Wilhelm Mild – IT Architect
10 / 2013

Effectively running Linux on IBM System z in a virtualized environment and cloud

Enterprise2013

Trademarks

- This presentation contains trade-marked IBM products and technologies. Refer to the following Web site:

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

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads 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).

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 SEs 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.

Why Linux on and System z is a great Linux

- Linux is Linux - but inherits System z – characteristics like extreme virtualization
- System z the largest scalable server
- Reliability and availability of System z
- High Availability of System z and zEnterprise - out of the box
- System z servers can virtualize everything with up to 100% utilization rates

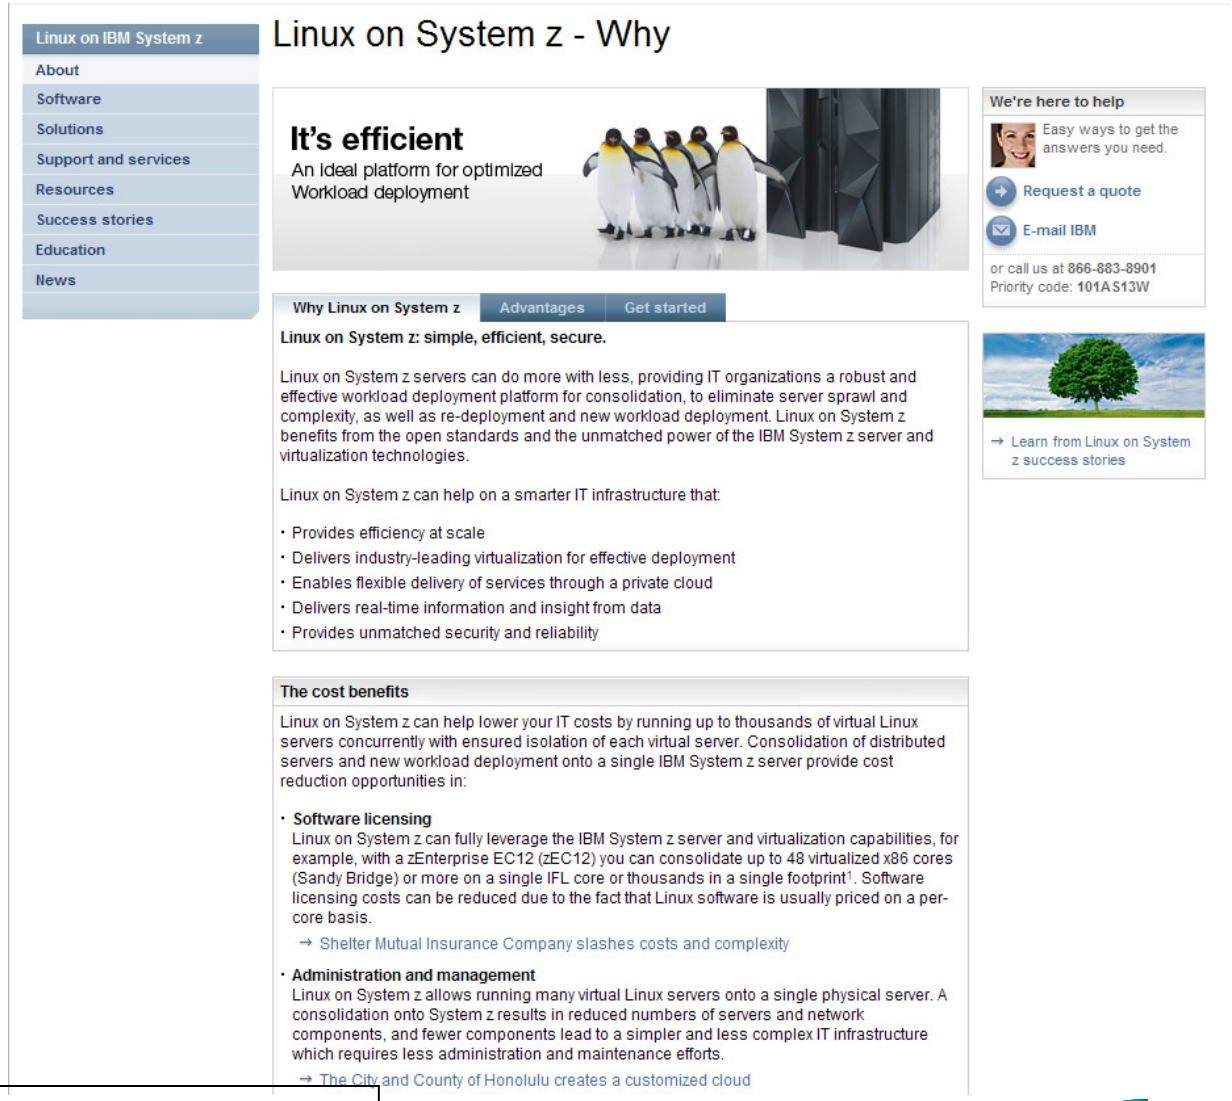

http://www-03.ibm.com/systems/z/os/linux/about/index.html

Virtualized environment on Linux - but what's best

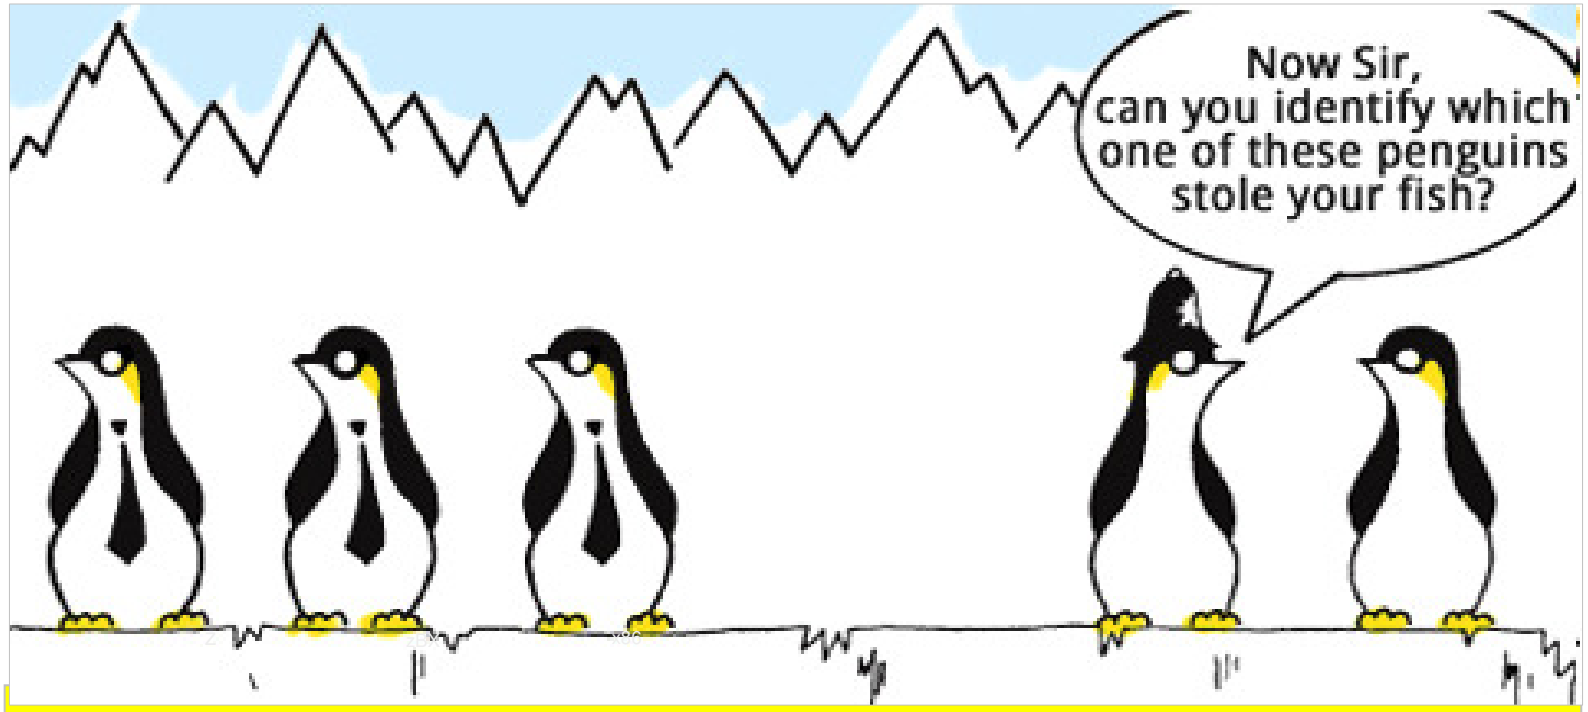

The **Linux's all look the same** (on different architectures) and have the same Linux kernel source.

But they have **different personalities**, **qualities**, **features** and **options** derived from the architectures.

Linux on System z has a Continuous Focus on Characteristics the Workload Benefits from

Security Capabilities:

- Centralized Authentication,
- Cryptographic Acceleration,
- Regulatory requirements, Identity management, Common Criteria Certification, Image Isolation,
- Physically secure communications with HiperSockets™ and Guest LANs

Consolidation Capabilities:

- Server, Network, Storage, Applications, hosting of different workloads at the same time

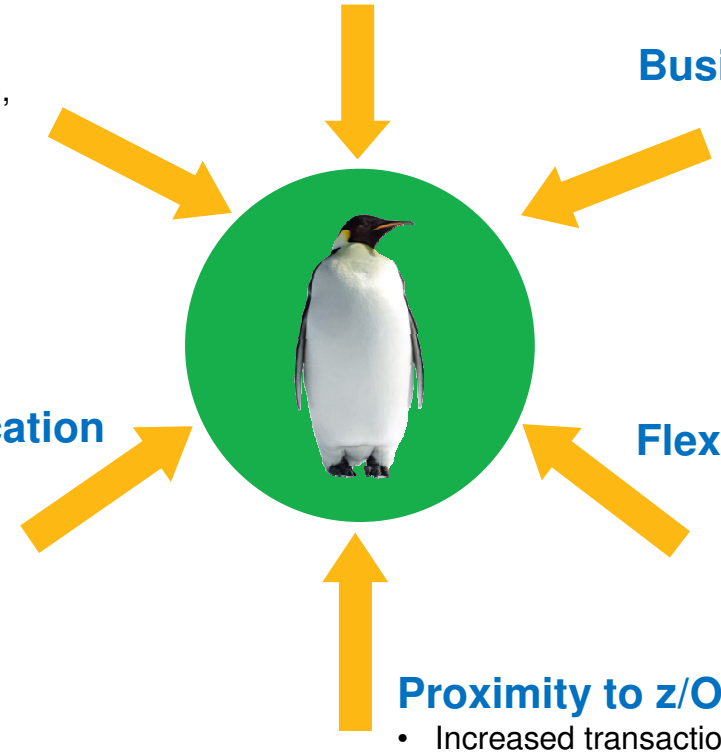

Business Resiliency Capabilities:

- High Availability,
- Disaster Recovery, Serviceability, Reliability,
- Storage failover (HyperSwap™), Data replication (XRC, PPRC)

Operational Simplification Capabilities:

- Virtualization,
- Resource Sharing
- Single Point of Control,
- Single System Image,

Flexibility / On demand Capabilities:

- Mixed Workloads: Scale-up & scale-out,
- Rapid server (de)commissioning,
- Idle Servers don't consume resources

Proximity to z/OS managed Data:

- Increased transaction throughput, HiperSockets
- Shared data access
- Integrated storage management

What is Virtualization and why is it important ?

Virtualization is the logical representation of resources not constrained by physical limitations

- Enables user flexibility
- Centrally manage many resources as one
- Dynamically change and adjust across the infrastructure
- Create many virtual resources within single physical device
- Eliminates trapped capacities

IBM Virtualization

A comprehensive platform to help virtualize the IT infrastructure

Terminology – 'Virtualization'

- If it's there and you can see it - it's real.
- If it's there and you can't see it - it's transparent.
- If it's not there and you can see it - it's virtual.
- If it's not there and you can't see it - you erased it.

Accumulated USED Distributed Server capacity

RPE2

(Relative Performance Estimate

from Ideas International™)

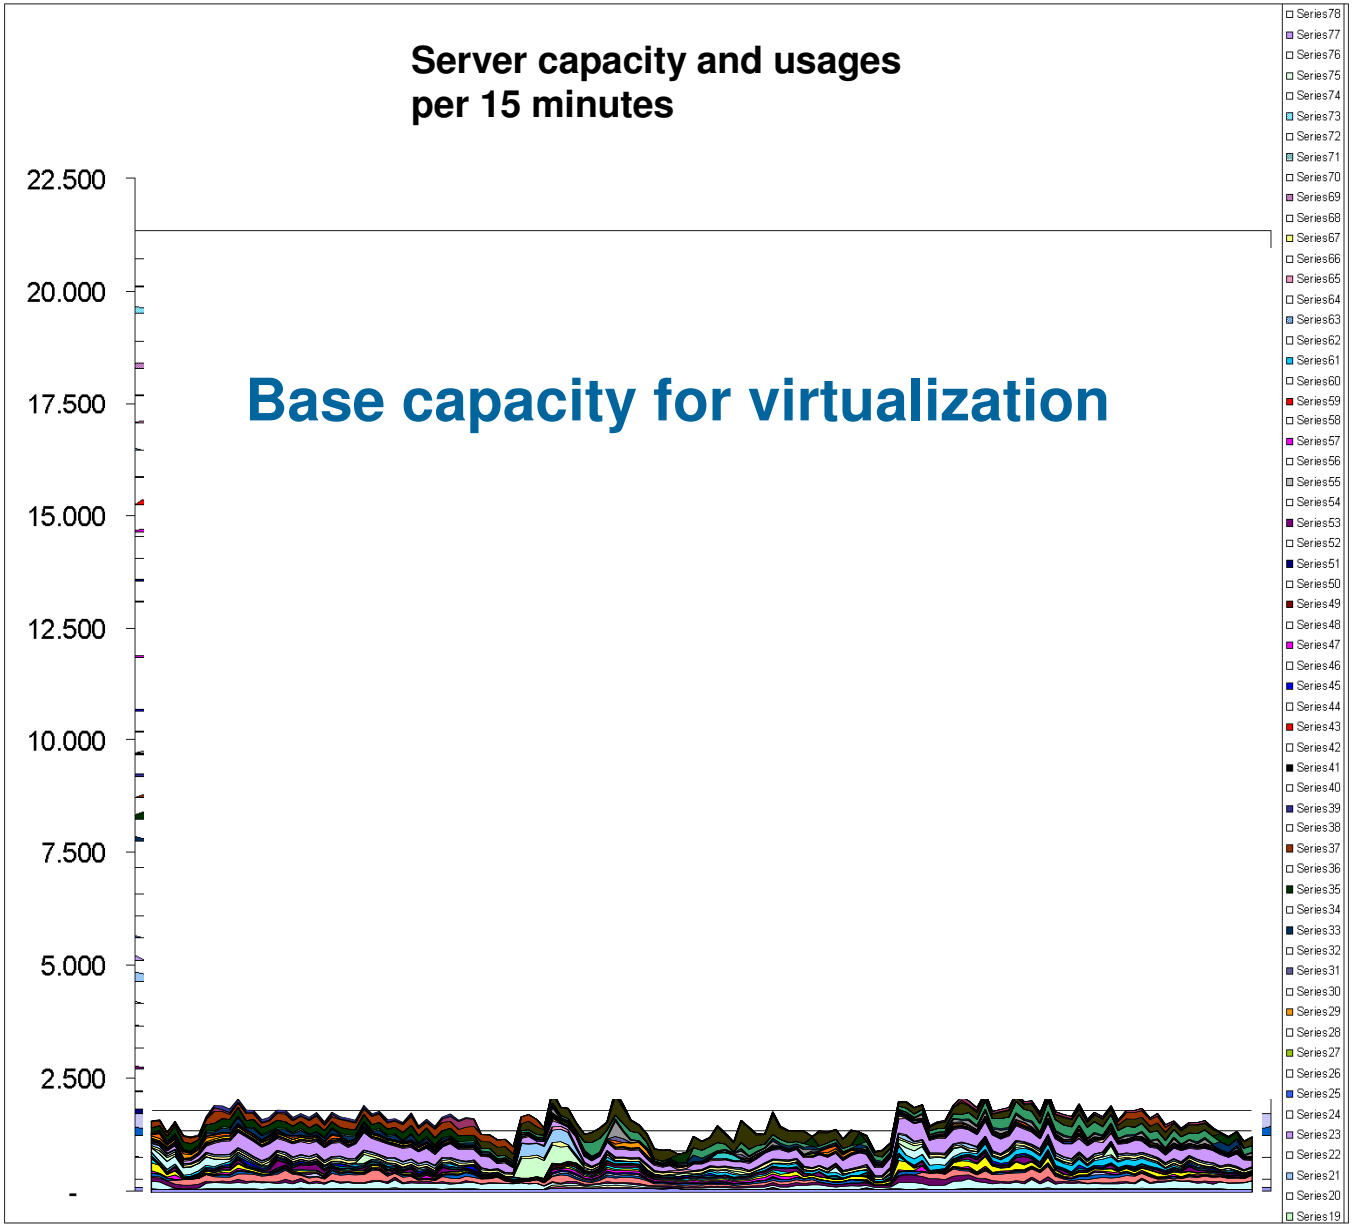

Global Virtualization in System z and zEnterprise
z/VM Technology: Share everything

- z/VM simulates the existence of a dedicated real machine, including processor functions, storage, and input/output resources.
- z/VM includes network Virtualization, high availability and integrated security between VMs
- It supports uniquely, **over commitment on all levels**.
- z/VM Virtualizes in a single virtualization Layer different workloads

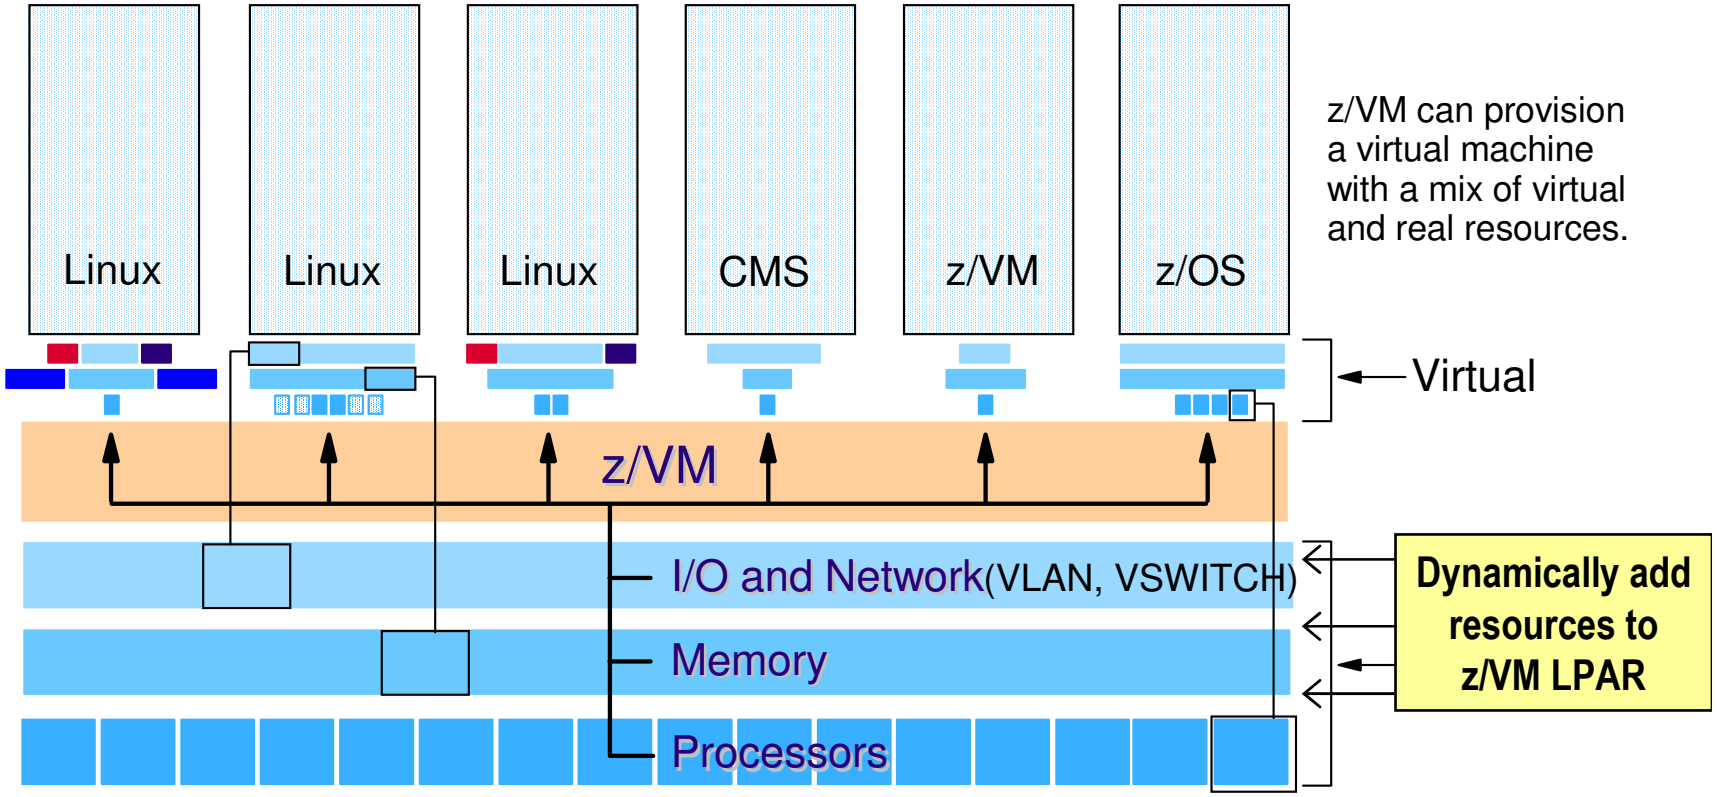

Linux on z/VM is the industry's most advanced virtual solution

A Brief History of IBM and Virtualization

IBM has 40 years of experience in virtualizing our servers. Virtualization was originally developed to make better use of critical hardware. Hardware support for virtualization has been critical to its adoption.

