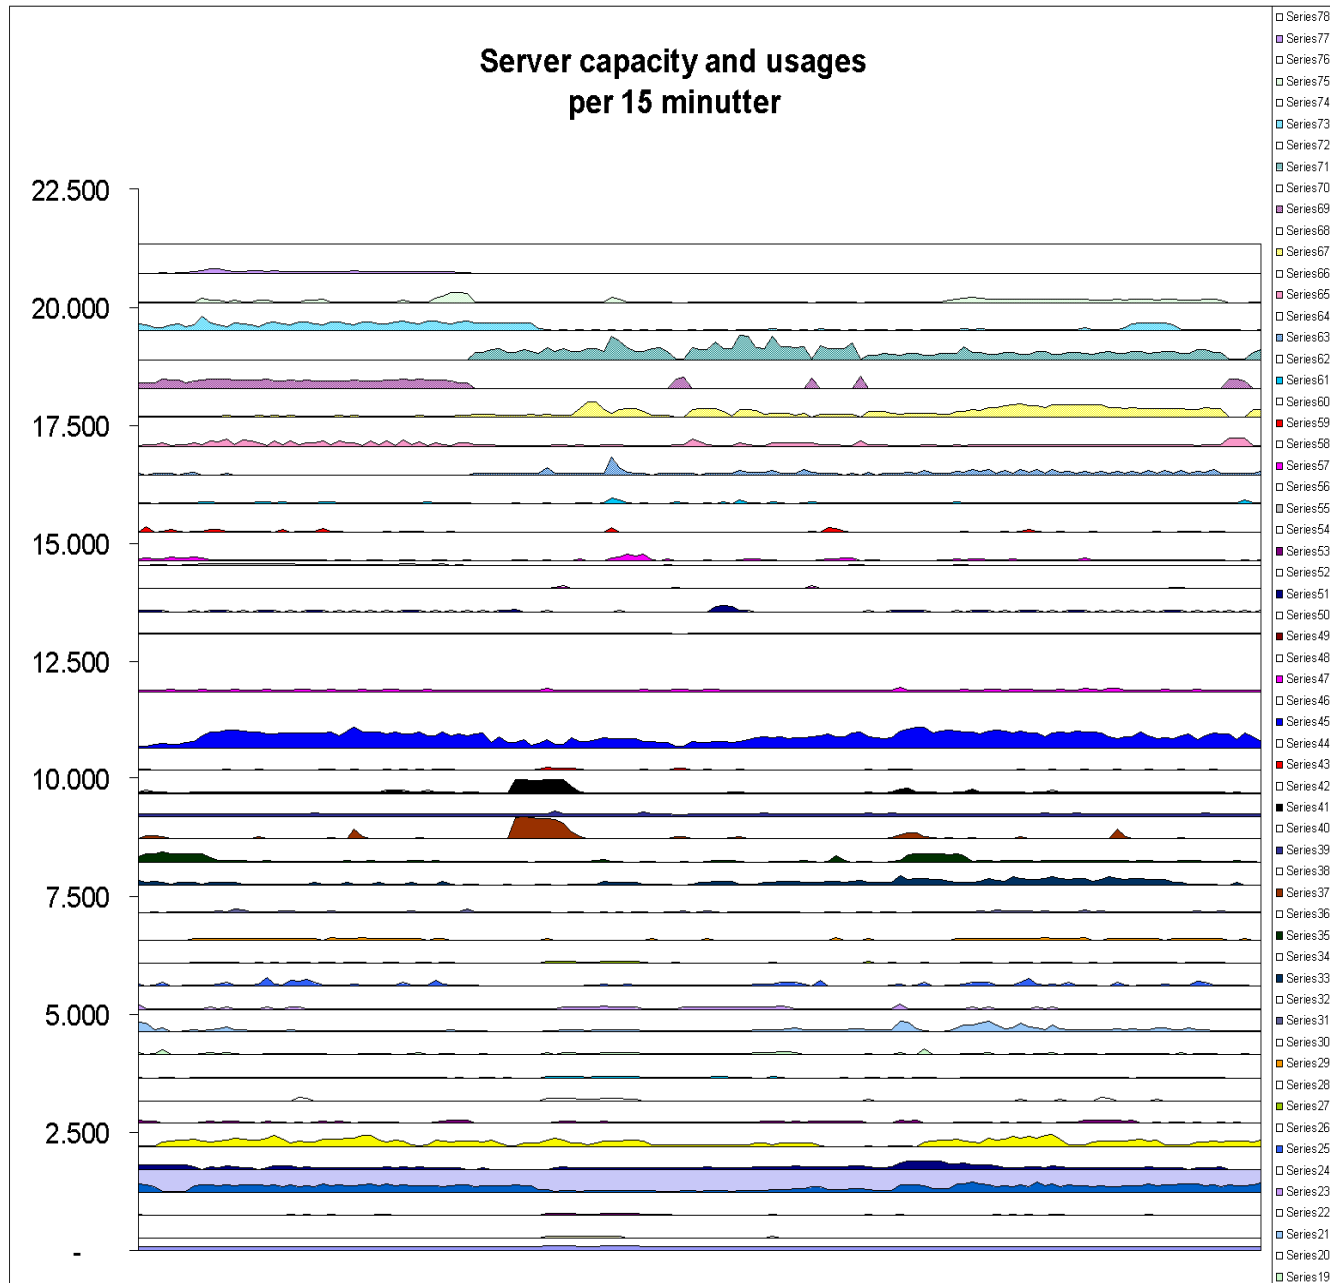
Mixed Utilization on IBM High End Servers



