IBM System z Technology Summit

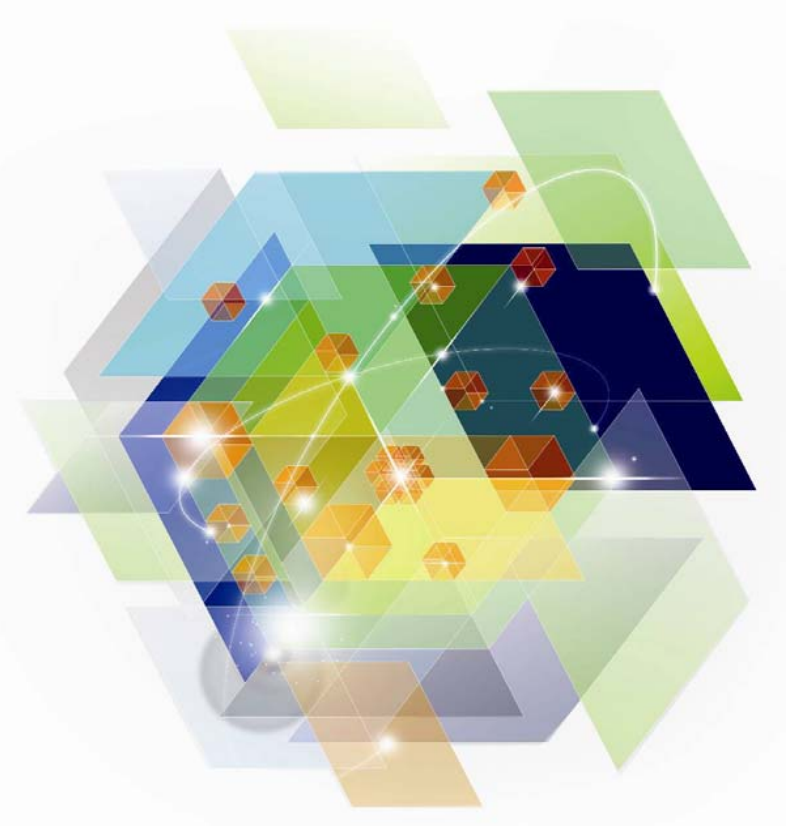

Creating an Infrastructure Tuned to the task with z/VM

Presenter –

Date –

Creating an Infrastructure Tuned to the Task with z/VM

System z + Virtualization + Linux = SYNERGY

- **Why IBM System z® with z/VM® is the ultimate virtualization resource**
- **Why everyone is talking about z/VM 6.2 and what are they saying**
- **A little bit about what's coming next**

System z with z/VM – The ultimate virtualization resource
Consolidation and Scalability

- **Massive consolidation platform**
  - Up to 60 logical partitions
  - 100s to 1000s of Linux® virtual servers under z/VM
  - HiperSockets™ for memory-speed communication
  - Most sophisticated and complete hypervisor function available – Virtualization that is built-in, not added-on
- **Scalability**
  - IBM zEnterprise™ 196 (z196) scales to 80 configurable processors
  - IBM zEnterprise 114 (z114) scales to 10 configurable processors
  - IBM System z10™ Enterprise Class (z10 EC) scales to 64 configurable processors
  - IBM System z10 Business Class (z10 BC) scales to 10 configurable processors
  - Dedicated I/O processors
    - Up to 14 (z196), up to 11 (z10 EC), or 2 (z10 BC and z114)
- **Meeting Customer Expectations for Service Levels**
  - The most reliable hardware platform available
  - I/O bandwidth
  - System availability

Consolidate 3x or more* servers per core than virtualized x86 offerings: spend less on software, energy, floor space and disaster recovery

* Based on IBM "Friendly Bank" workload scenario

System z with z/VM – The ultimate virtualization resource

Maximizing Resource Utilization

- **Software Hypervisor integrated in hardware**
  - Sharing of CPU, memory and I/O resources
  - Virtual network – virtual switches/routers
  - Virtual I/O (mini-disks, virtual cache, …)
- **Shared everything infrastructure through hardware allows for maximum utilization of resources**
  - CPU, Memory, Network, Adapters, Cryptography, Devices
- **Designed to support diverse mixed workloads – not just more of the same**
  - Allows consolidation while maintaining one virtual server per application
  - Complete work load isolation
  - High speed inter-server connectivity
- **Handles peak workload utilization of 100% without service level degradation**
  - Utilization often (usually?) exceeds 90%