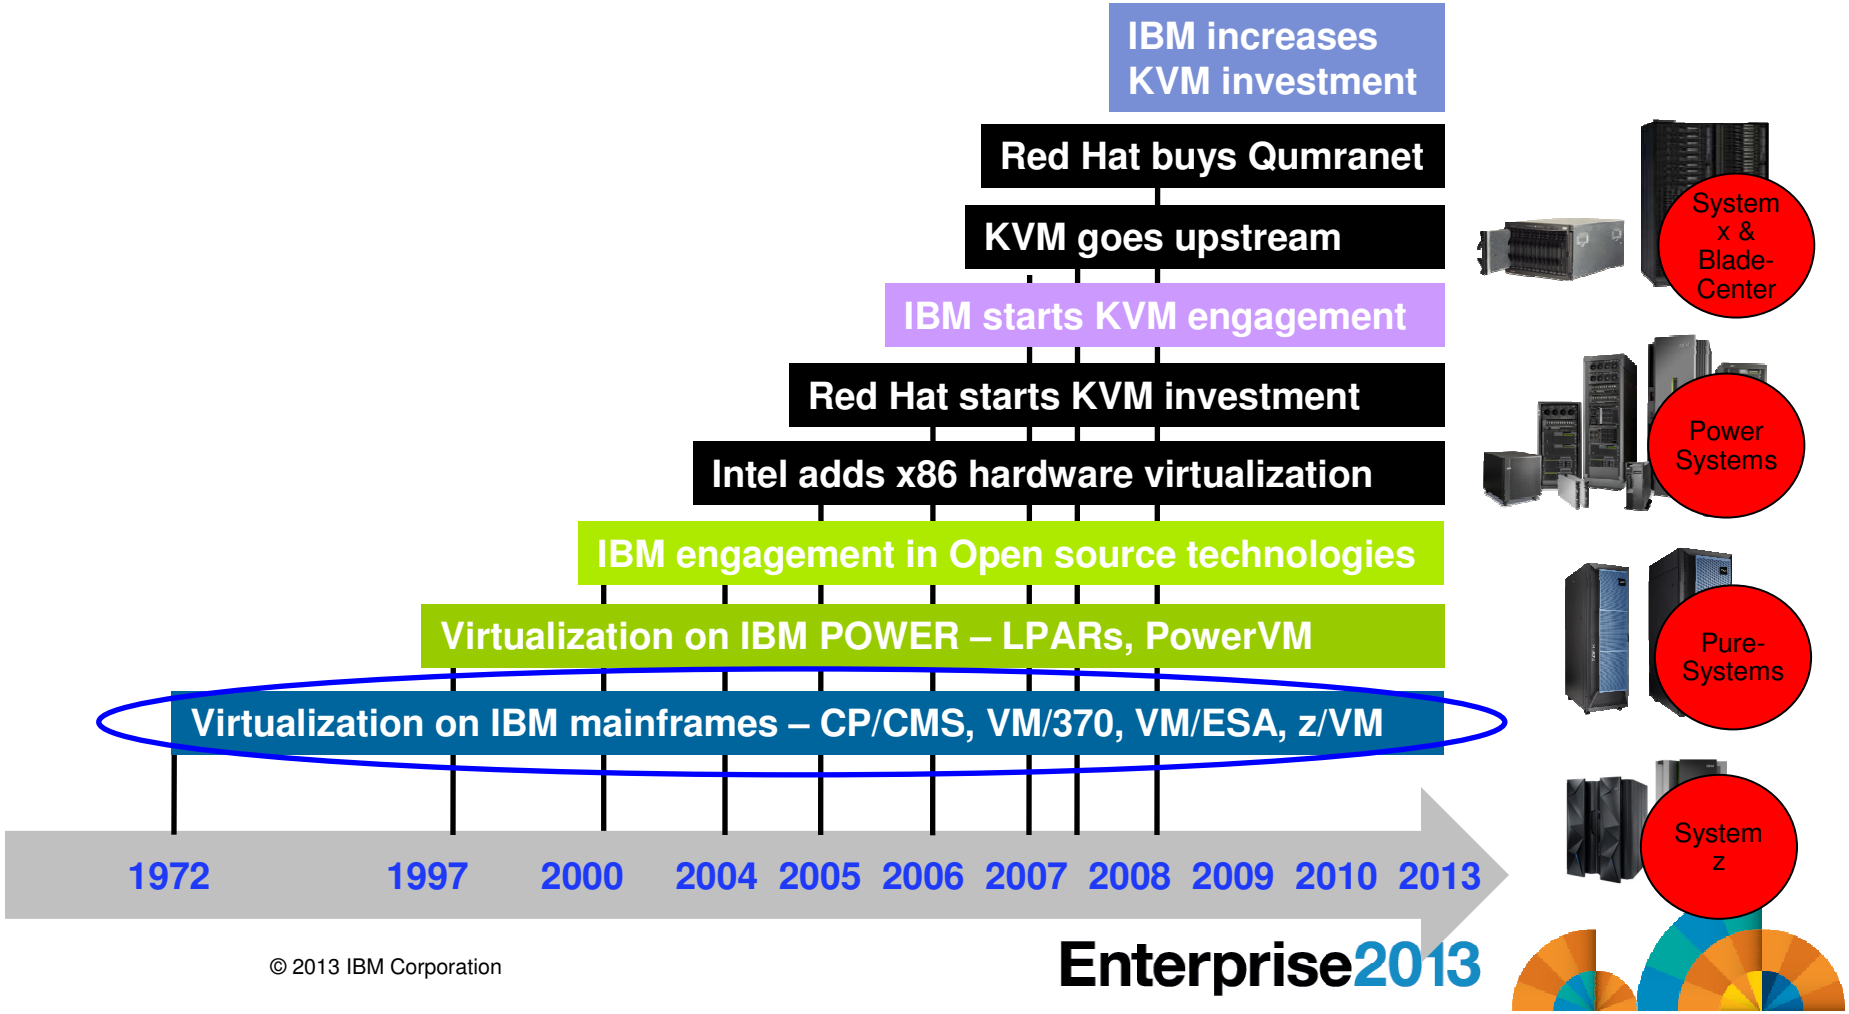

System z Multidimensional Virtualization Technology

Build-in and Shared Everything Architecture

Start Interpretive Execution (SIE)

Hardware assisted virtualization

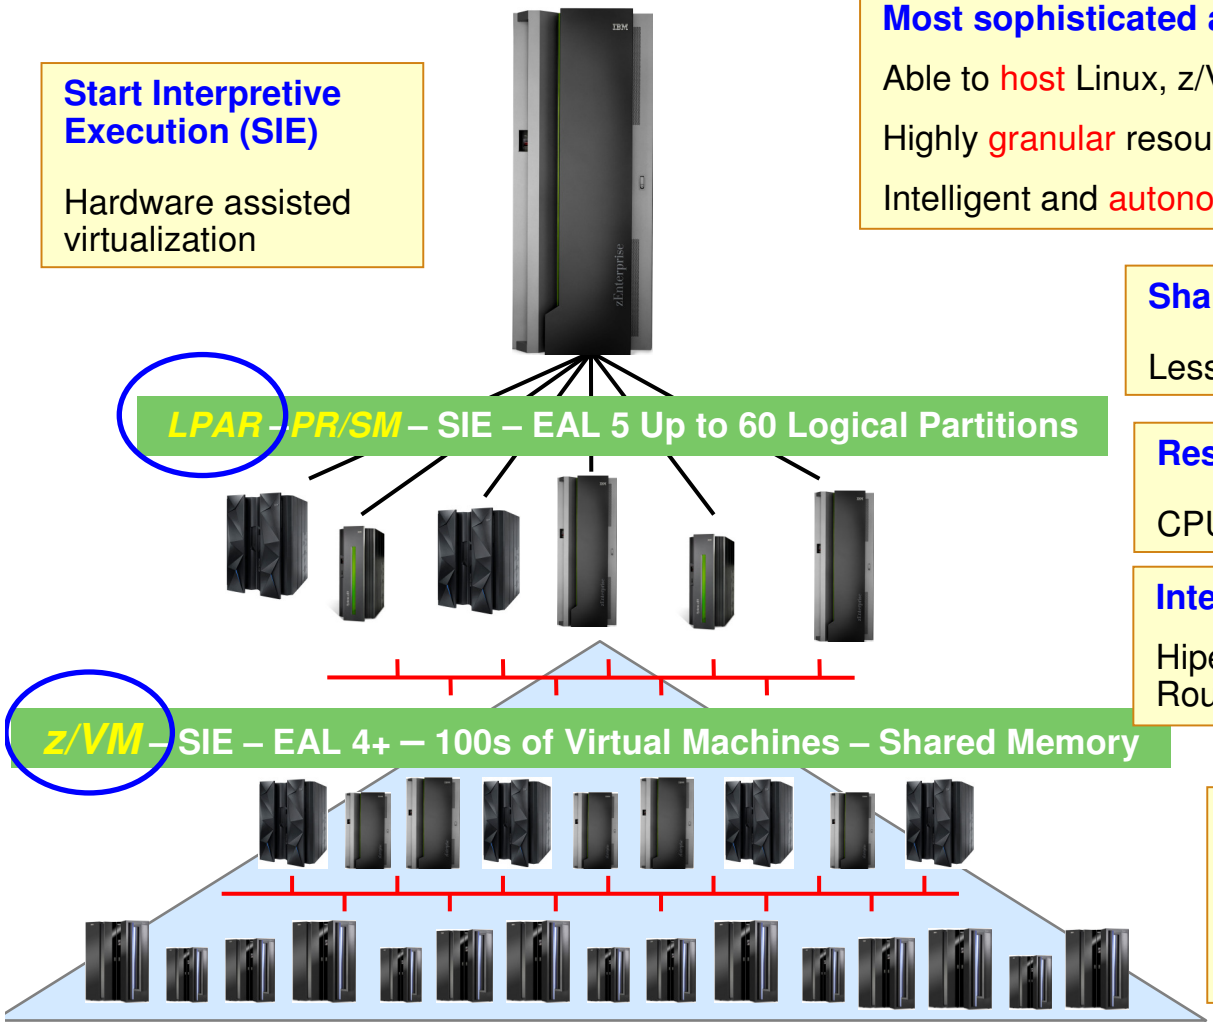

Most sophisticated and functionally complete hypervisors

Able to host Linux, z/VM and z/VM-on-z/VM, z/OS, z/VSE, zTPF

Highly granular resource sharing (less than 1% utilization)

Intelligent and autonomic workload management

Shared executables and filesystems

Less administration and reduced memory

Resources can be over-committed

CPU: up to 3:1 Memory 2 to 4:1

Internal high speed TCP/IP communication

Hipersocket, VLAN, Virtual Switches, Virtual Routers, Virtual Firewalls, CQ

HW (LPAR) and SW (z/VM) hypervisors

Linux virtual servers runs on bare metal but under control of hypervisor.

I/O is transparent to hypervisor

Maximum security and capacity isolation between virtual Linux servers

Effective via System z Virtualization

Note: There are typically dozens or hundreds of Linux servers in a z/VM LPAR.

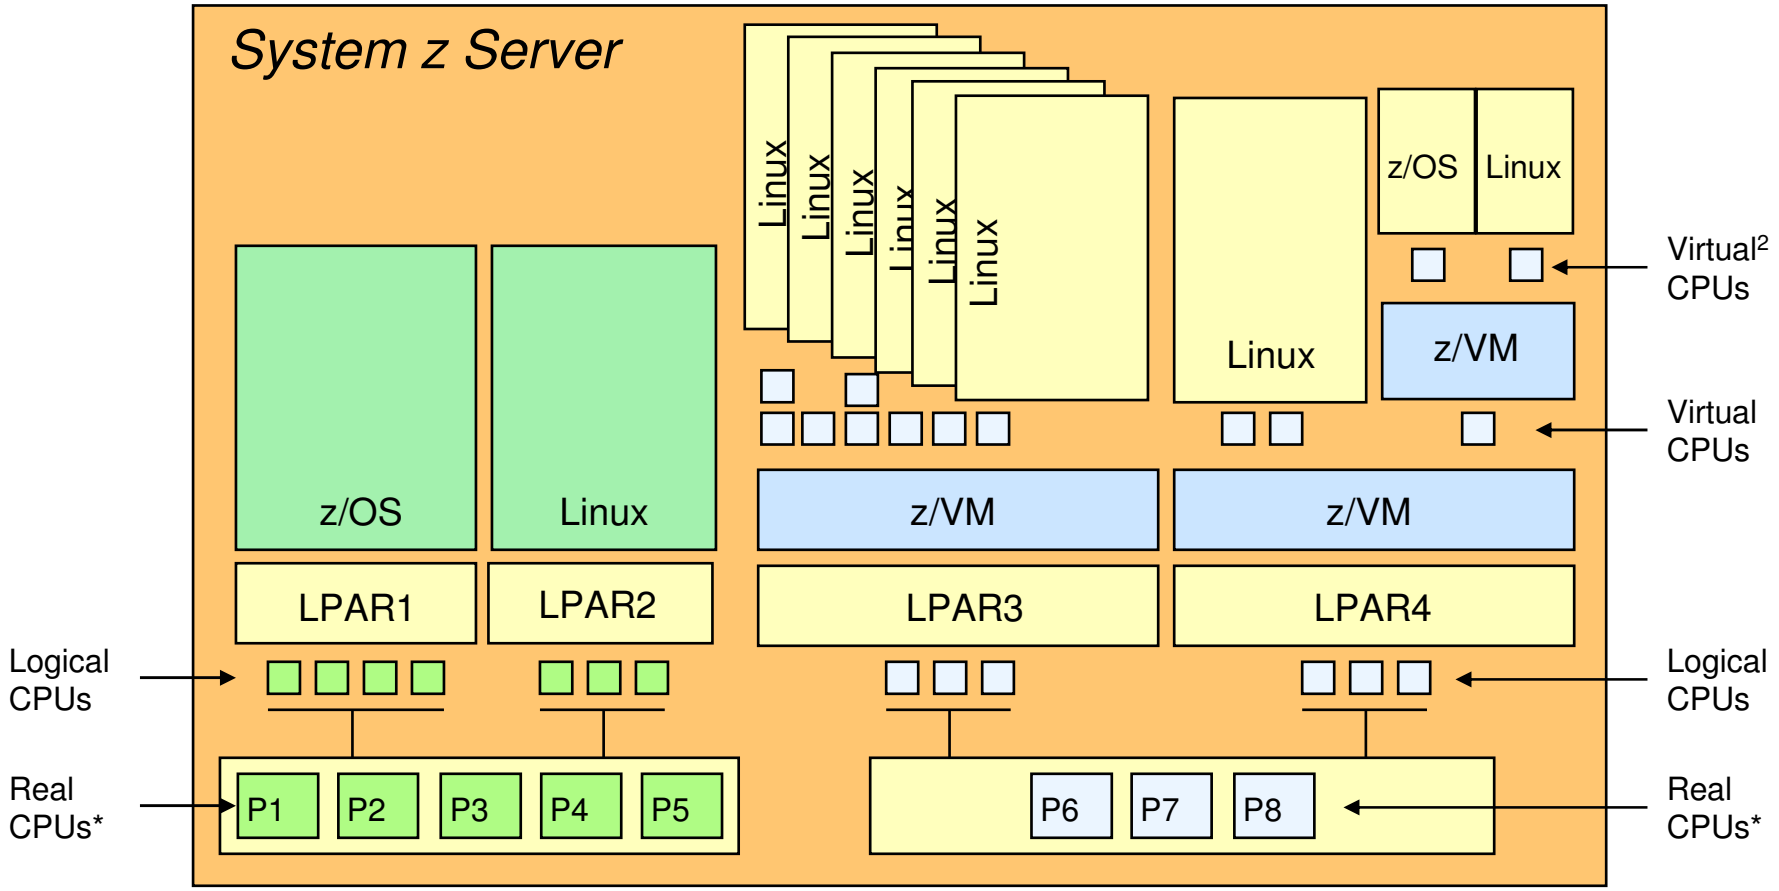

P1 – P8 are Central Processors (CP) or Integrated Facility for Linux (IFL) Processors

* - One shared Pool of CPUs per machine (CEC) only

IBM Virtualization on System z, with Linux on System z -Scalable to thousands of Linux guests

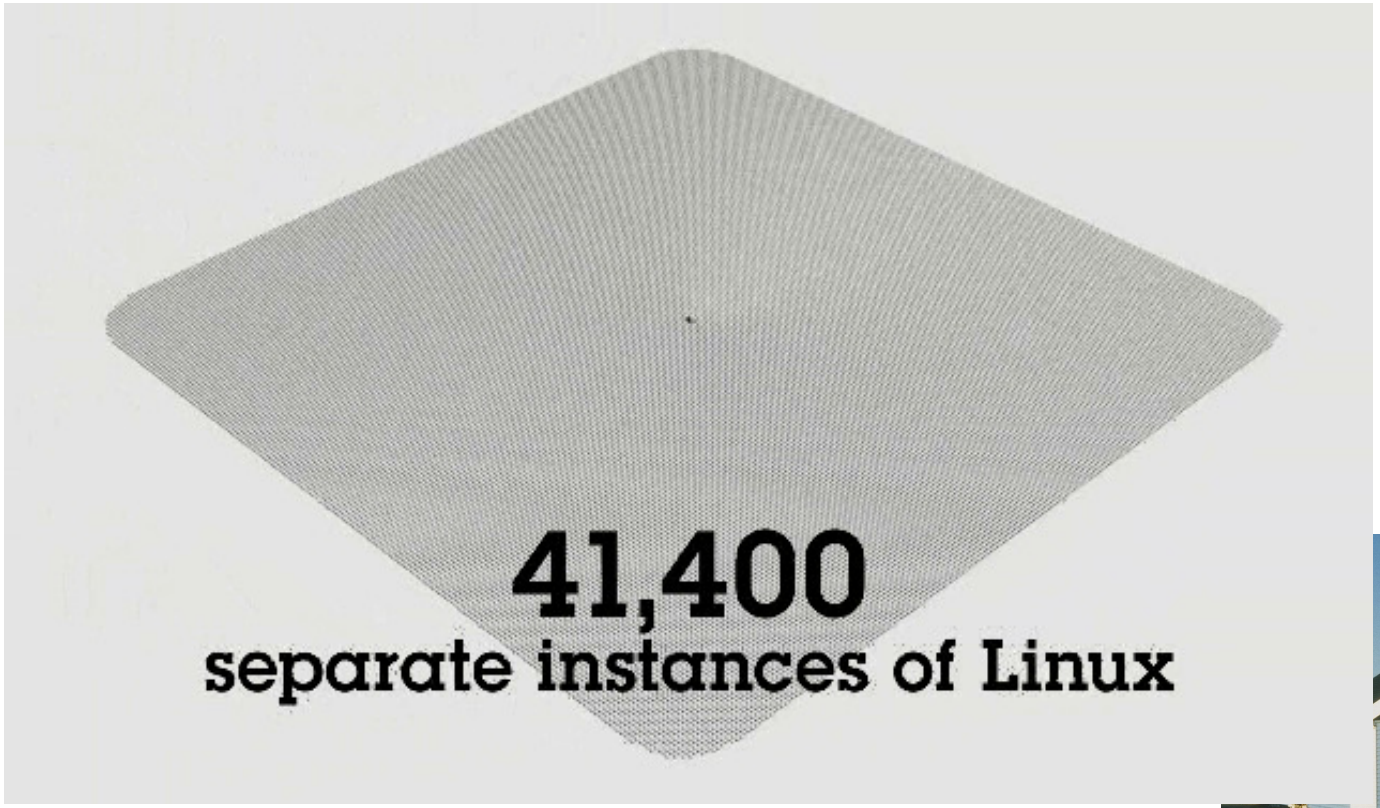

Increased flexibility in virtualization

Single System Image (SSI) Feature
Clustered Hypervisor with Live Guest Relocation

- Single System Image (SSI)
  - Connect up to four z/VM systems as members of a Single System Image cluster
  - Cluster members can be run on the **same or different System z servers**
- Simplifies management of a multi-z/VM environment
  - Single user management
  - Cluster management from any member
    - Apply maintenance to all members in the cluster from one location
    - Issue commands from one member to operate on another
- Live Guest Relocation (LGR)
  - Non-disruptively move Linux guests from one z/VM member to another
  - Reduce planned outages
  - Enhance workload management