IBM High End Server: Up to 100% utilization

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






RPE2
(from Ideas International)



Server capacity and usages per 15 minutter



- Series78
- Series77
- Series76
- Series75
- Series74
- Series73
- Series72
- Series71
- Series70
- Series69
- Series68
- Series67
- Series66
- Series65
- Series64
- Series63
- Series62
- Series61
- Series60
- Series59
- Series58
- Series57
- Series56
- Series55
- Series54
- Series53
- Series52
- Series51
- Series50
- Series49
- Series48
- Series47
- Series46
- Series45
- Series44
- Series43
- Series42
- Series41
- Series40
- Series39
- Series38
- Series37
- Series36
- Series35
- Series34
- Series33
- Series32
- Series31
- Series30
- Series29
- Series28
- Series27
- Series26
- Series25
- Series24
- Series23
- Series22
- Series21
- Series20
- Series19

RPE2
(from Ideas International)



The Met Office forecasts a bright outlook for Linux on z Systems

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



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




Enterprise Reliability



Macro Efficiency



IBM MicroLatency™



* IBM internal test results with IBM FlashSystem 820 and FiconExpress4s

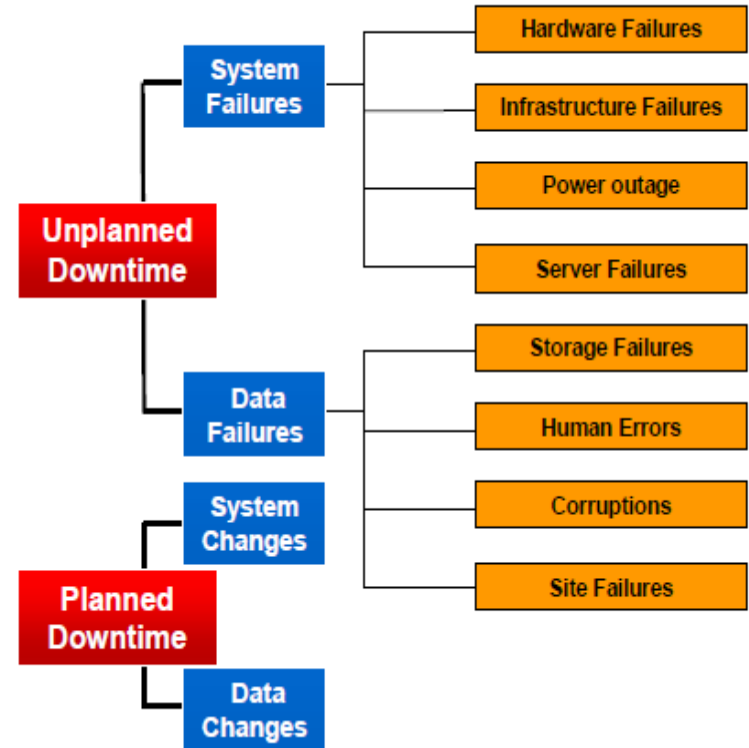


What is about High Availability ?

- **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.