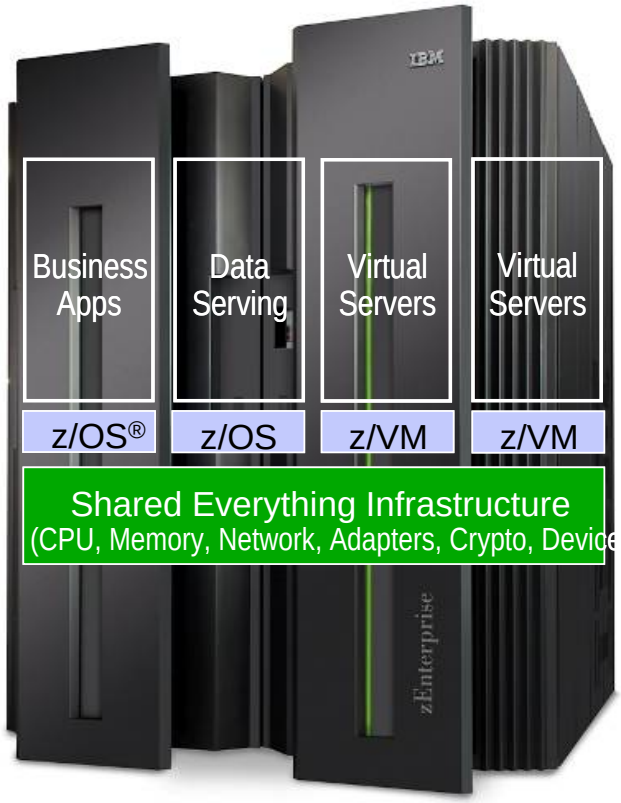

Do more with less

* Based on IBM “Friendly Bank” workload scenario

System z with z/VM – The ultimate virtualization resource

Hypervisor and Virtual Server Management

- **z/VM facilities for life cycle management of virtual servers are extensive**
  - Provisioning, automation, monitoring, workload management, capacity planning, security, charge back, patching, backup, recovery, …
- **Self-optimizing workload management**
- **Centralized Linux systems can be easier to manage**
- **Intelligent and autonomic management of diverse workloads and system resources based on business policies and workload performance objectives**

Manage more server images with fewer people

System z and z/VM – The ultimate Virtualization Resource

Flexible, Granular, Efficient growth of Virtual Servers

- **Smart economics: non-disruptively scale the z/VM environment by adding hardware assets that can be shared with every virtual server**
  - Start small with Linux on System z and non-disruptively grow your environment as business dictates
  - Add physical resources dynamically without taking the system down - add CPUs, memory, I/O adapters, devices, and network cards to a running z/VM LPAR
  - Highly granular resource sharing (<1%)
  - Capacity-on-demand capability on System z
- **More virtual servers per core; share more physical resources across servers**
- **z/VM virtualizes the additional resources for the Linux guest machines**