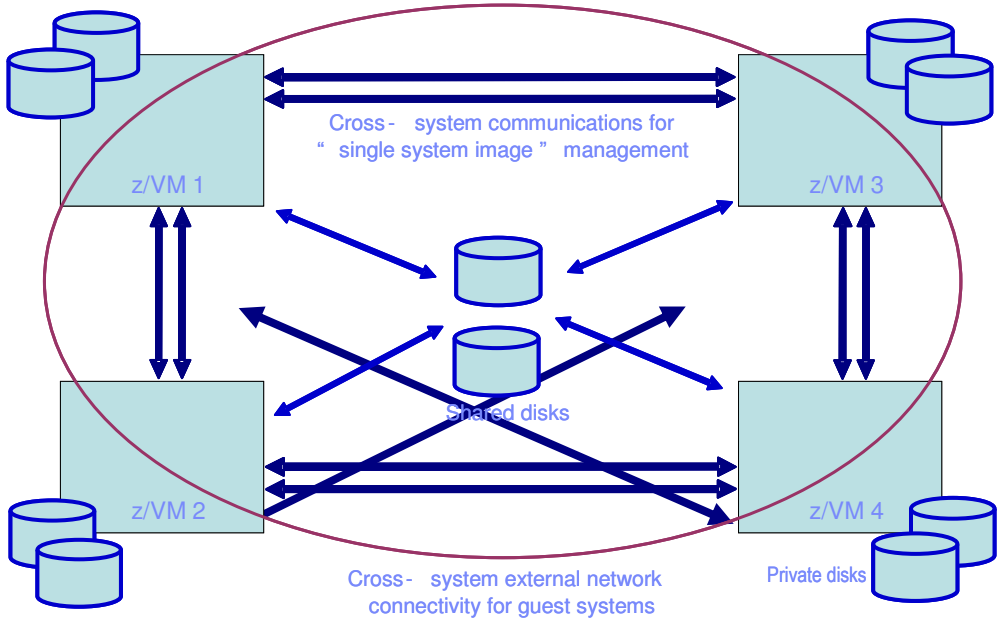

Client feedback from IBM Systems Directions Study

Consolidation and Virtualization together increase the effectiveness

Reasons to run Workloads on Linux on System z

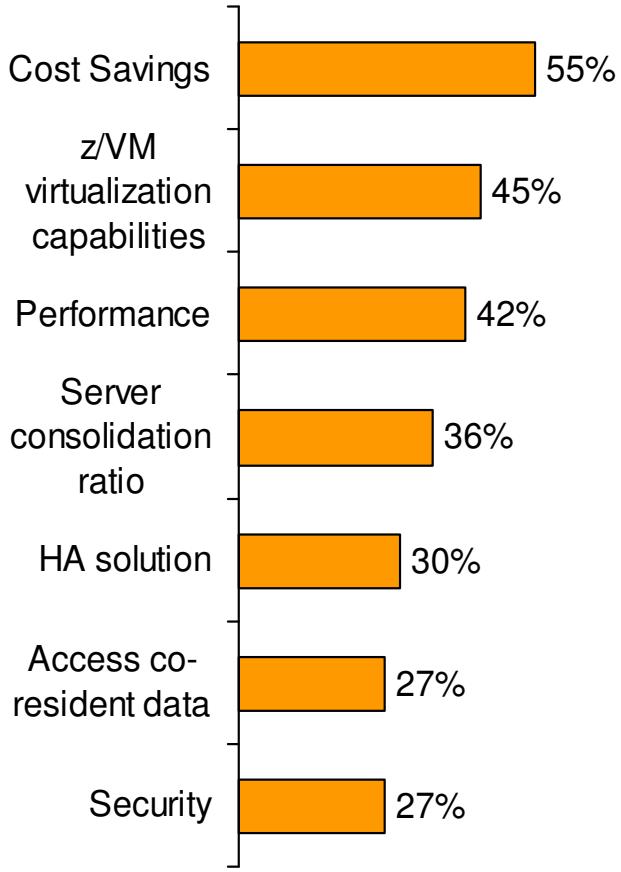

Previous Platform for Consolidated Workloads on Linux on System z

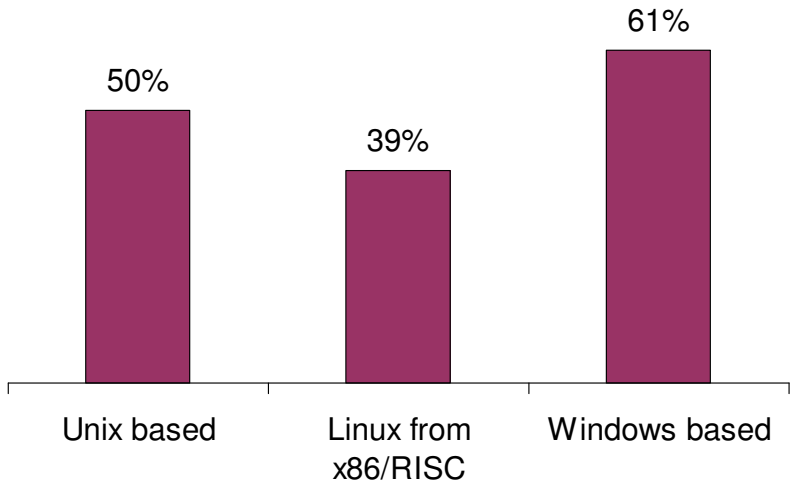

Source: March 2012 IBM Market Intelligence, Percentage of survey respondents

Effectively running Linux on z virtualized with z/VM

z/VM Multi-zone Network VSWITCH (red - physical isolation)

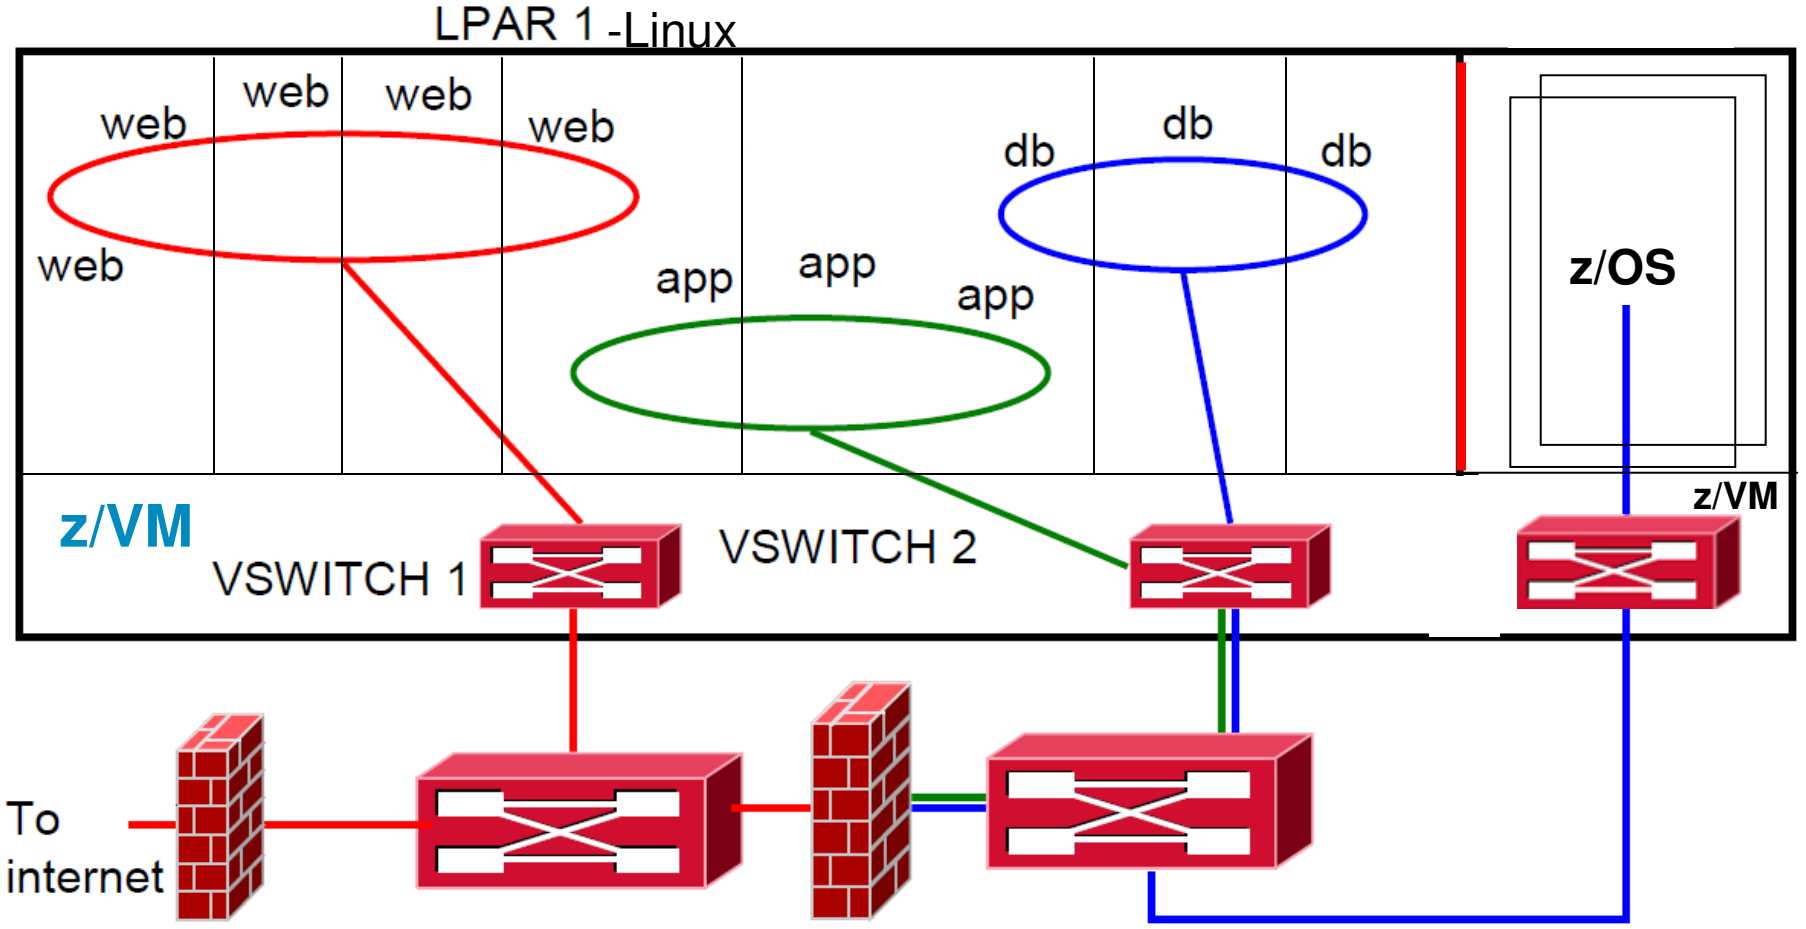

With 2 VSWITCHes, 3 VLANs, and a multi-domain firewall

Network Virtualization - in z/VM

- z/VM virtualizes network connectivity
  - VLAN
  - VSWITCH
- Virtual connectivity uses memory to memory connections
  - Very high bandwidth
  - Low latency
  - Internal – no physical devices

Guest LAN

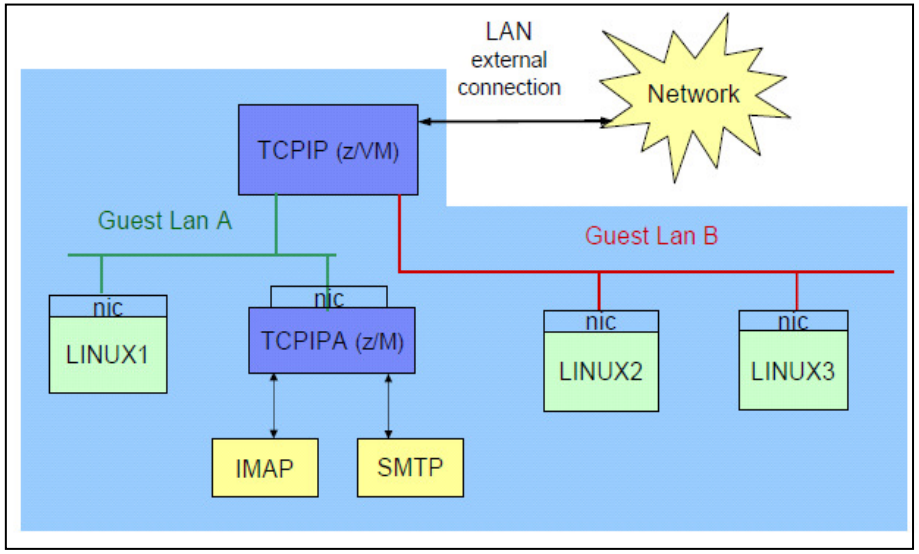

Virtual SWITCH

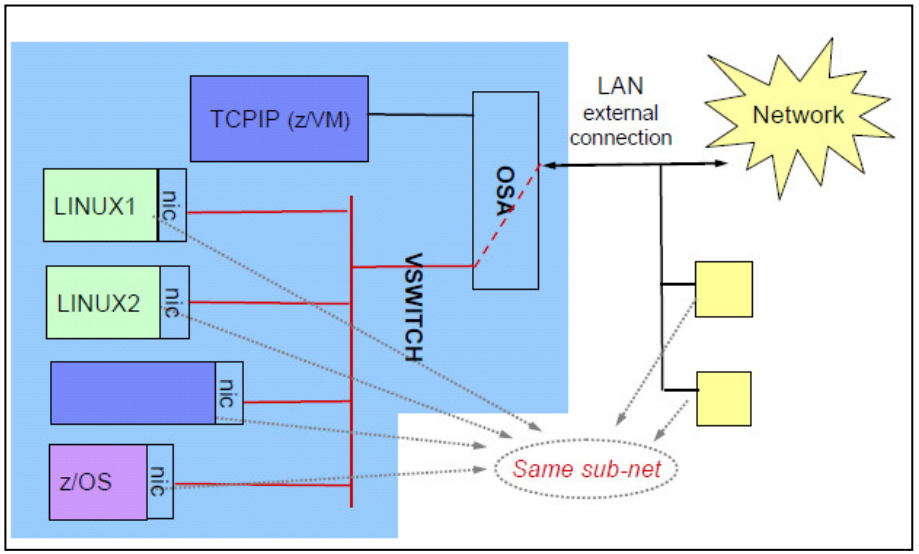

System z virtual network bandwidth enhancement and Automated Failover

Resource Virtualization:

OSA Channel Bonding in Linux on z

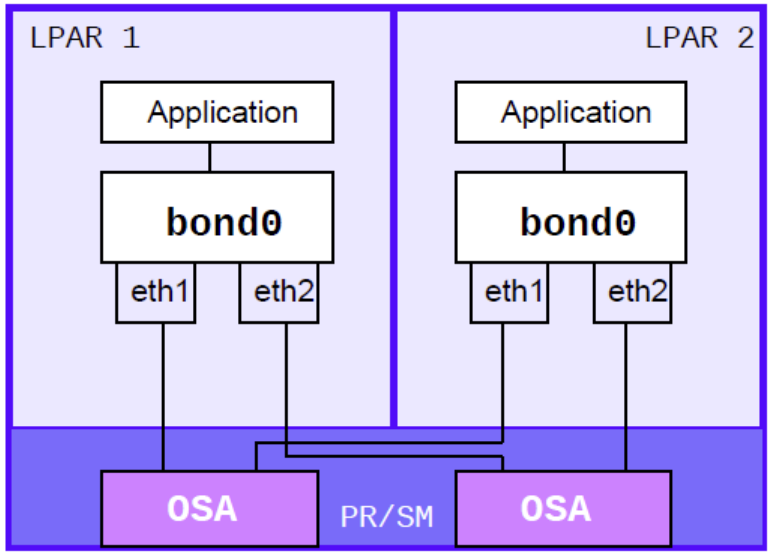

- Linux *bonding* driver enslaves multiple OSA connections to create a single logical network interface card (NIC)
- Detects loss of NIC connectivity and automatically fails over to surviving NIC
- Active/backup & aggregation modes
- **Separately configured for each Linux**

Network Virtualization:

z/VM Port aggregation

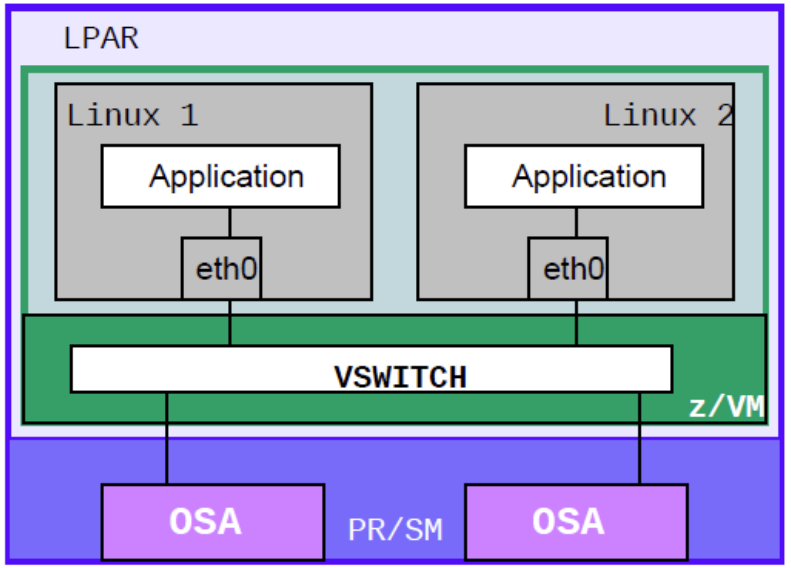

- z/VM *VSWITCH* enslaves multiple OSA connections. Creates virtual NICs for each Linux guest
- Detects loss of physical NIC connectivity and automatically fails over to surviving NIC
- Active/backup & aggregation modes
- **Centralized configuration benefits all guests**

Insurance Company Consolidated 292 Servers to one System z

Before

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

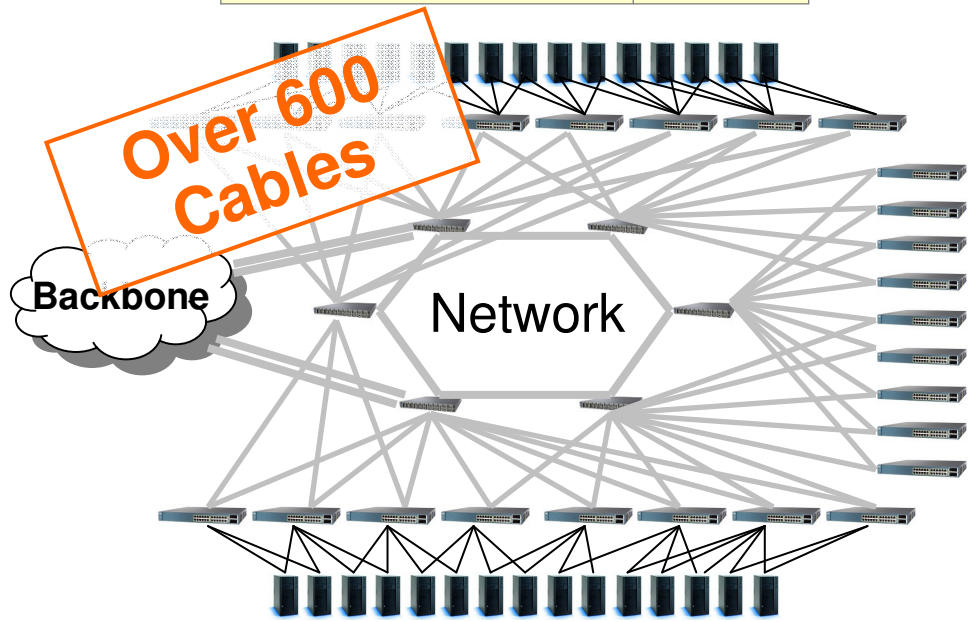

The diagram only shows **30** of **292** servers

After

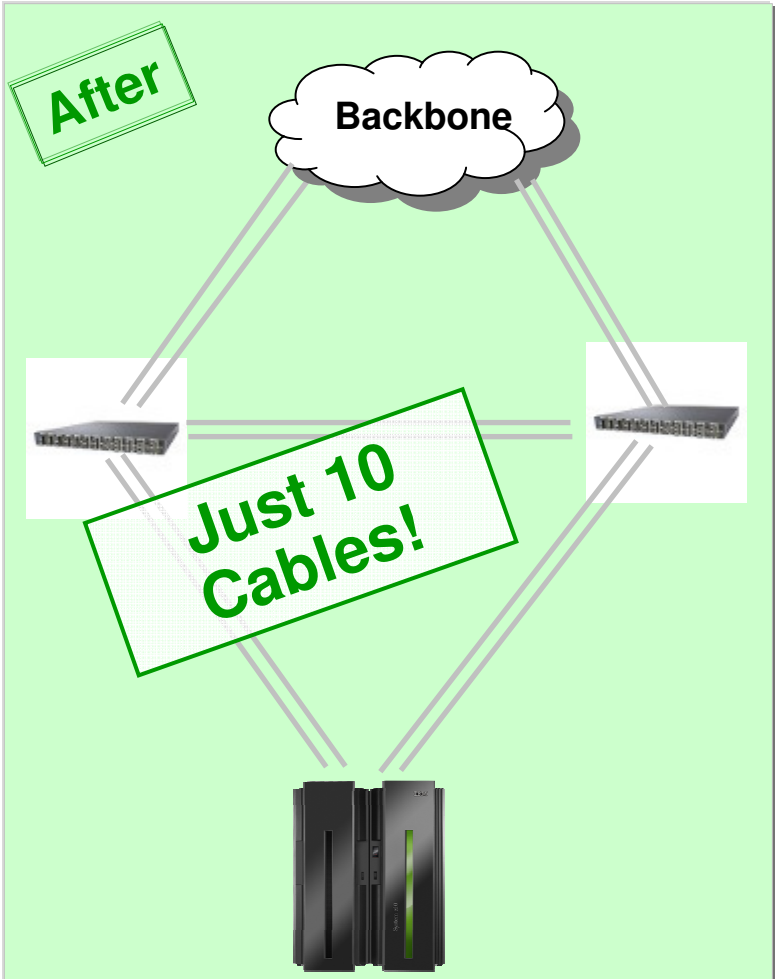

Data is based on real client opportunity and on internal standardized costing tools and methodologies.
Client results will vary by types of workloads, technology level of consolidated servers, utilization factor, and other implementation requirements. Savings will vary by client.