IBM GDPS for Linux on z Systems




- **IBM GDPS can be used for Linux on z Systems, providing high availability in case of system, application or network failure**
- GDPS/Peer to Peer Remote Copy (GDPS/PPRC) multiplatform resiliency capability



- z/VM HyperSwap between sites
- GDPS integration provides for automation
- Requires IBM GDPS with z/OS or new GDPS appliance offering



Glitch Freezes NYSE Trading for Hours



Source: Factset (volume); NYSE (closings)

THE WALL STREET JOURNAL.

Nobody wants to make the news with such a message!

July 8, 2015 – The Wall Street Journal:

A glitch forced the New York Stock Exchange to halt trading for nearly four hours Wednesday, an outage that unnerved Wall Street and revived concerns about the fragility of the technological systems that underpin financial markets.


The NYSE accounts for a far smaller slice of overall trading than it once did, and **investors were able to route buy and sell orders through the Nasdaq Stock Market and other platforms that continued to function normally.**

It took about an hour and a half to find the issue, the person familiar with events said. It was a problem in the way computers at the NYSE data center were communicating with one another. The precise cause is still under investigation, but it may have been related to a software update performed the night before, the person said.



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.)



Virtualization and Cloud Portfolio for Linux on z Systems and z/VM



Virtualization Infrastructure & Virtualization Management

IBM z/VM 6.3

- Support more virtual servers than any other platform in a single footprint
- Integrated OpenStack support



IBM Wave for z/VM

- A graphical interface tool that simplifies the management and administration of z/VM and Linux environments

Differentiation

Entry Level Cloud Standardization & Automation



IBM Cloud Manager with OpenStack

- A simple, entry level cloud management stack
- Based on OpenStack
- Supports System z environments
- Formerly known as SmartCloud Entry

Standardization

Advanced Cloud Orchestration & Optimization



IBM Cloud Orchestrator

- Builds on functionality of **IBM Cloud Manager with OpenStack** and adds runbook automation and middleware pattern support for workload deployment
- Based on OpenStack
- System z support as “managed-to”
- Formerly known as SmartCloud Orchestrator

Service Lifecycle Management



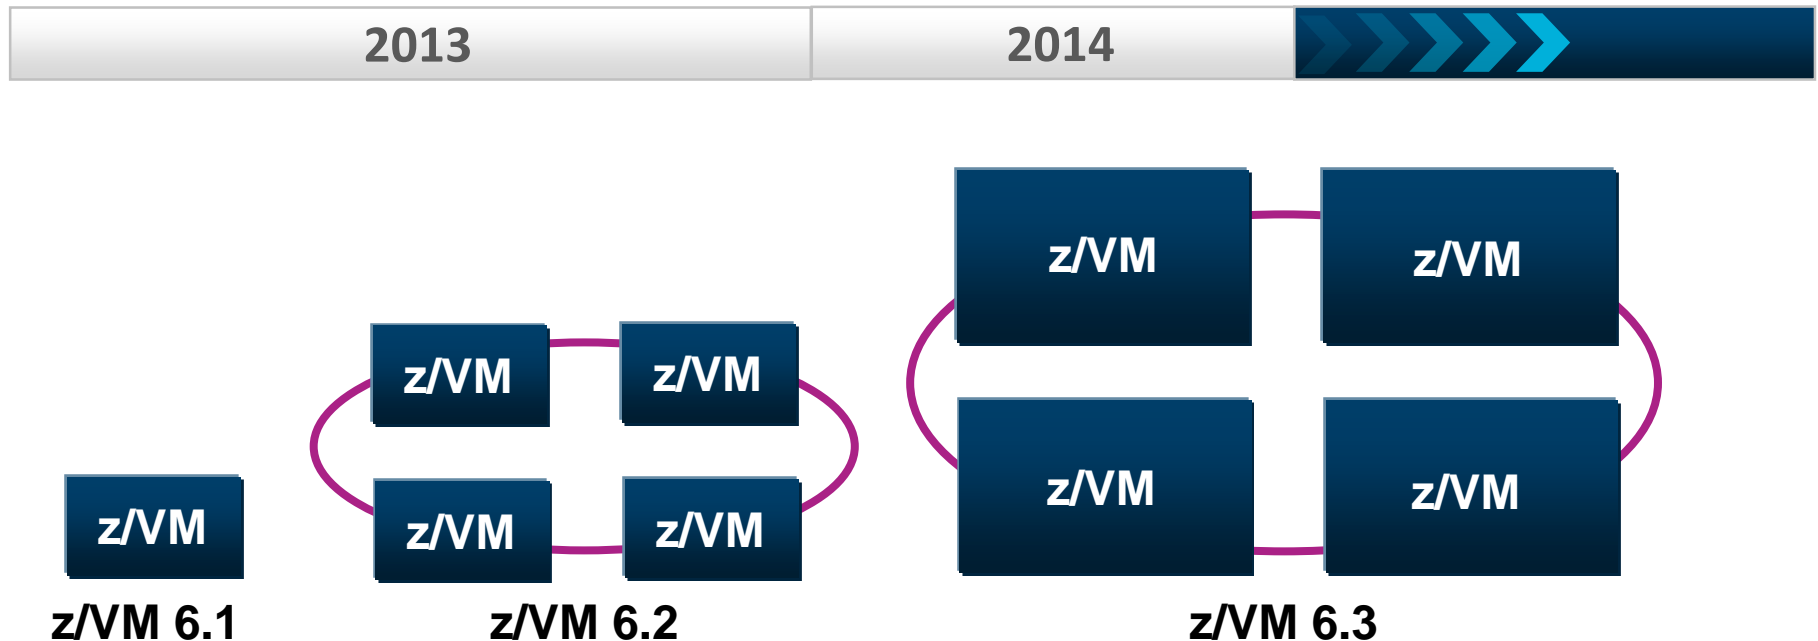
z/VM Version 6 Release 3

Making Room to Grow Your Business

Continued Evolution

- Improved TCO
- Scaling and efficiency improvements
- Ease of usability

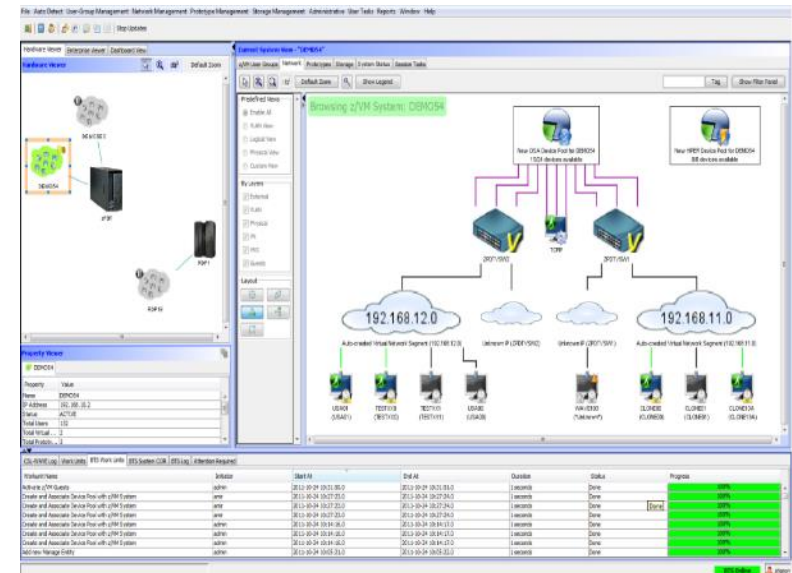
See <http://www.vm.ibm.com/zvm630/>



IBM Wave for z/VM (IBM Wave)

IBM Wave simplifies and helps automate management and administration of z/VM and Linux virtual servers, jumpstarting the steps needed to get to cloud. With its content rich interface IBM Wave extends the reach of your staff and lets you manage z/VM and Linux intuitively and cost effectively, reducing reliance on deep expert skills.

- Monitors and manages virtual servers and resources from a single interface
- Simplifies and automates administration and management tasks
- Provisions virtual resources (Guests, Network, Storage)
- Supports advanced z/VM capabilities such as Single System Image and Live Guest Relocation
- Allows delegation of administrative capabilities to the appropriate teams



A simple, intuitive graphical management, provisioning, and automation tool to help you fully leverage the power of System z virtualization on z/VM.