z/VM Version 6 Release 3
HiperDispatch: Dispatching Affinity

- Processor cache structures become increasingly complex and critical to performance
- Re-dispatch virtual CPU near where its data may be in cache based on where the virtual CPU was last dispatched
- Keep virtual CPUs of the same virtual machine near one another

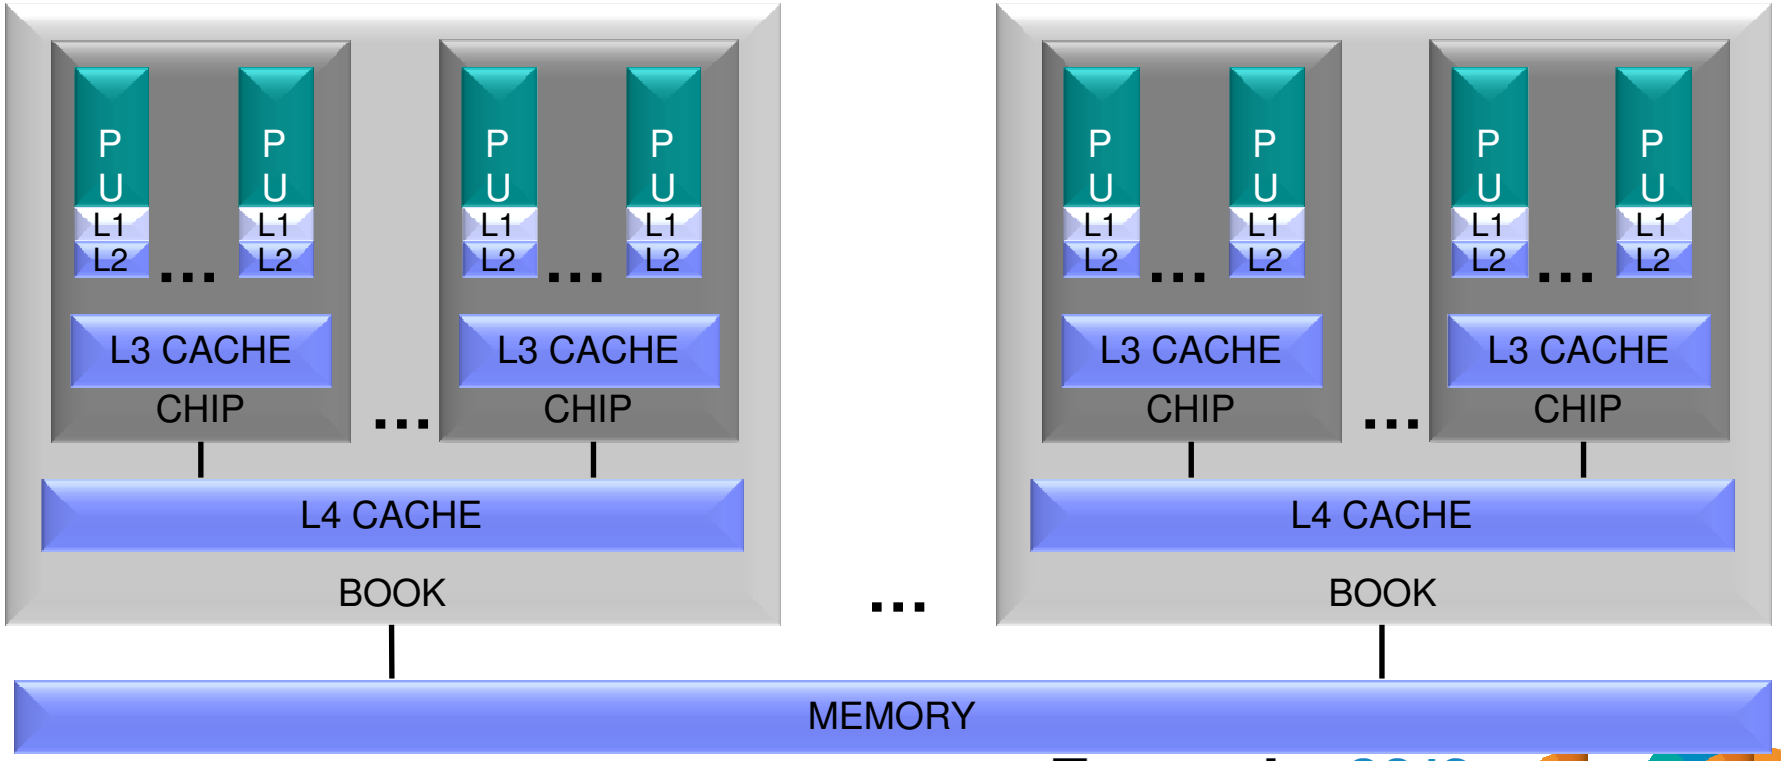

Horizontal vs. Vertical CPU Management

• Today's "horizontal" management distributes the LPAR weight evenly distributed across the logical processors of the z/VM LPAR

Horizontal:

•The logical processors are all created/treated equally.
•z/VM dispatches work evenly across the 7 logical processors

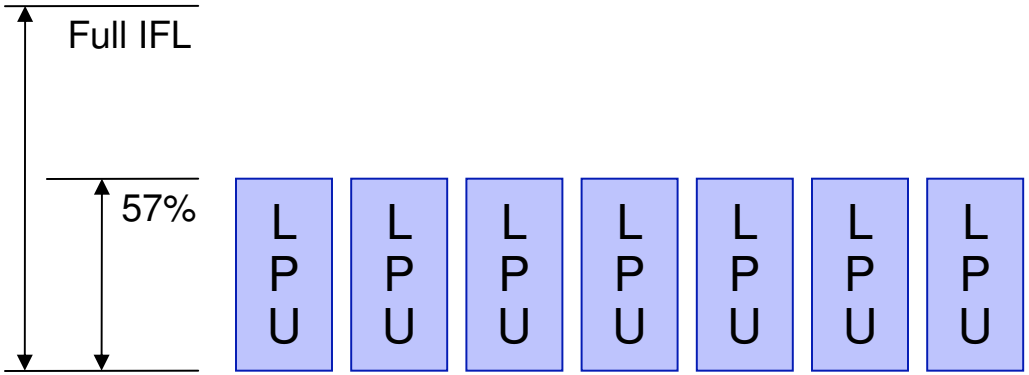

• "Vertical" management attempts to minimize the number of logical processors, allowing LPAR to similarly manage logical CPUs

Vertical:

•The logical processors are skewed to where some get greater share of the weight.
•z/VM dispatches work accordingly to the heavier weighted workload.

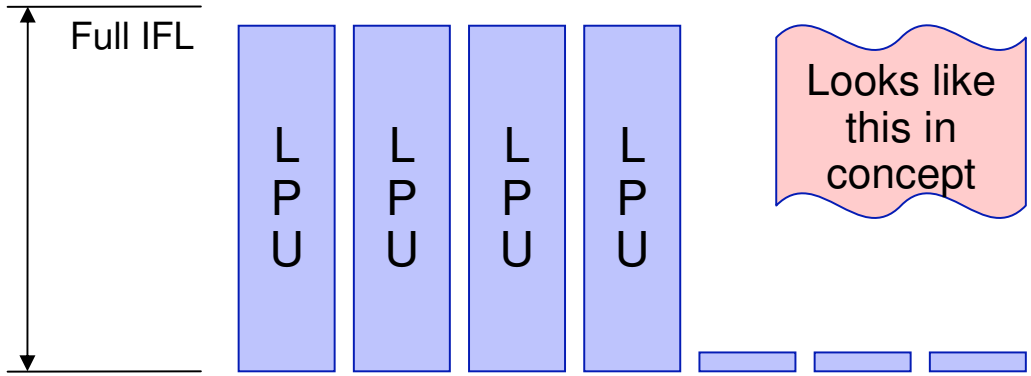

Linux interaction with the z/VM Hypervisor
CMSFS user space file system support

- Allows to mount a z/VM minidisk to a Linux mount point
- z/VM minidisk needs to be in the enhanced disk format (EDF)
- The cmsfs fuse file system transparently integrates the files on the minidisk into the Linux VFS, no special command required

```
# cmsfs-fuse /dev/dasde /mnt/cms
# ls -la /mnt/fuse/PROFILE.EXEC
-r--r----- 1 root root 3360 Jun 26  2009 /mnt/fuse/PROFILE.EXEC
```

- By default no conversion is performed
  - Mount with '-t' to get automatic EBCDIC to ASCII conversion

```
# cmsfs-fuse -t /dev/dasde /mnt/cms
```

- Rread & Write support available
- use “vi” to edit PROFILE.EXEC anyone ?
- Use fusermount to unmount the file system again

```
# fusermount -u /mnt/cms
```


Automated z/VM actions via Linux

Deliver z/VM CP special messages as uevent

11.2

6.1

- Allows to forward SMSG messages to user space programs
  - Message needs to start with “APP”
- The special messages cause uevents to be generated
- See “Writing udev rules for handling CP special messages” in the Linux on z Device Drivers Book

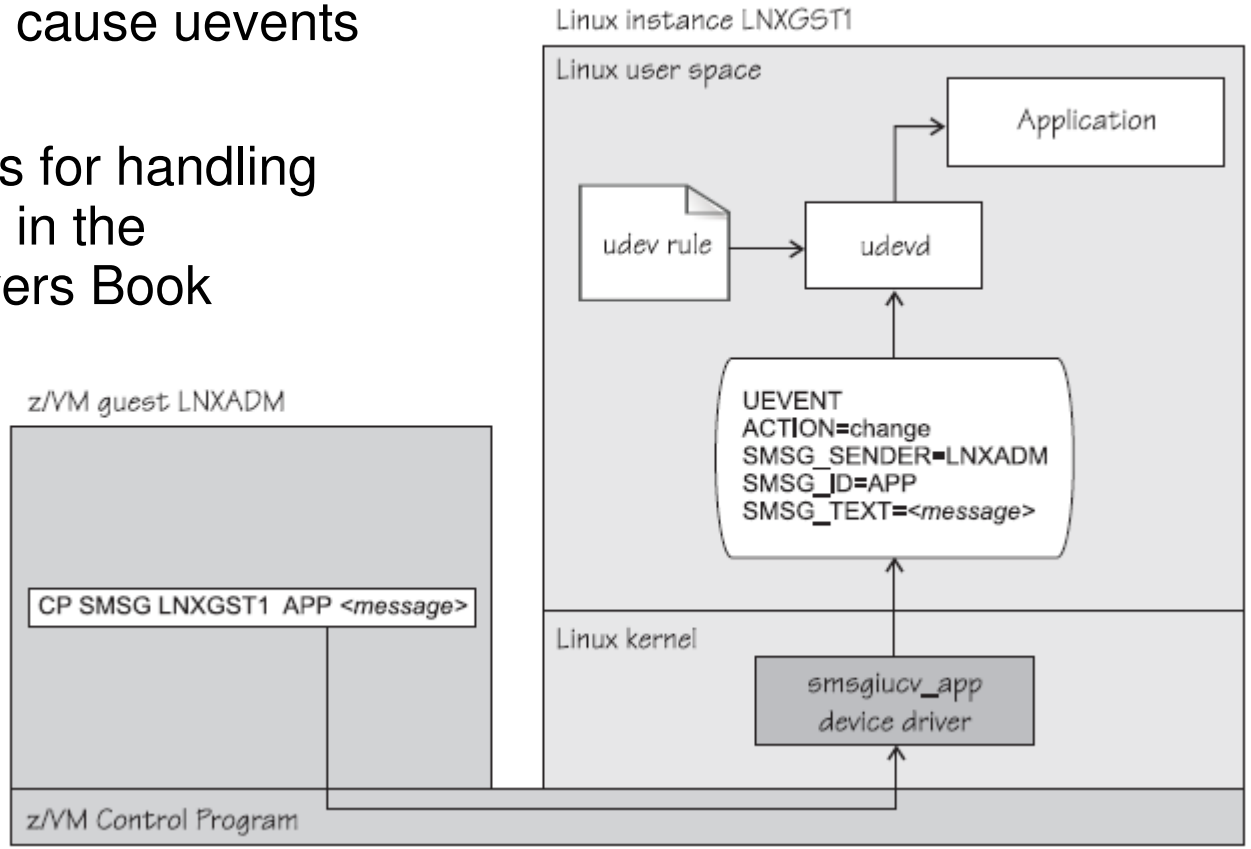

Manage z/VM resources within the range of the guest definition - Tool cpuplugd

- Sizing Linux z/VM guests made effective with **cpuplugd**
- Oversized guests often cause additional management effort by the Hypervisor and undersized guests often have performance-related issues with workload peaks
- A large amount of guests with large ratios of resource overcommitment (more virtual resources than are physically available) and changing workload characteristics over time makes a correct sizing even more challenging
- The cpuplugd daemon available with SUSE Linux Enterprise Server (SLES 11) SP2 or Red Hat Enterprise Linux (RHEL) 6.2, greatly enhances the capability:
  - to define rules
  - define the performance parameters for the rule set.
- This tool now enables the operating system of the guest to manage the resources within the range of the guest definition.
- The Linux cpuplugd daemon, can be used to automatically adjust CPU and memory resources of a Linux z/VM guest.
- The description ZSW03228-USEN-00 is available at the IBM Information Center
- http://publib.boulder.ibm.com/infocenter/lnxinfo/v3r0m0/index.jsp?topic=%2Fliaag%2Fl0cpup00_2012.htm

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, DB2, etc)**
- **The DCSS becomes a consistent unit of one software level**
- **NSS – Named Saved System – for a bootable Linux image**

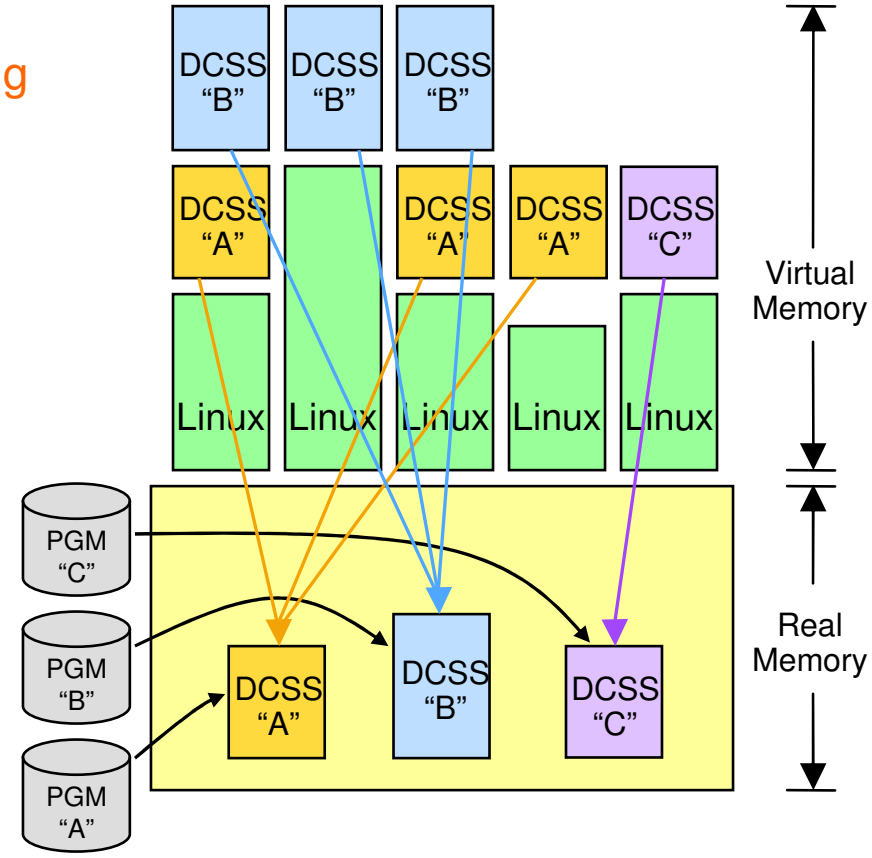

http://public.dhe.ibm.com/software/dw/linux390/perf/ZSW03186USEN.PDF

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

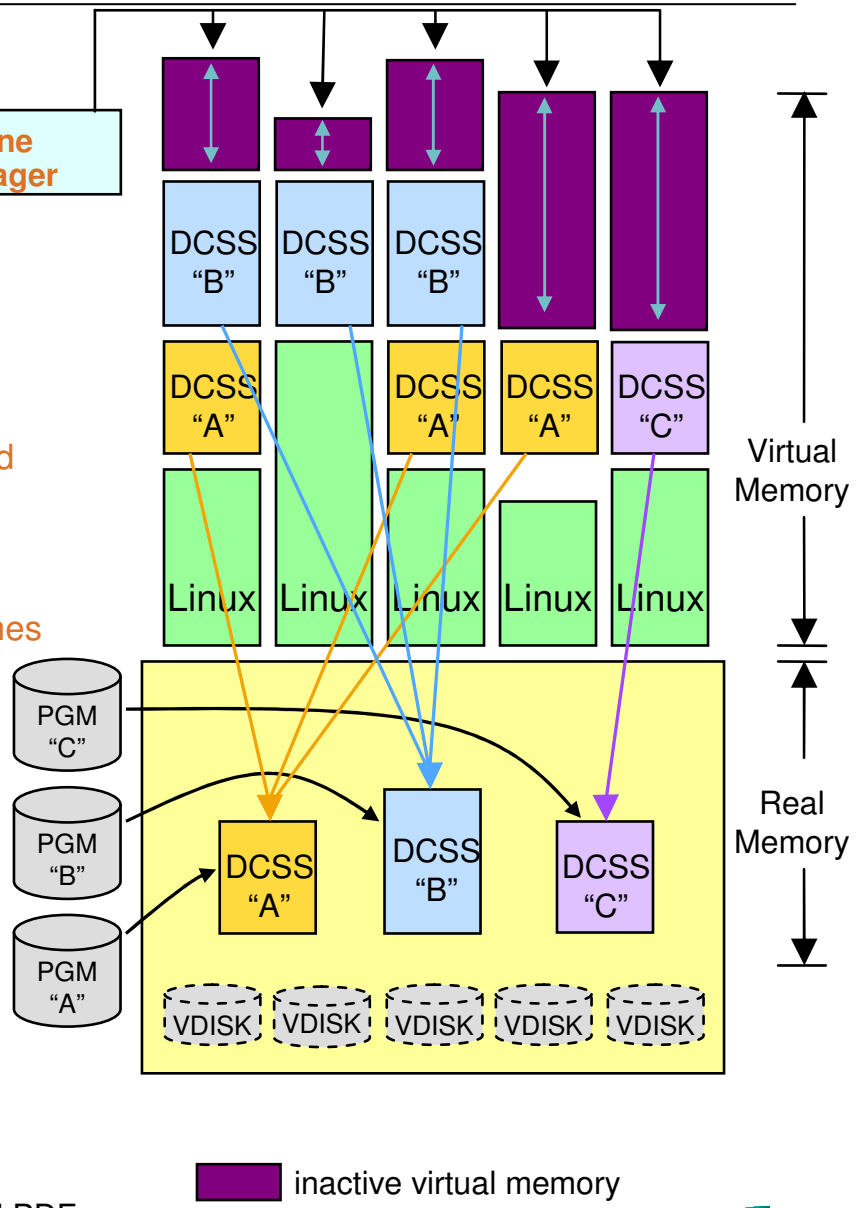

http://public.dhe.ibm.com/software/dw/linux390/perf/ZSW03186USEN.PDF

Memory overcommitment

- **Memory overcommitment,** often mentioned as major **benefit of virtualized environments**
  - Memory overcommitment is the ability to use more virtual memory as physically available memory. It is based on the assumption that not all guests use their memory at the same time and/or some guests are over-sized in their memory setup.
- The most **common mistake made with Linux guests under z/VM is over-configuring Linux** memory:
  - In a virtualized environment under z/VM, oversized guests place unnecessary stress on the VM paging subsystem:
    - Real memory is a shared resource, caching pages in a Linux guest reduces memory available to other Linux guests.
    - Larger virtual memory requires more kernel memory for address space management.
- Very different definitions exist for the **level of memory overcommitment**
  - they are independent of the used middle ware
  - "Common" System z ratio is from 1.5 : 1 to 3: 1 virtual (guest memory) to physical memory E.g. run guests defined with a total of 3GB on 1 GB real memory
  - Performance can be heavily degraded when memory overcommitment level is to high
- Identify "rules"
  - Determine the minimum amount of physical memory required to run with an acceptable performance
  - Identify a dependency on the used middle ware / applications