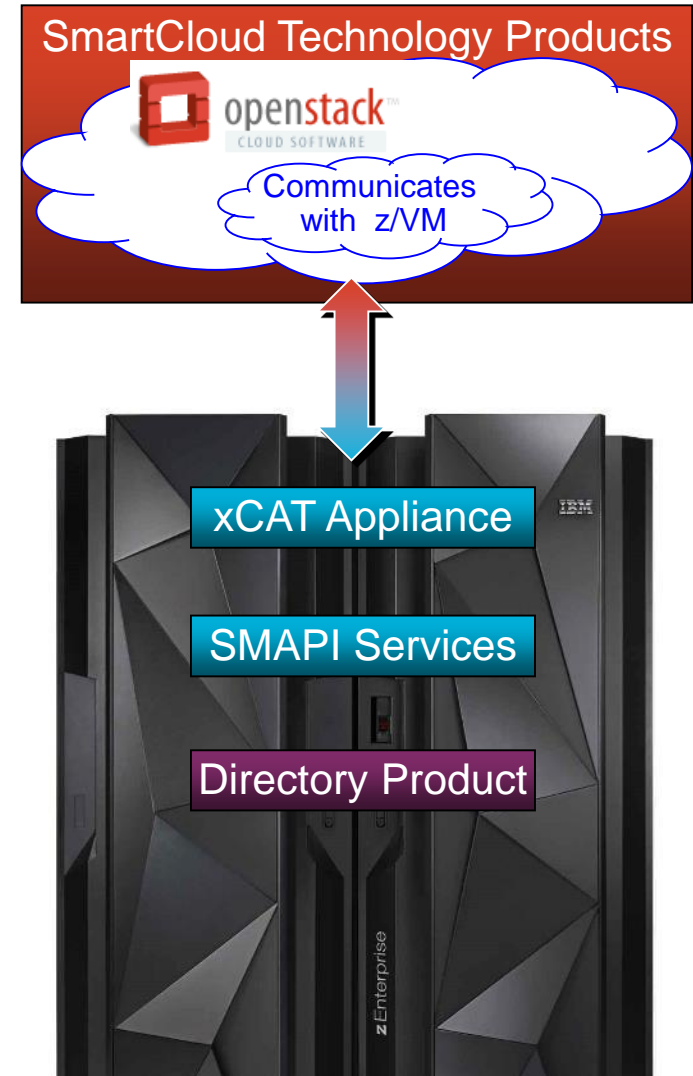
IBM Cloud Manager with OpenStack

Benefits:

- Full access to OpenStack APIs, backed with IBM support
- Cloud management solution that is easy to use
- Self service portal for workload provisioning and virtualized image management
- Heterogeneous support for IBM PowerVM®, z/VM, IBM PowerKVM and x86, and more
- Deploy, resize and capture
- Linux server backup and restore

Requires z/VM 6.3 with appropriate service

- *The xCAT Appliance utilizes new and existing Systems Management APIs (SMAPI) to interact with the z/VM system*




Linux Your Way

Distributions	Hypervisors	Languages	Runtimes	Management	Database	Analytics
 Supported by Canonical		python™ Ruby		docker CHEF™ JUJU	MariaDB mongoDB PostgreSQL	Spark
		php R ERLANG	RAILS ZF	openstack™ vmware vRealize Cloud Manager	cassandra CouchDB™ ORACLE™ Diamond Partner	hadoop IBM InfoSphere BigInsights
redhat.	LPAR	Scala Clojure JS OCaml Java	OpenJDK 	puppet labs	DB2	BLU Acceleration



IT Optimization with z Systems

Improved Effectiveness and Efficiency




- ✓ Operational and management reduction
- ✓ Software acquisition and licensing cost reduction
- ✓ Maximizing utilization
- ✓ Network reduction
- ✓ Collocation of data and applications
- ✓ Floor-space and energy reduction
- ✓ Growth inside a server
- ✓ Improving security
- ✓ Disaster recovery cost reduction



Linux on IBM z Systems



The real alternative to x86 server sprawl



A full room of servers **versus** One footprint with the size of a refrigerator

Unmatched Linux capacity

It's easy and cost-effective.

- Dundee City Council

Great degree of flexibility and scalability.

- Halkbank

Quickly and cost-effectively deploy innovative services.

- Banca Carige

Maintenance and support effort reduced by at least 65%.

- Algar Telecom

Operates even when resources are at 100% utilization.

- Bank of Tokyo-Mitsubishi UFJ

Differentiates in level of service and quality of service.

- L3C LLP



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



YOUR OPINION MATTERS!



Submit four or more session evaluations by 5:30pm Wednesday to be eligible for drawings!

*Winners will be notified Thursday morning. Prizes must be picked up at registration desk, during operating hours, by the conclusion of the event.



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.

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