Test Scenario

- Test environment
  - Running a mix of server types as Linux guests on z/VM:
    LPAR with 28GB central storage + 2 GB expanded storage

Guest workload	Guest Memory
WebSphere Application Server	**13.5 GB (Java heaps 8GB)**
Database DB2	**12.0 GB (memory pools about 2 GB)**
Tivoli Directory Server (ITDS)	**1.5 GB**
Idling guest	**1.0 GB**

- **Test scenarios**
  - Leave the guest size fix
  - Decrease the LPAR size in predefined steps to scale the level on memory overcommitment
  - Measure the execution time of a predefined workload (TPM)

Test Results

Memory – less is better		Performance – more is better
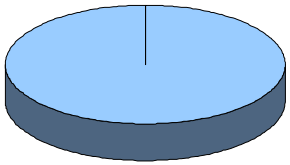 100%	**BASE settings = 100%** • Sum of guest size definition • Base performance	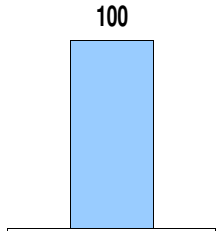
8% saved 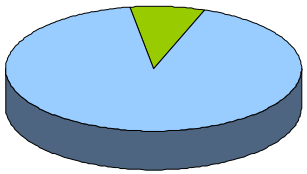	**OPTIMAL settings** + Reduce memory by 8% + Improved performance by 6%	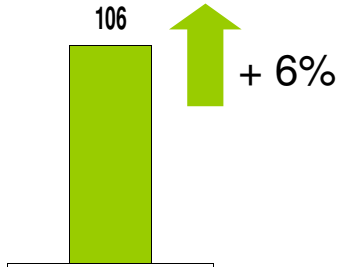
64% saved 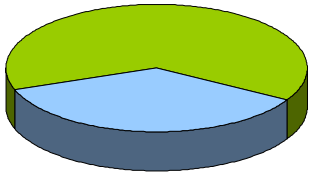	**CHEAPEST settings** + Reduce memory by 64% – Decreased performance by 7% – effective for test & development env.	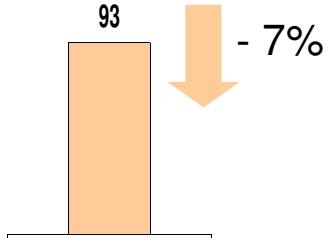

Virtualization Considerations for Linux guests

- Default SRM settings (q srm) - CMS guests are typically Q1 guests and Linux Q3 guests
  - **LDUBUF** : Q1=100% Q2=75% Q3=60%
  - **STORBUF**: Q1=125% Q2=105% Q3=95%
- LDUBUF – Defines amount of paging "capacity" to be used in scheduler algorithms
  - 60% for Q3 means:
    All Q3 guests together can use maximum 60% of paging resources - if already used → eligible list
  - Recommendation:
    SET SRM LDUBUF 100 100 100
    to allow all Q3 guests to allocate the whole paging space
- STORBUF – to partition host storage (central storage)
  - 95% for Q3 means:
    All Q3 guests together can use only 95% of the system storage
    → This prevents memory overcommitment when running Linux guests
  - Recommendation:
    SET SRM storbuf 300 250 200
    to allow all Q3 guests to allocate twice the amount of real storage
  - Depending on the level of overcommitment and amount of active/inactive guests, it might be necessary to go even higher, e.g. SET SRM storbuf 300 300 300
  - Ensure to have sufficient paging space!

The economics of virtualization and consolidation with Linux on z

- Consolidate an average of ***30 distributed servers*** or more on a single core, or ***hundreds*** in a single footprint.
- Deliver a virtual Linux server for approximately ***$500 per year*** or as little as a ***$1.45 per day per virtual server*** (TCA)[1]

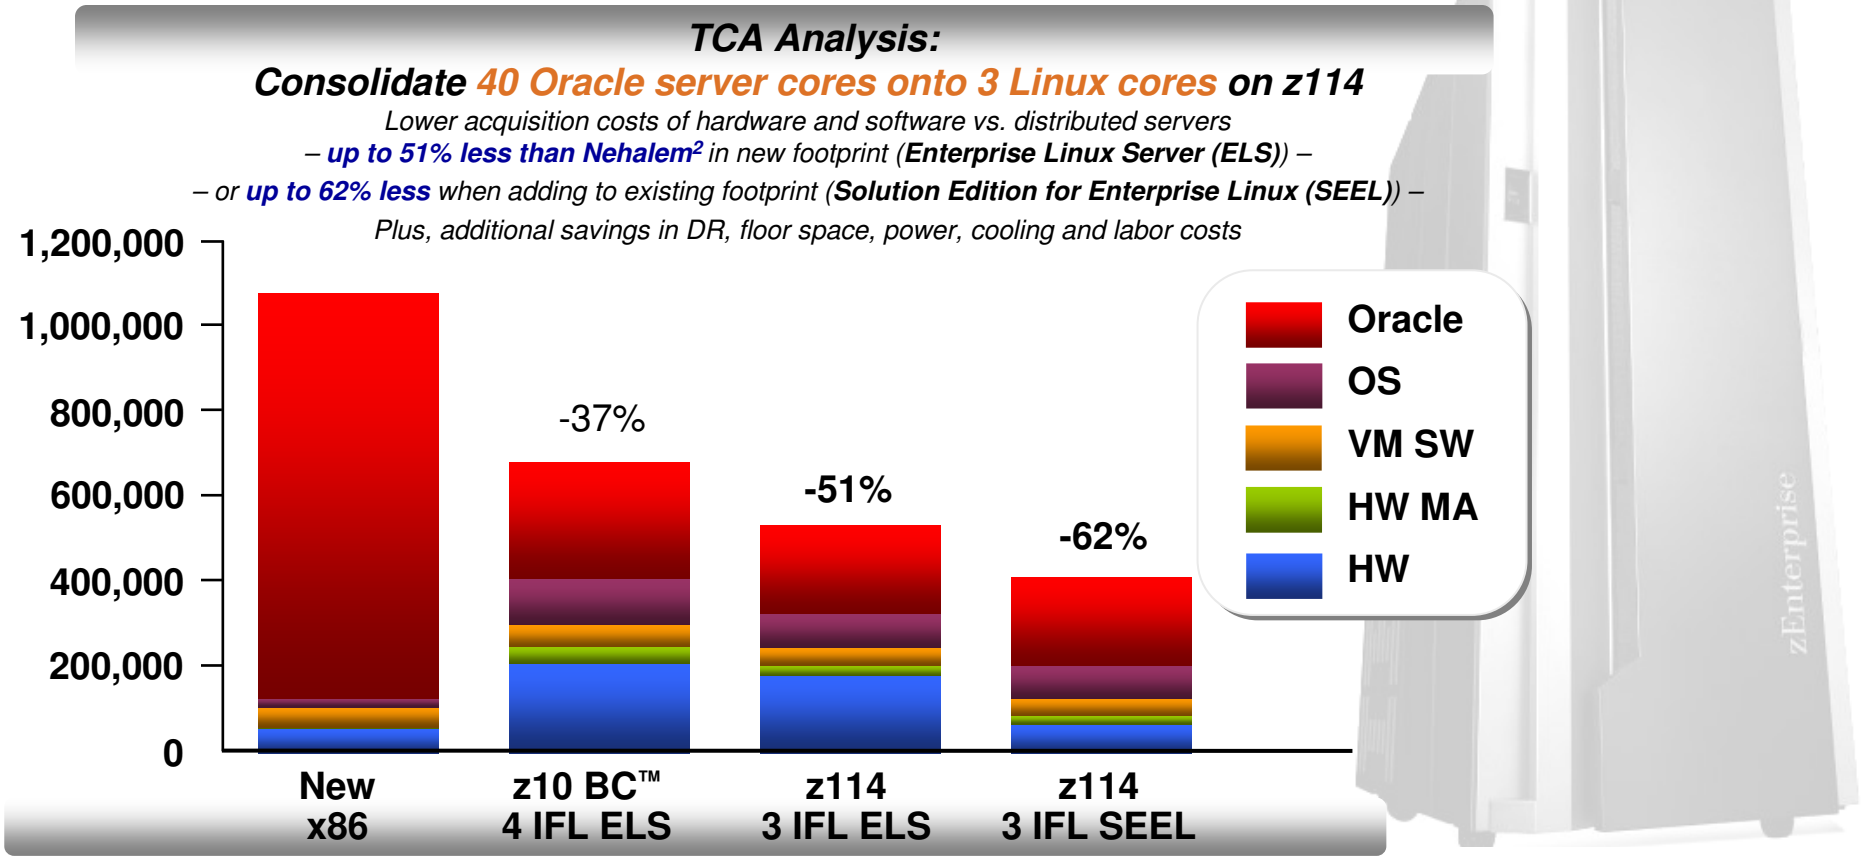

[1] Based on US Enterprise Linux Server pricing. Pricing may vary by country. Model configuration included 10 IFL cores running a mixed workload averaging 31 virtual machines per core with varying degrees of activity. Includes zEnterprise hardware and z/VM virtualization software. Does not include Linux OS or middleware software.

[2] Distributed server comparison is based on IBM cost modeling of Linux on zEnterprise vs. alternative distributed servers. Given there are multiple factors in this analysis such as utilization rates, application type, local pricing, etc., savings may vary by user.

What is Different About Cloud Computing?

Without cloud computing

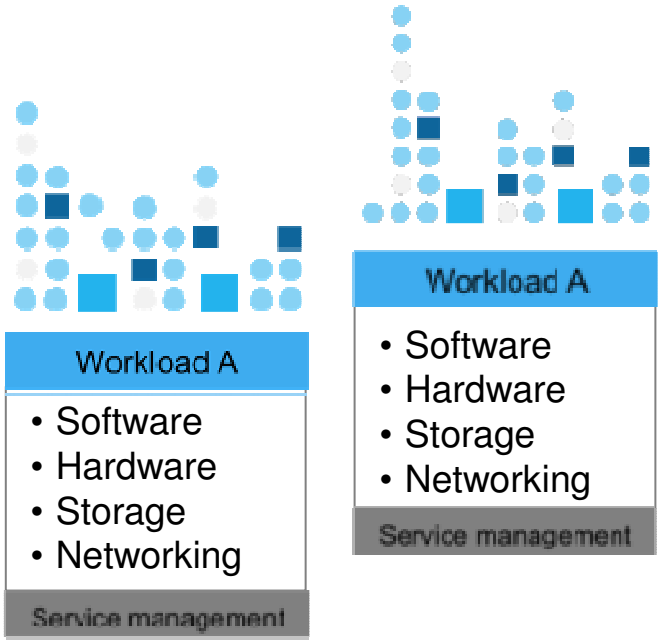

With cloud computing

- **Virtualized resources**
- Automated service management
- Standardized services
- Location independent
- Rapid scalability
- Self-service

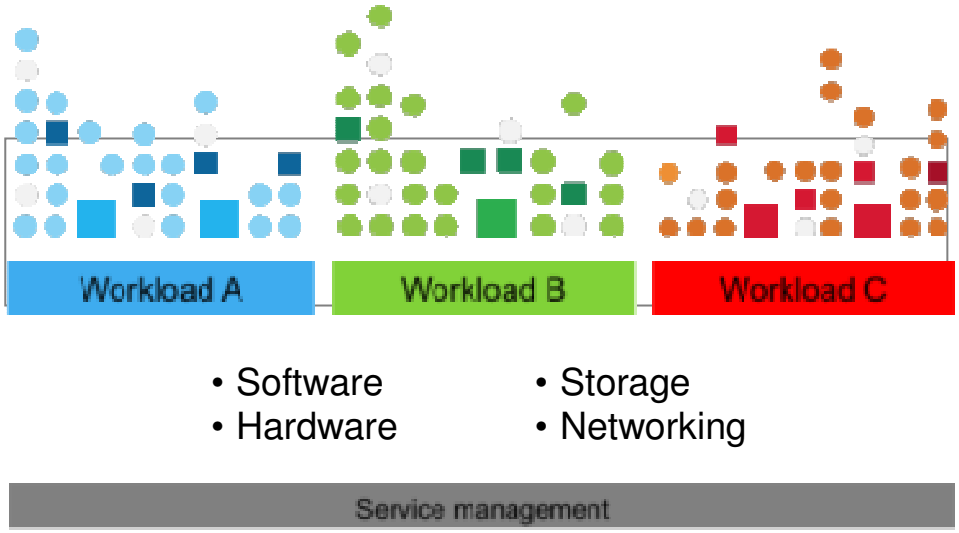

Definition – National Institute of Standards and Technology

SP 800-146	May 2012	Cloud Computing Synopsis and Recommendations sp800-146.pdf
SP 800-145	Sept. 2011	The NIST Definition of Cloud Computing SP800-145.pdf

http://csrc.nist.gov/publications/PubsSPs.html#800-145

IBM Cloud Computing Reference Architecture (CC RA) – Cloud Livecycle Management 'Best Practices'

Publically available RA whitepaper on ibm.com:
http://public.dhe.ibm.com/common/ssi/ecm/en/ciw03078usen/CIW03078USEN.PDF

- The IBM CC RA is based on Best Practices from:
  - IBM Cloud Projects with clients
  - IBM Public Cloud Offerings like Smart Cloud Enterprise
- The CC RA consists of 21 detailed Documents, with best-of-industry knowledge regarding Cloud Architecture, Design and Implementation

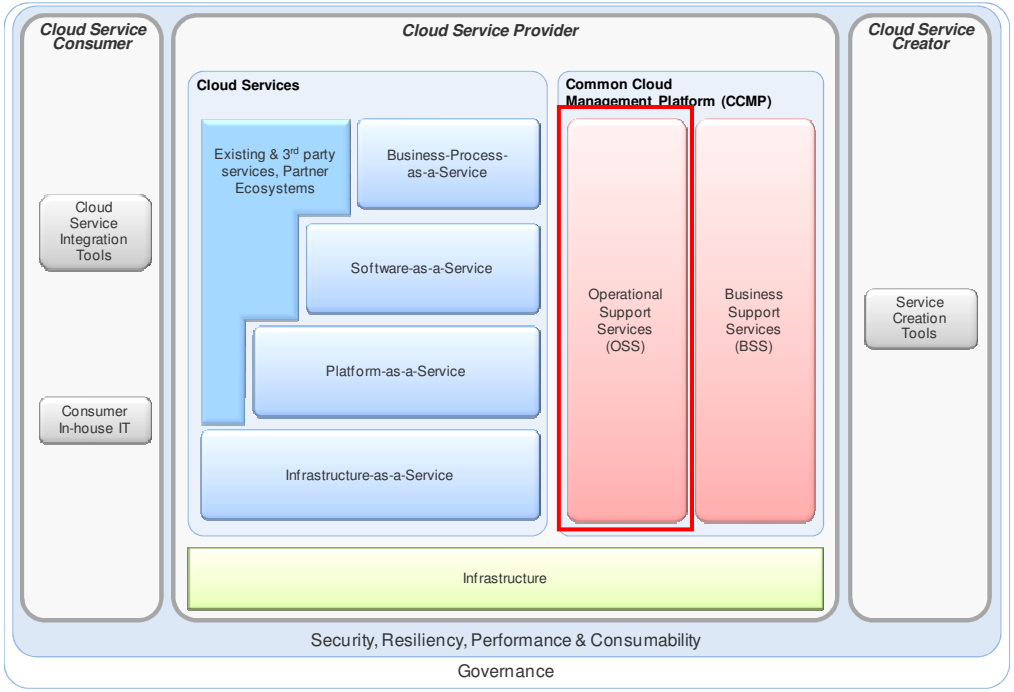

CCRA OpenGroup submission:
http://www.opengroup.org/cloudcomputing/uploads/40/23840/CCRA.IBMSubmission.02282011.doc

http://www.opengroup.org/cloudcomputing/uploads/40/23840/CCRA.IBMSubmission.02282011.doc

Cloud Computing – Deployment, Service, Characteristics

4 Deployment Models	3 Service Models	5 Characteristics
▪ Private cloud ▪ Community cloud ▪ Public cloud ▪ Hybrid cloud	▪ Software as a Service (SaaS) ▪ Platform as a Service (PaaS) ▪ Infrastructure as a Service (IaaS)	▪ Rapid elasticity ▪ Broad network access ▪ Resource pooling ▪ Measured service ▪ On-demand self-service

Decision criteria

http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

Source: National Institute of Standards and Technology (NIST)

IBMSmartCloud Foundation

- System z Offerings

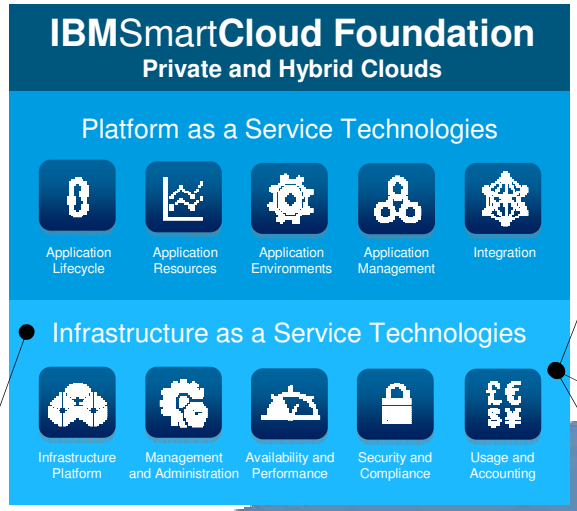

Integrated Systems

IBM PureFlex System

Private Cloud Management Solutions

IBM SmartCloud Entry
IBM SmartCloud Provisioning
IBM Tivoli Service Automation Manager
IBM Solution Edition for Cloud Computing (System z)
IBM Service Delivery Manager (ISDM)
IBM SmartCloud Virtual Storage Center
IBM Tivoli System Automation
IBM SmartCloud Cost Management
IBM Endpoint Manager Solutions
IBM SmartCloud Patch Management
IBM SmartCloud Monitoring
IBM Service Management Extension for Hybrid Cloud
IBM SmartCloud Application Performance Management

Infrastructure Systems and Storage

IBM System x with ex5 technology
IBM Power Systems with POWER7
IBM System z with Unified Resource Manager
CSLWave, xCat
IBM Scale out NAS Storage Systems
IBM BladeCenter Foundation for Cloud
IBM Storwize V7000 Unified, IBM XIV Storage Systems Gen 3, IBM iDataplex

Infrastructure Management Solutions

IBM Cloud Ready for Linux on System z
IBM SmartCloud Control Desk
IBM Security Virtual Server Protection for VMware
IBM Systems Director
Tivoli Storage Manager for Virtual Environments

System z Cloud Roadmap

	1 Integrate	2 Automate	3 Orchestrate
IBM Products & Offerings	▪ zEnterprise: zEC12, z196, zBC12, z114 ▪ Linux on System z	▪ Tivoli Provisioning – Cloud Ready	▪ Tivoli Service Automation Manager
openstack CLOUD SOFTWARE	▪ z/VM 6.3 **New**	▪ SCE/SCP	▪ SCO *System z support currently in development*
ISV Solutions	**CSL Wave** ▪ Provided by **CSL International** ▪ Hypervisor Manager **IBM recent acquisition**	**zPRO** ▪ Provided by **Velocity Software**	**MOAB** ▪ Provides a policy based cloud management based on xCAT
Open Source Options	**xCAT** ▪ Extreme Cloud Administration Tool	**OpenStack** ▪ Being enabled for z/VM first and eventually zManager ▪ Being used as a code base for SCE / SCP / SCO	

Openstack – Industry Accepted IaaS Cloud Computing Model

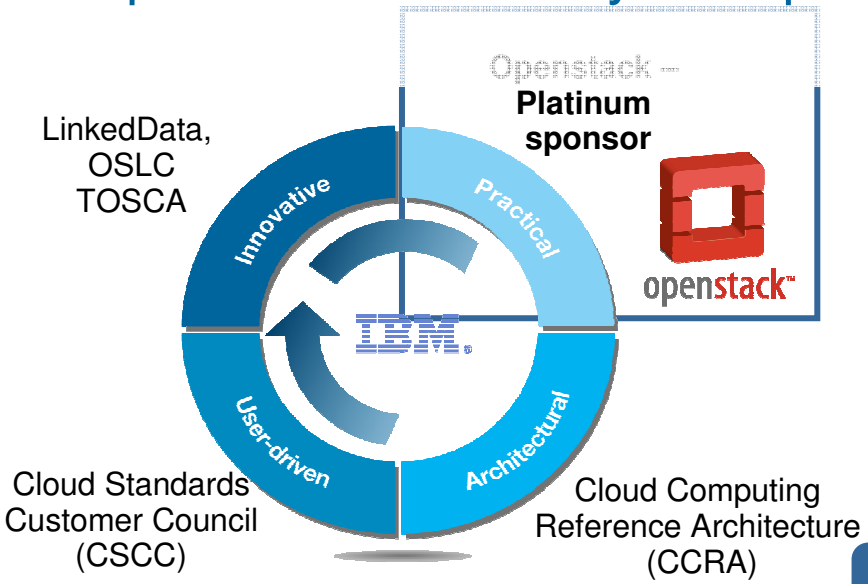

IBM joins the new OpenStack Foundation as Platinum Sponsor

Objectives

- OpenStack is an Infrastructure as a Service (IaaS) cloud computing project that is free open source software released under the terms of the Apache License.
- Managed by the OpenStack Foundation, a non-profit corporate entity established in September 2012
- More than 150 companies have joined the project among which are AMD, Intel, Canonical, SUSE Linux, Red Hat, Cisco, Dell, HP, IBM, NEC, VMware and Yahoo
- It is portable software, but is mostly developed and used on the Linux operating system.

http://en.wikipedia.org/wiki/OpenStack

System z - Participation in OpenStack

Self-Service

Services Catalog

Cloud Management Application

(including SmartCloud technologies)

Multi-Tenant

Billing and Charge-back

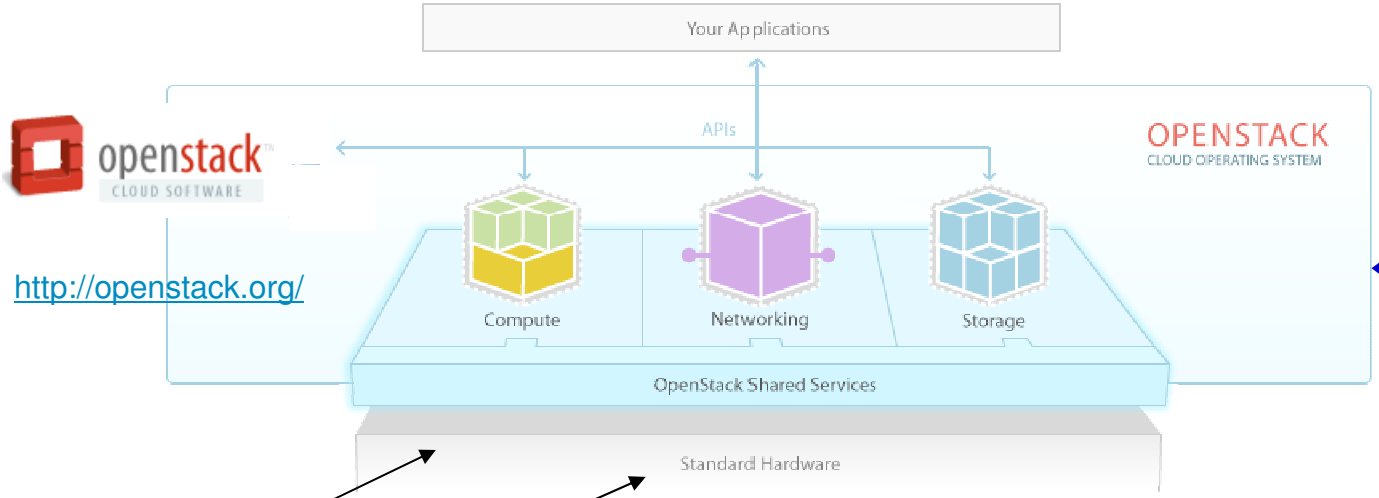

http://openstack.org/

Compute (Nova)
Storage (Cinder)
Network (Quantum)
Provision and manage virtual resources such as servers, networking, and storage.

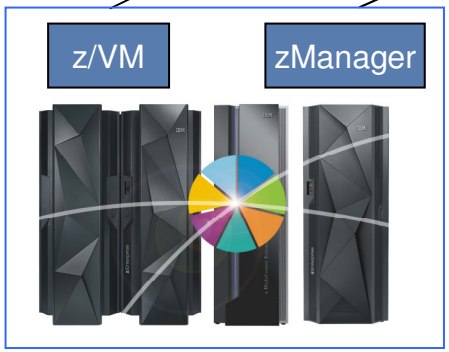

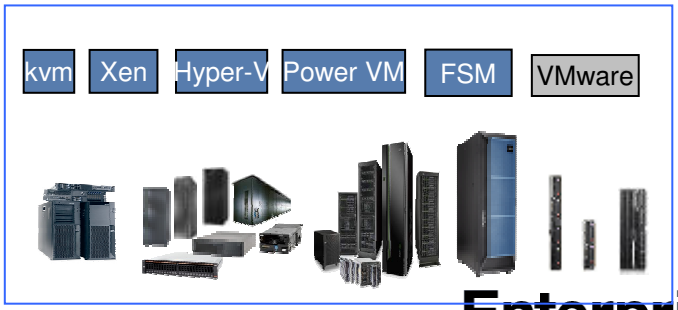

Direct Hypervisor Integration
System z support (z/VM followed by zManager) being added to OpenStack, in the same manner as other hypervisors

A z/VM Cloud Management with Open Source multi-platform tool xCAT

xCAT Architecture on System z

http://www.vm.ibm.com/sysman/

- **Stands for Extreme Cloud Administration Toolkit - xCAT**
- **Tool to manage, provision, and monitor physical and virtual machines including IBM servers**
- **Open sourced in 2007 and licensed as EPL (Eclipse Public License)**
- **Used by NASA, University of Toronto, IBM, Adaptive Computing, Los Alamos Laboratory, and more!**

xCAT MN: Central management server running on normal Linux

zHCP: Runs on privileged VM and manages other VMs via SMAPI and CP

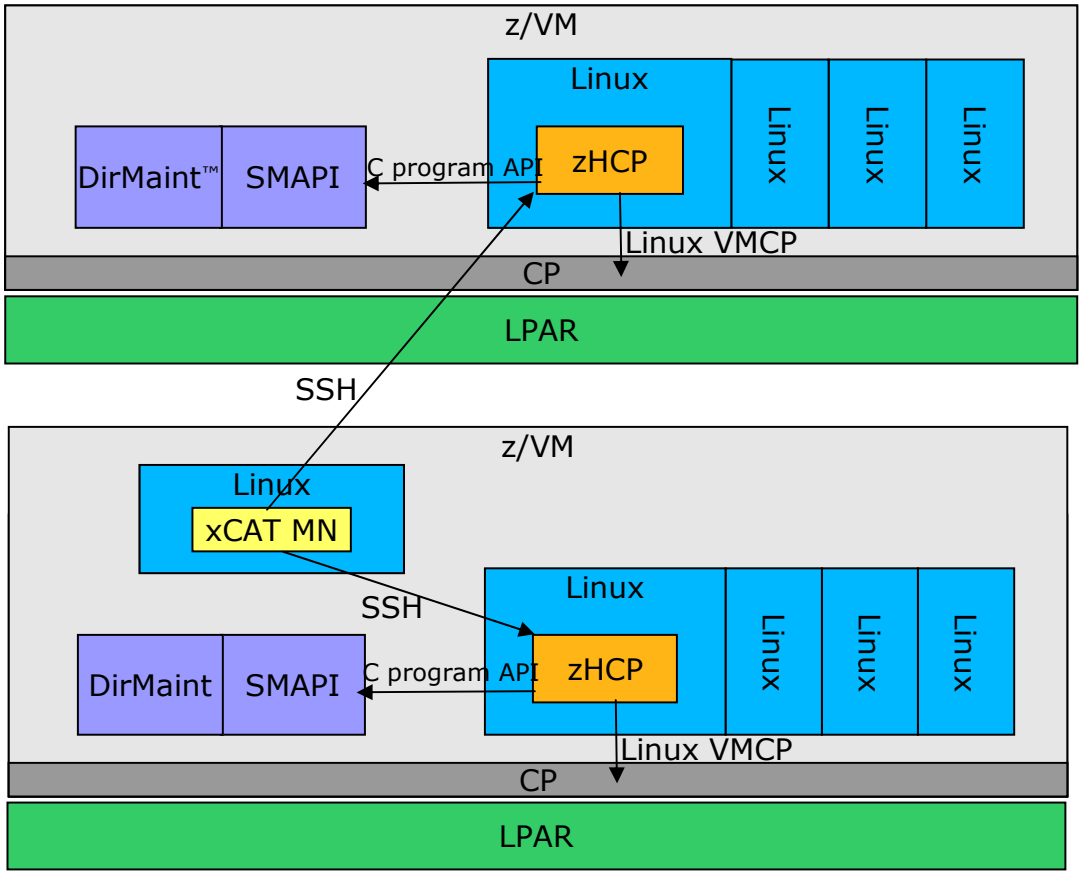

xCAT GUI – Nodes, Actions, incl. Migrate

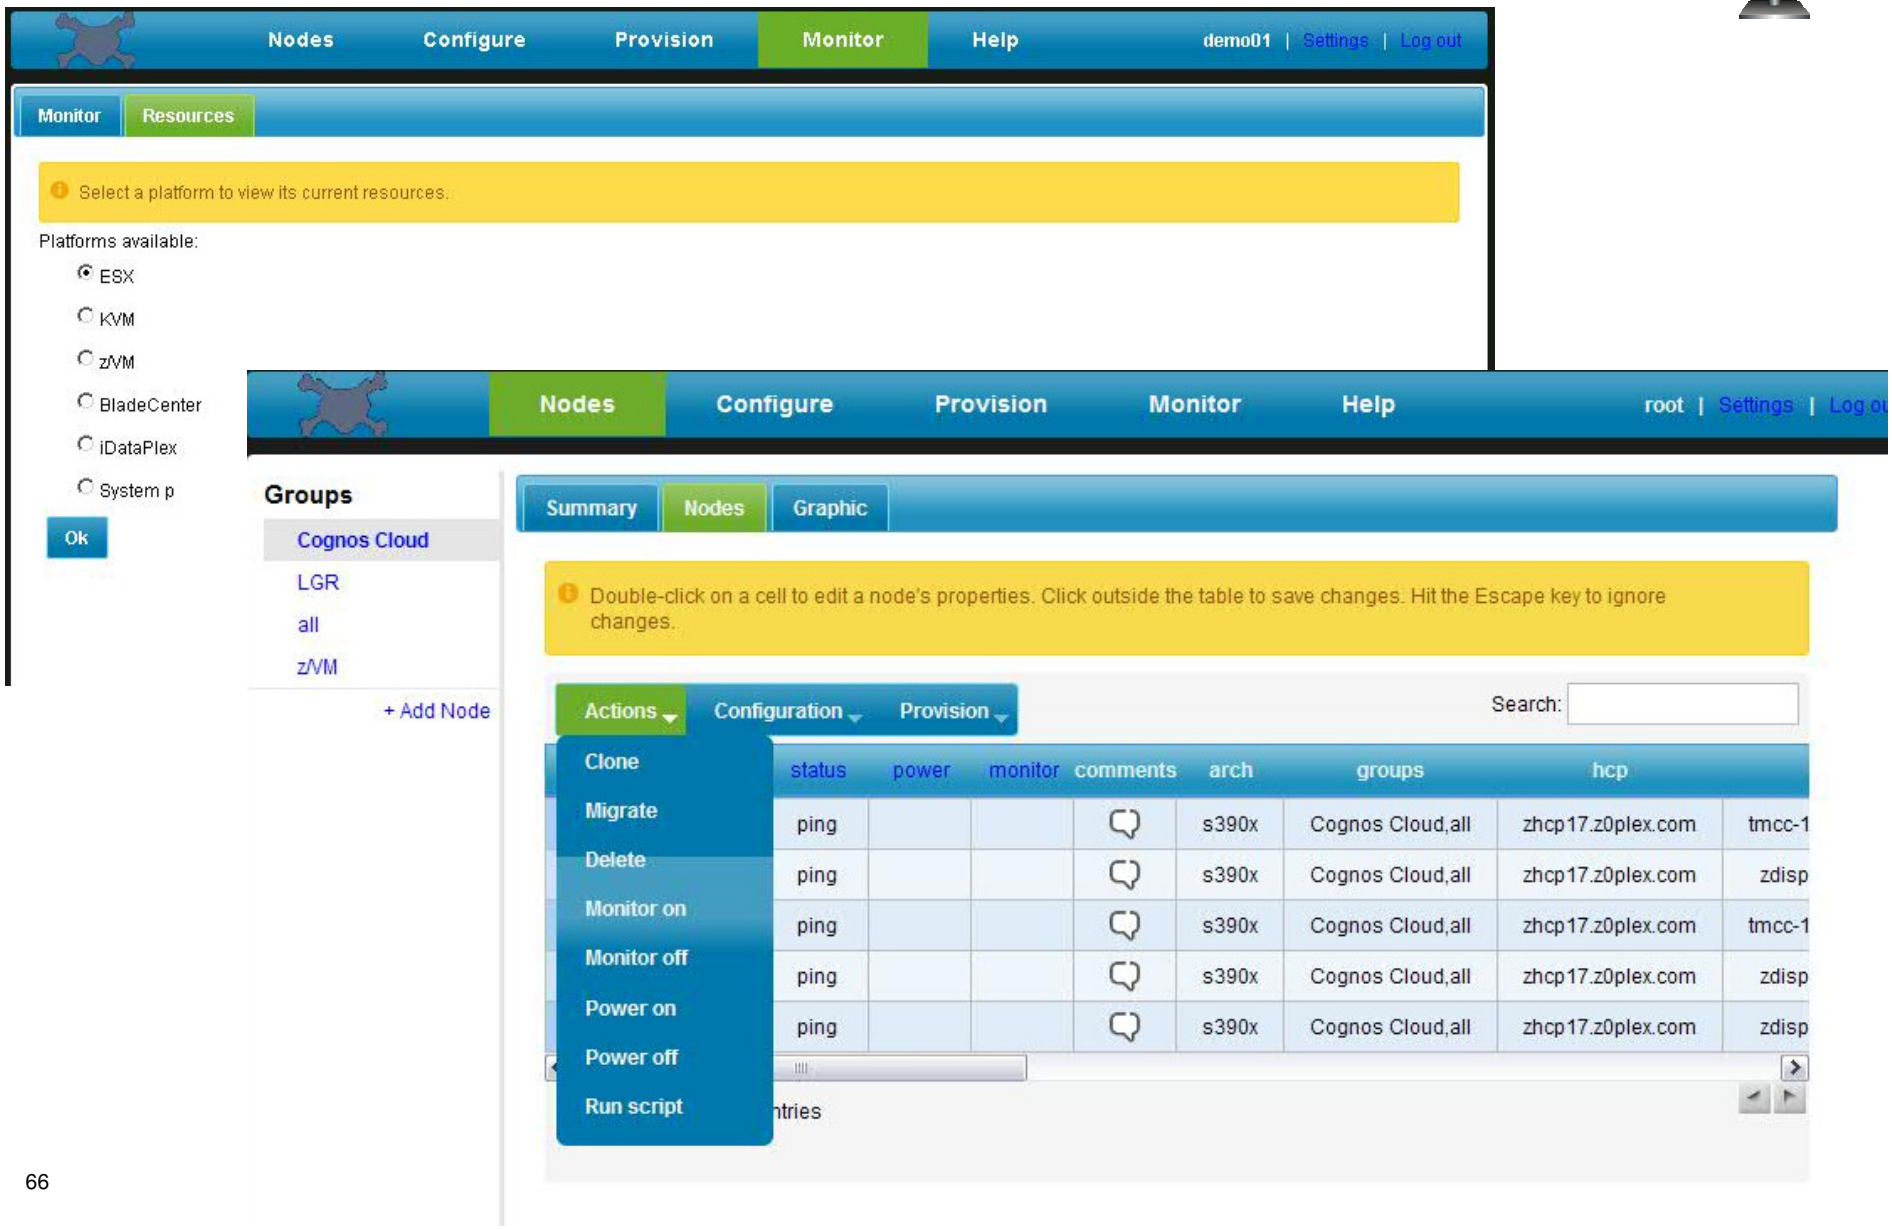

xCAT GUI – Nodes, Summary

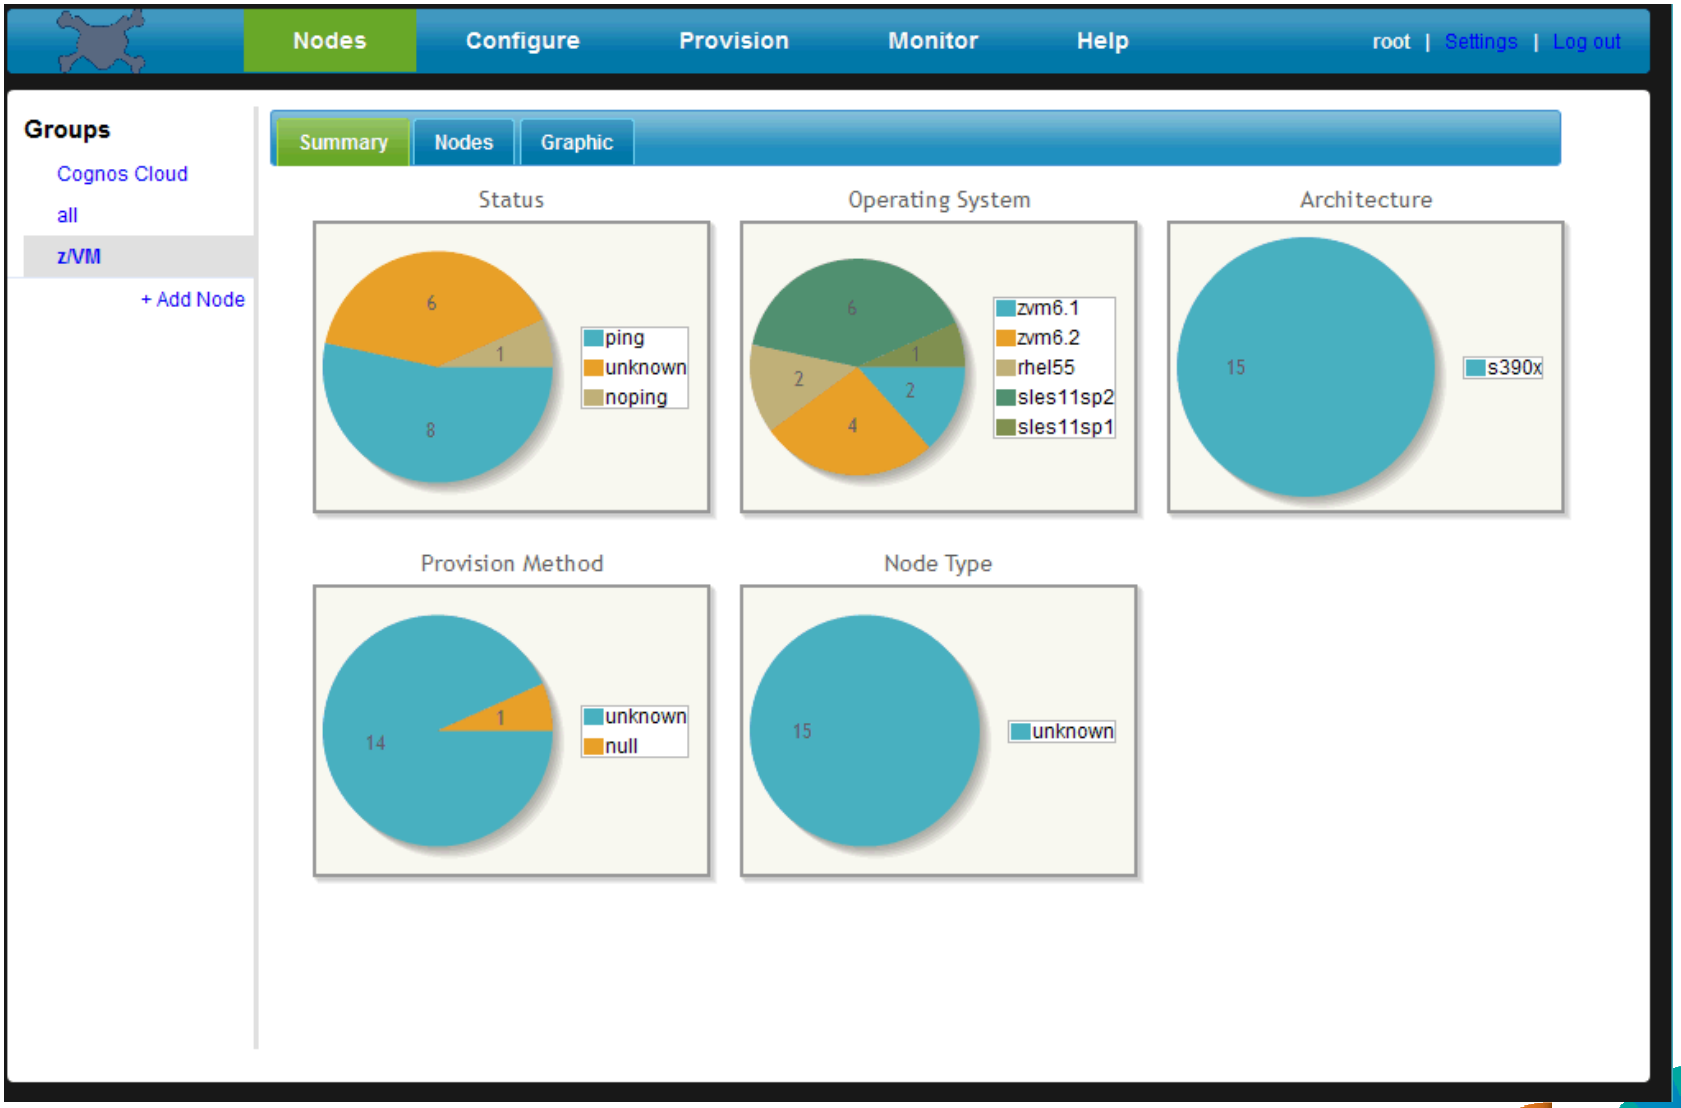

New IBM tool: CSL-WAVE – z/VM and cloud management

Hardware Viewer – the z/VM Users group

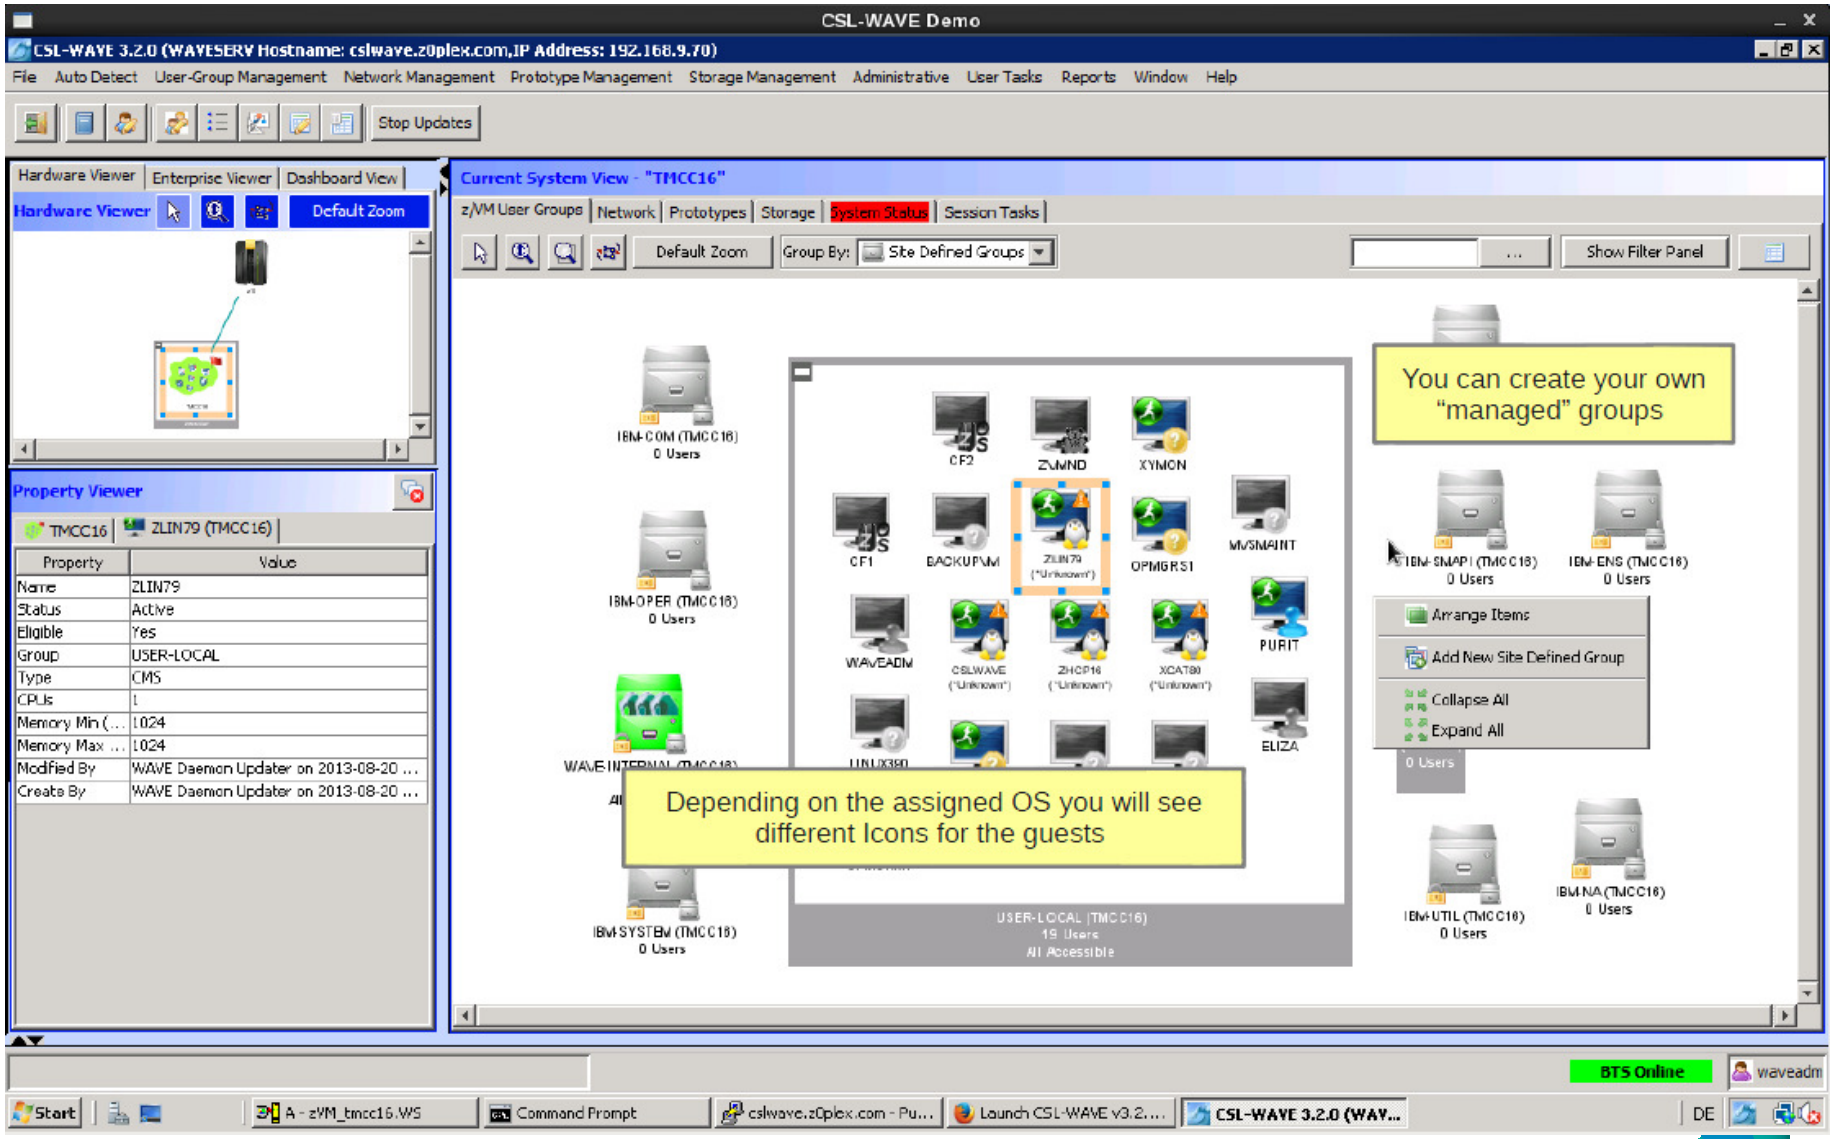

Hardware Viewer – the Network

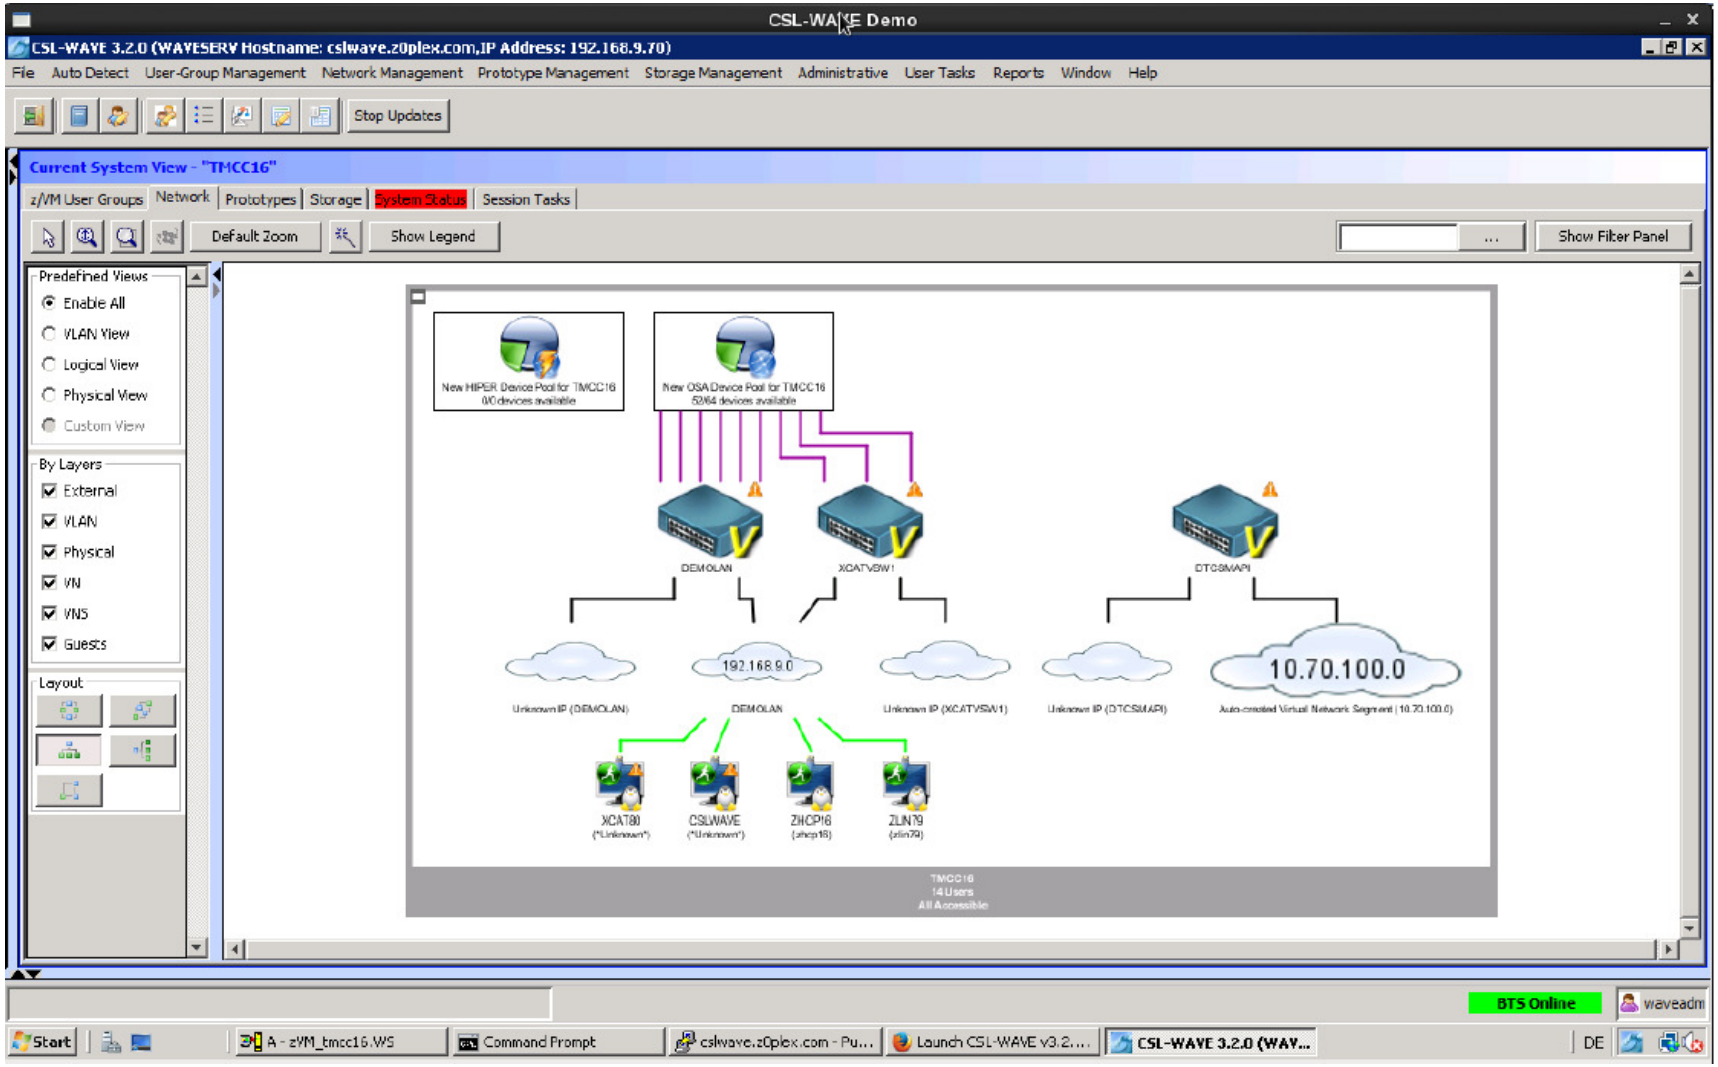

Dedicated Device Pools

- Drilling into the Device Pool Manager, we get to the actual Device Pools which list information about each device in the pool
- Information such as:

  Device Address

  ONLINE status

  Usage

  Owner

  LOCK status

  IAN (Intelligent Active Note) status

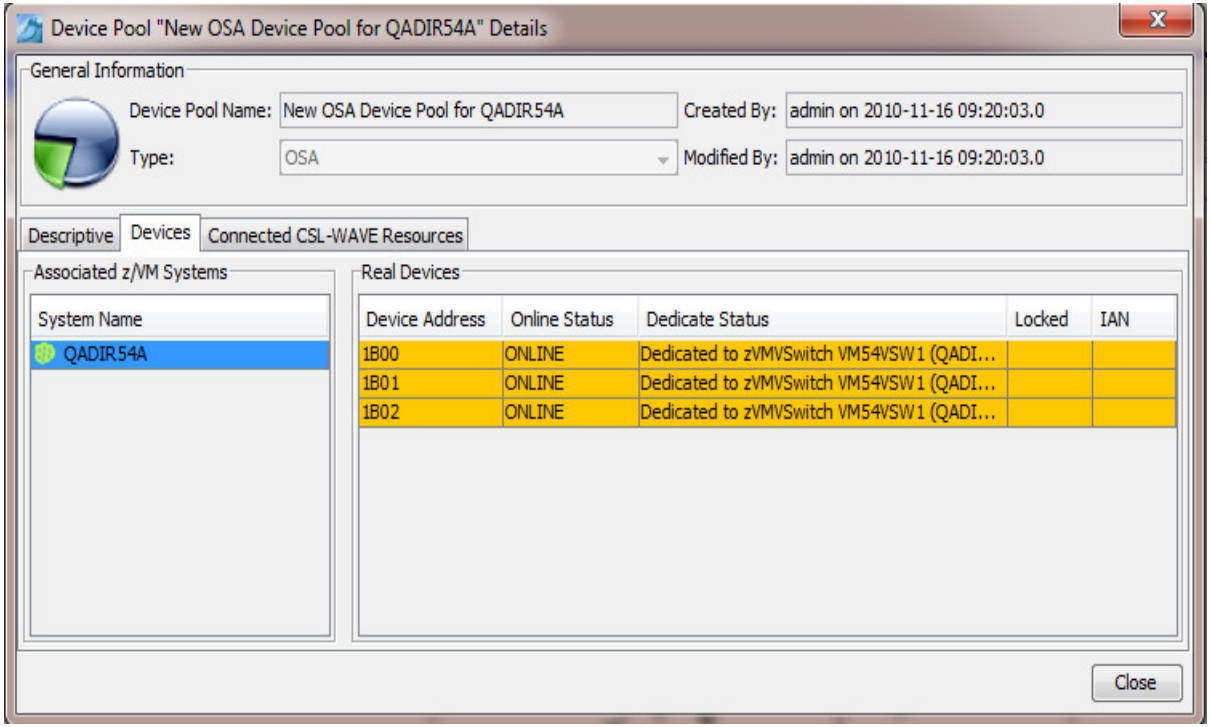

System Viewer – the system activity

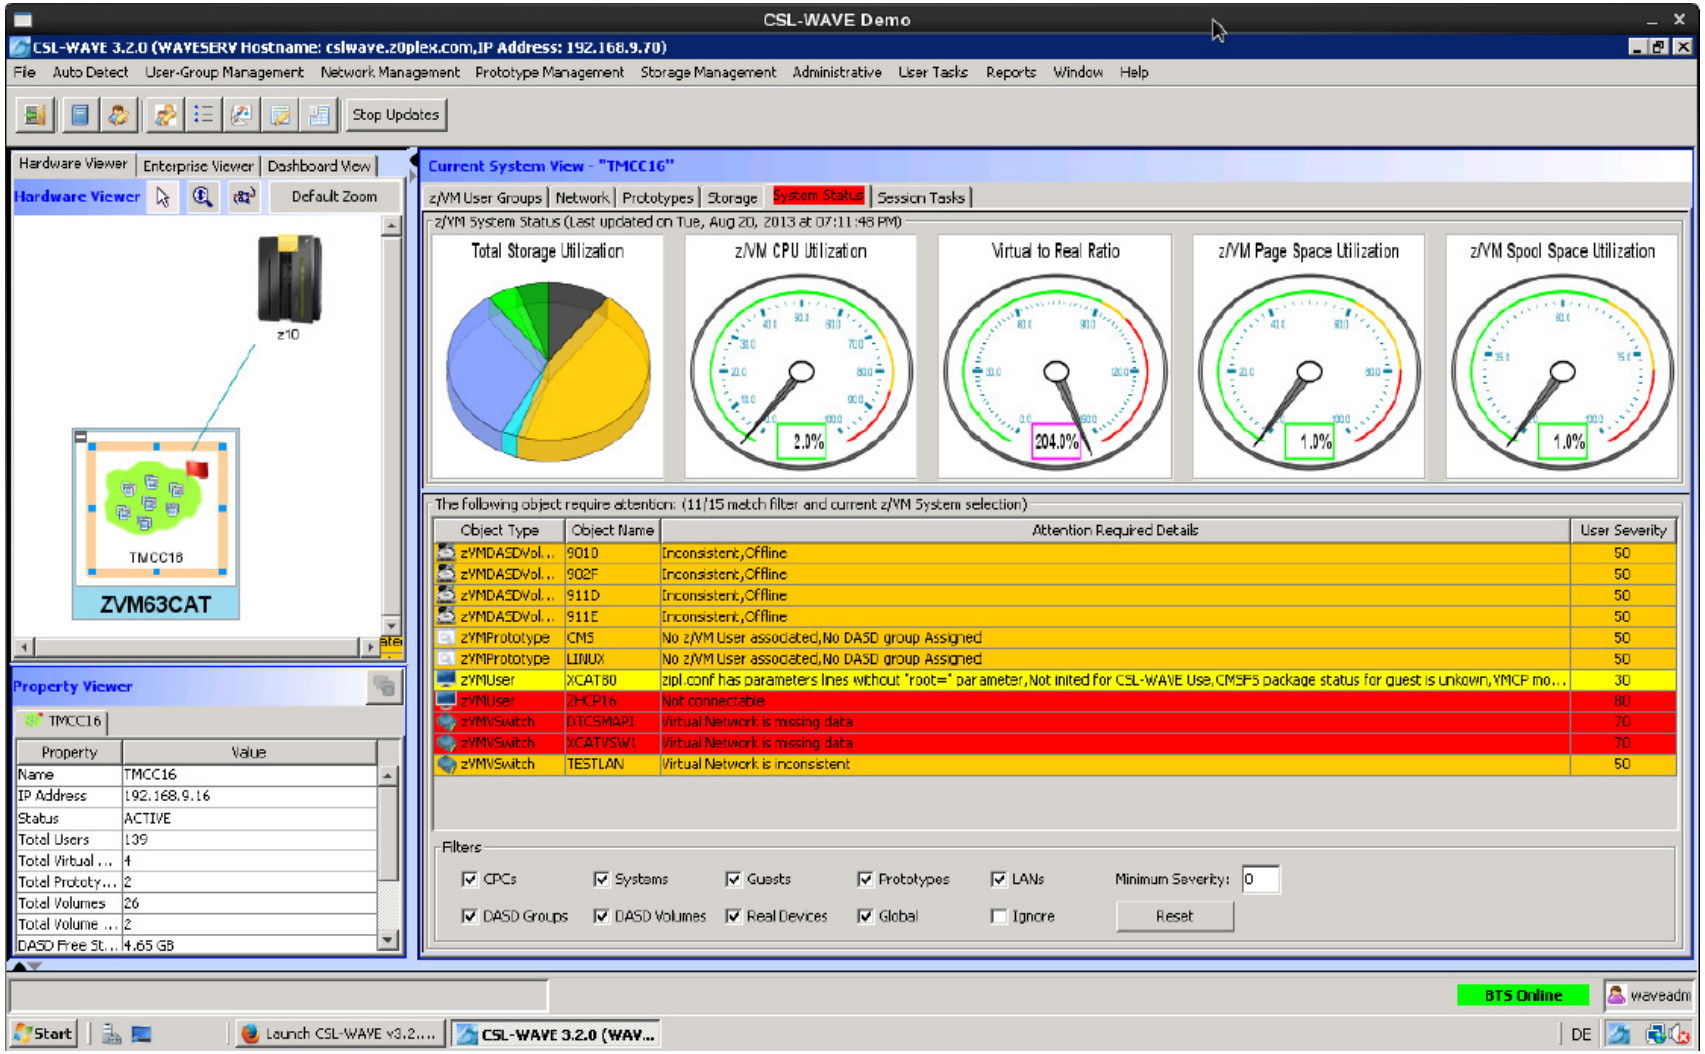

MobileFirst: Linux on z and IBM Worklight Server – A Cloud environment on System z for mobile Devices

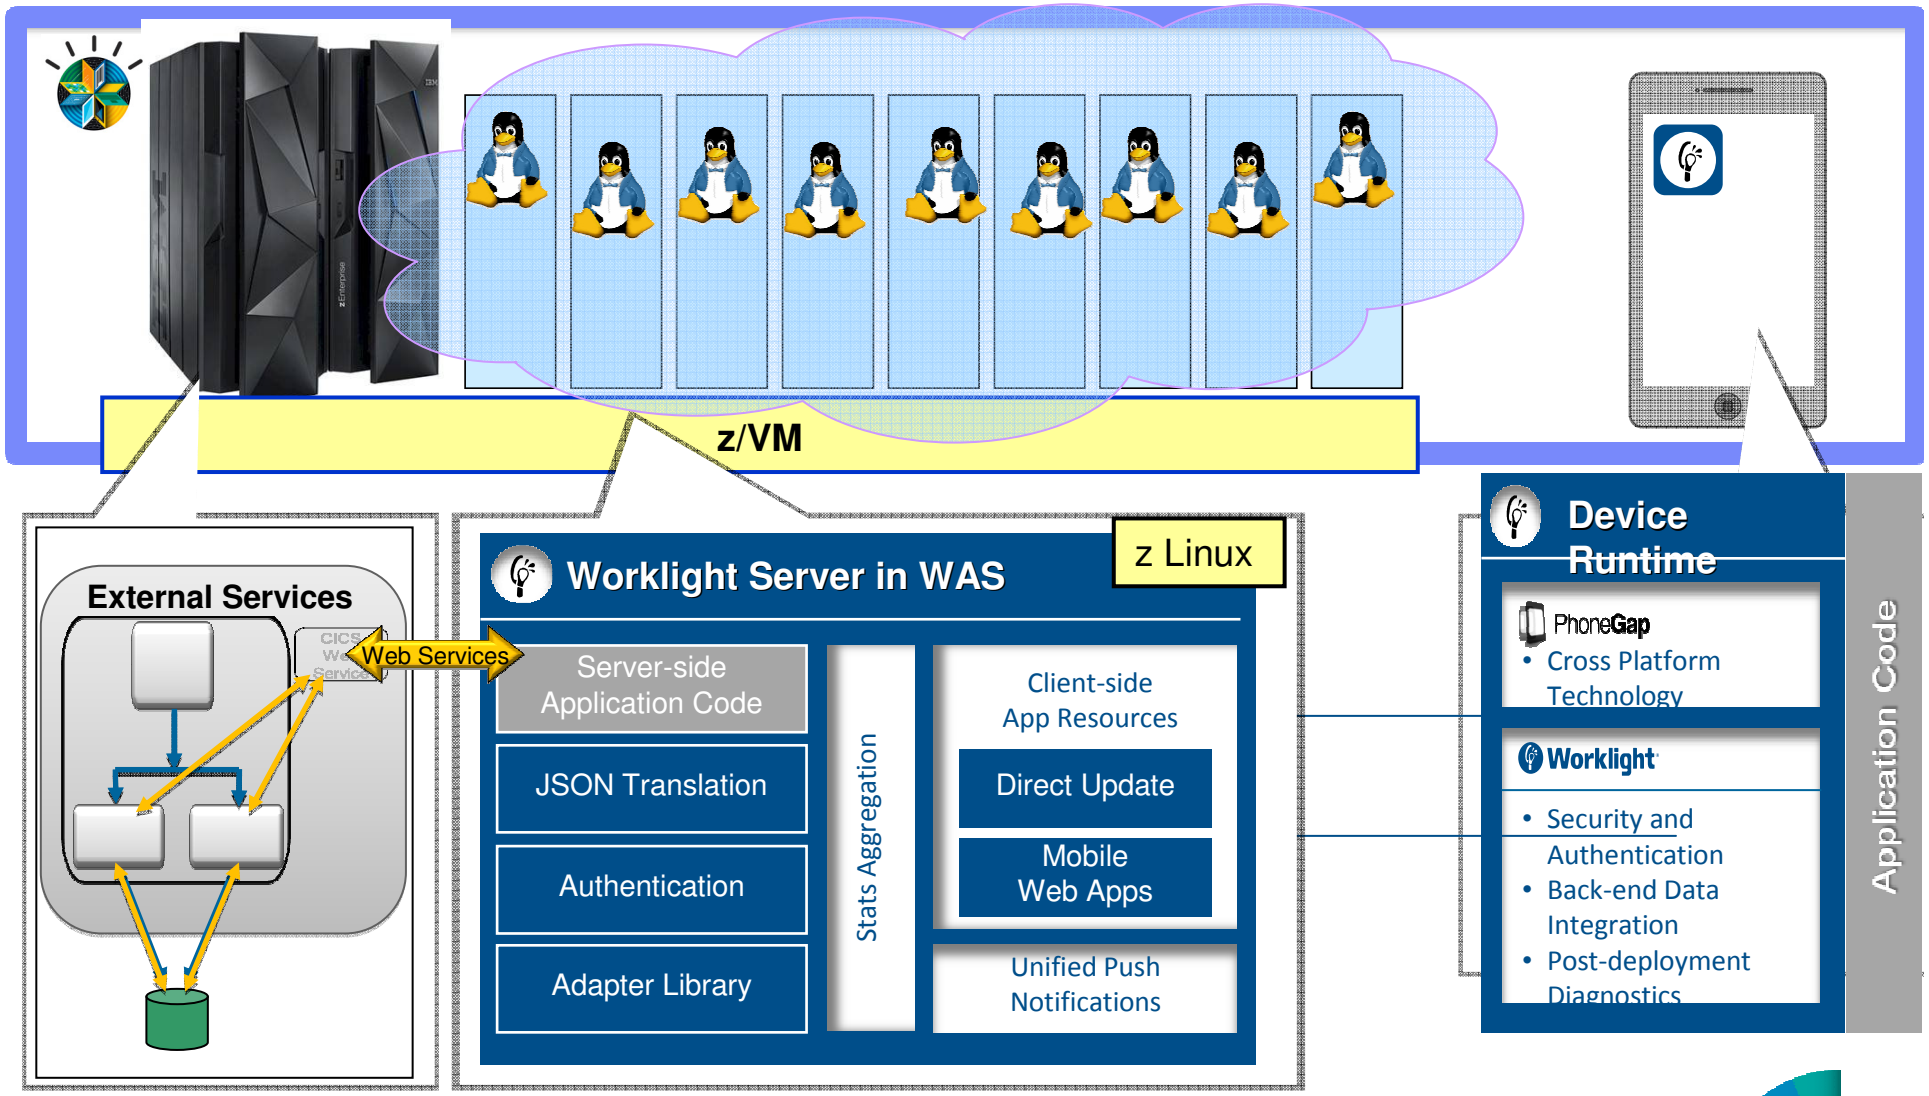

Cloud Test Drive with Linux on System z - Effective Virtualized environment

- **Up to 90 days, free of charge, access to up to 3 Linux on z servers under z/VM**
- Hands-on experience with Cloud, Linux on z, z/VM, Tivoli Provisioning Manager (TPM), and a selection of 5 system images based on SUSE or Red Hat
  - SLES 11 SP1 Base
  - RHEL 5.8 Base
  - SLES 11 SP1, DB2 9.7 Fixpack 5, WAS 8.5, IBM HTTP Server 8.5
  - RHEL 5.8, DB2 9.7 Fixpack 5, WAS 8.5, IBM HTTP Server 8.5
  - SLES 11 SP1, Oracle 11gR3, WAS 8.5, IBM HTTP Server 8.5
- Simple remote access over the internet to zEnterprise in the IBM Washington System Center in Gaithersburg, Maryland
- Customize your own Linux cloud with your own secure data
- Guided exercises for training

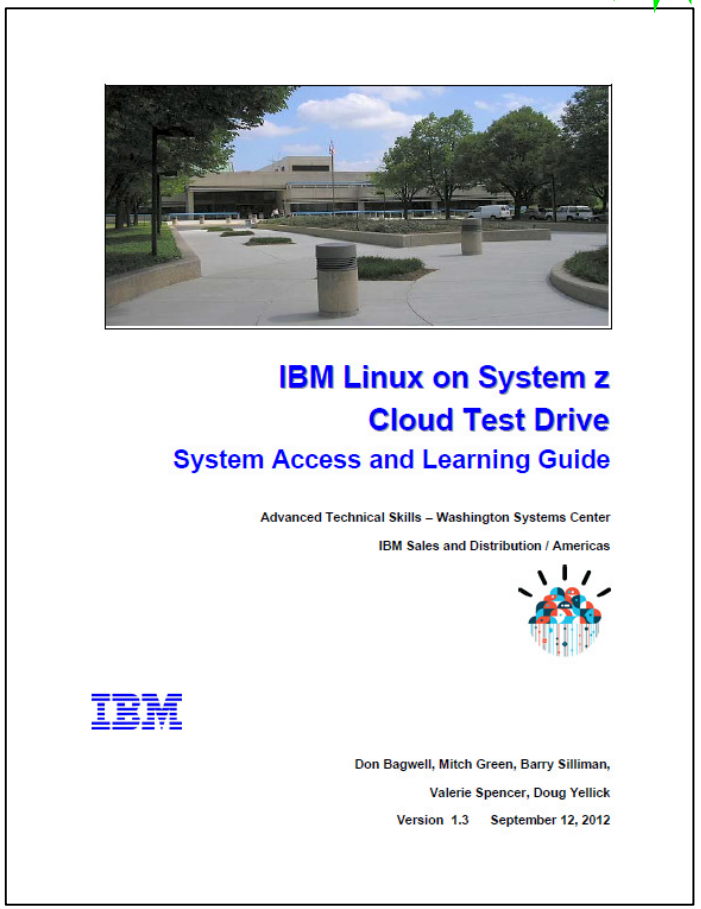

http://techsales5.austin.ibm.com/tsna/techxpress.nsf/request.html

IBM System z customers with effective virtualization and cloud solutions

http://www-03.ibm.com/systems/z/os/linux/solutions/

TRANSZAP

Moved to System z from Lintel to deliver the availability and security their clients demand of their e-Procure-to-Pay SAAS, while supporting the strong growth the company is experiencing

Casas Bahia centralized operations on System z to support rapid growth and reduce IT costs

Nationwide On Your Side

Consolidated Windows-based systems to Linux on z to achieve substantial cost efficiencies and introduce cloud solutions

Satyam — The Official IT Services Provider of the 2010 FIFA World Cup™

Satyam has positioned the mainframe as a platform to reach the SMB audience in growth markets with hosted web business services

MARIST
www.marist.edu

Entering provider space for cloud services for universities, schools systems and other public entities

HOPLON infotainment

Their massive-multi-player game and virtual world application middleware runs on System z. (www.taikodom.com)

IBM Linux Technology Center, LTC

IBM Linux Technology Center (LTC)

Development team for all server and software platforms and other key initiatives, such as:

- Linux on POWER, Linux on Cell, Linux on System z and z/VM
- Security: EAL certifications, Trusted Computing, SELinux, sHype
- Virtualization: z/VM, Xen/KVM, APV support
- Systems Management: kdump, SystemTap, xCAT

Technical liaison to IBM's customers and Linux Distribution Partners

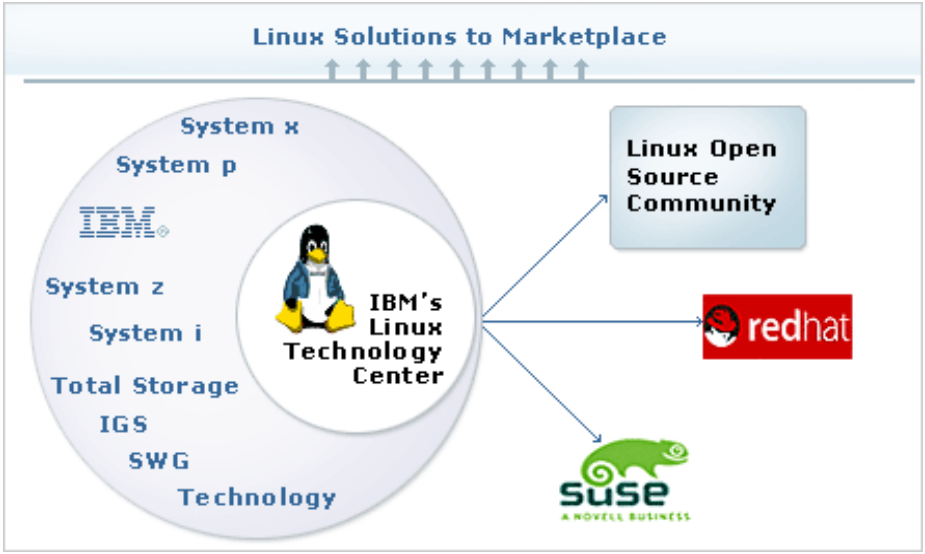

LTC:
Over 600 developers
40+ locations
100+ projects

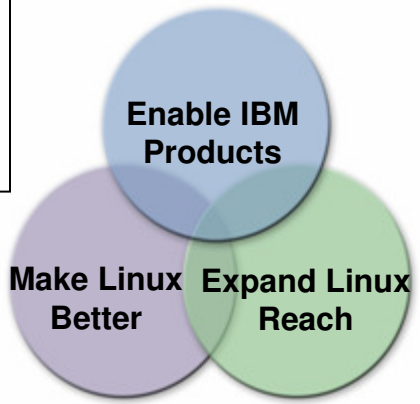

http://www-03.ibm.com/systems/z/os/linux/

Additional Information

- Enterprise Linux Server information:
  http://www-05.ibm.com/de/promotions/els/
- IBM Information Center - Linux on System z
  http://publib.boulder.ibm.com/infocenter/lnxinfo/v3r0m0/index.jsp?topic=%2Fliaag%2Fl0lgr00_2012.htm
- Linux on z Hints and Tips
- http://www.ibm.com/developerworks/linux/linux390/perf/index.html
- RedBooks
  http://www.redbooks.ibm.com/portals/linux

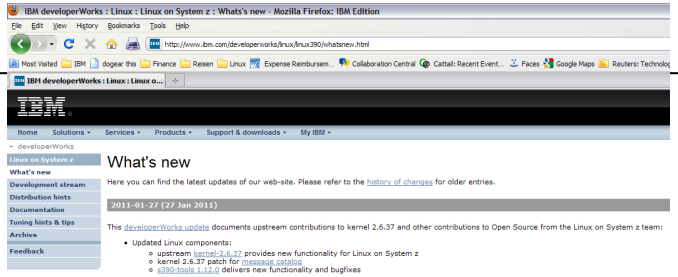

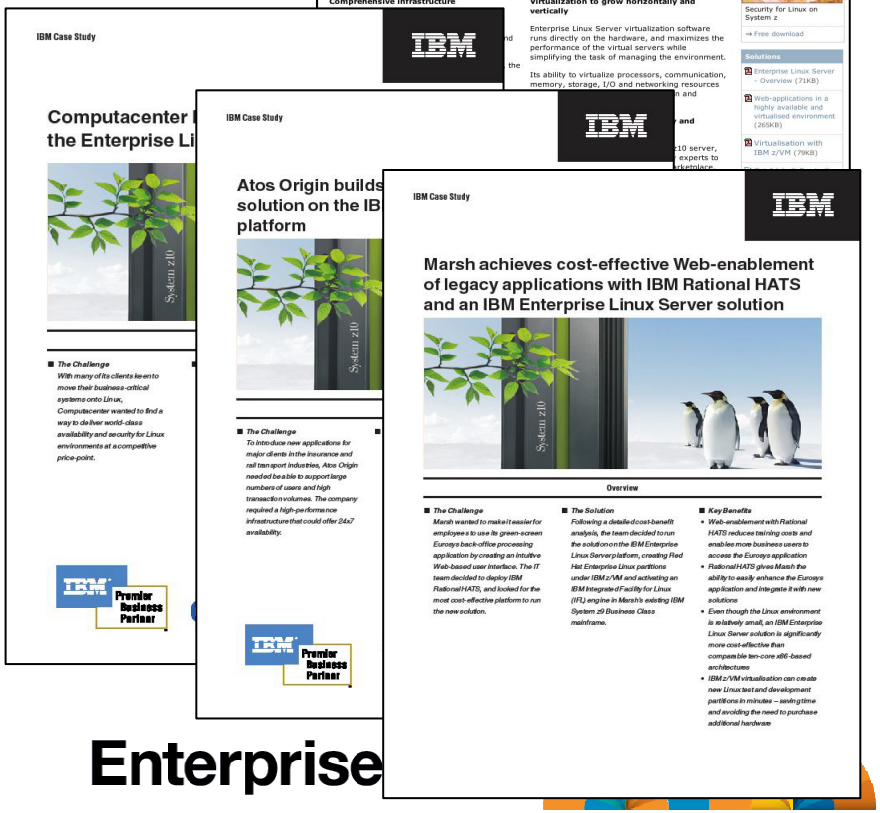

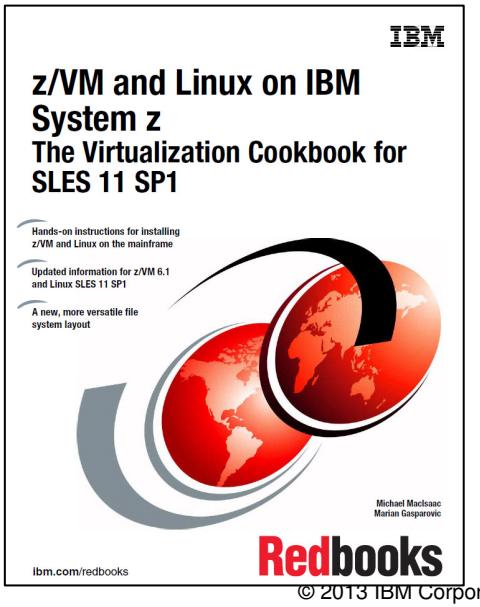

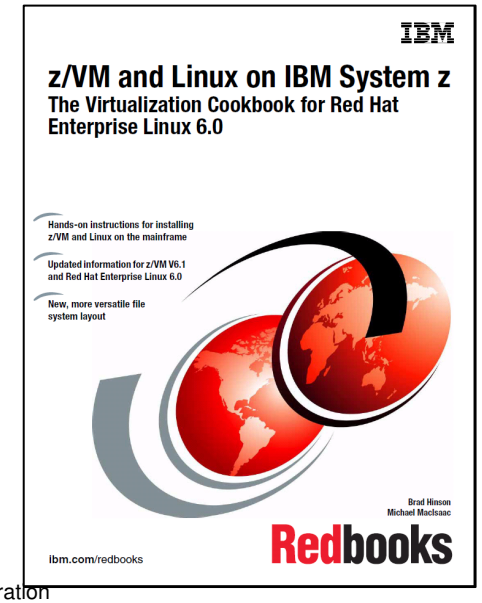

Questions?

Wilhelm Mild

IBM IT Architect

THE Open GROUP
Master
Certified IT Architect

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

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

Additional Information

- **IBM Systems - System z - z/VM, Home**
  http://www.vm.ibm.com/ | http://www.ibm.com/vm/

- **Linux on System z Information Center**
  http://publib.boulder.ibm.com/infocenter/lnxinfo/v3r0m0/index.jsp?topic=%2Fliaag%2Fl0lgr00_2012.htm

- **IBM developerWorks - Technical topics - Tuning hints & tips - Linux under z/VM**
  http://www.ibm.com/developerworks/linux/linux390/perf/tuning_vm.html

- **z/VM System Management**
  http://www.vm.ibm.com/sysman/

- **Extreme Cloud Administration Toolkit (xCAT) - Overview (IBM)**
  http://www-03.ibm.com/systems/software/xcat/index.html

- **Cloud Computing with xCAT on z/VM (by Thang Pham)**
  http://www.vm.ibm.com/sysman/xcatinfo.pdf