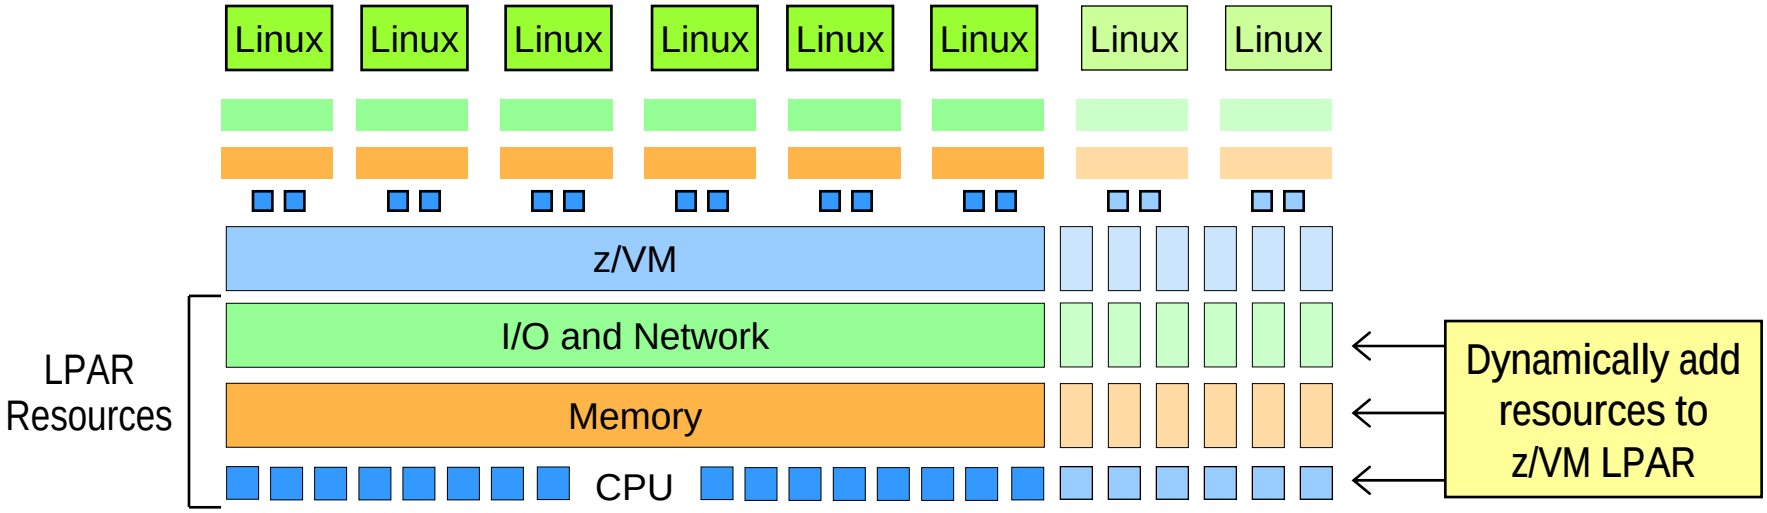

Deploy new servers and applications faster – in seconds instead of hours or days

Resource Utilization of Resources

Distributed Platforms

- Limited per-core virtual server scalability
- Physical server sprawl is needed to scale
- Operational complexity increases as virtual server images grow
- VMware, HyperV, Xen, KVM, Virtual Iron focus only on x86, not cross platform hardware management

Mainframe Platforms

- Achieve nearly 100% utilization nearly 100% of the time
- Reduced TCO
  - Environmental Savings with single footprint
  - Consolidation savings
  - Less physical resources to manage
  - Fewer software licenses
  - Centralized management of virtual servers

Moderate distributed servers

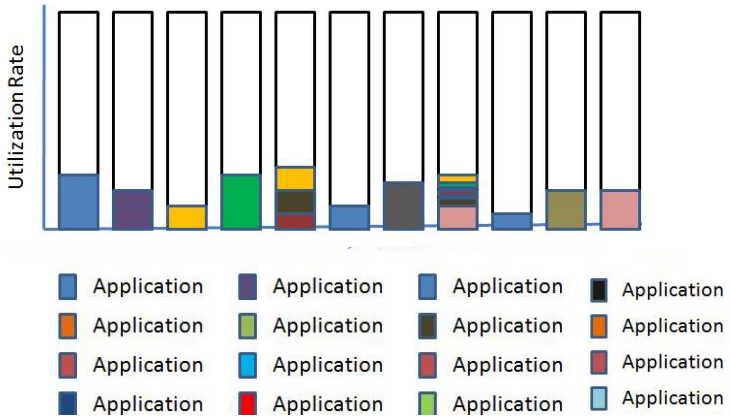

Typically single application per physical server

Up to 100% utilized System z server

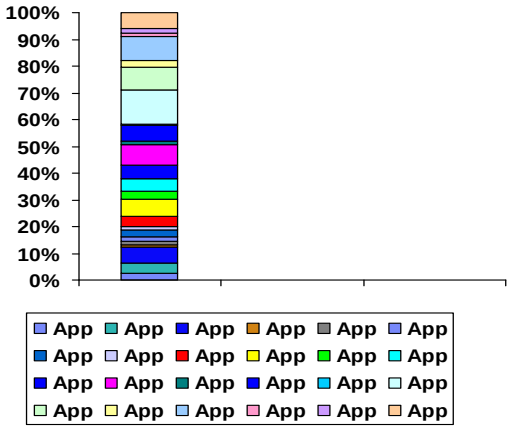

Multiple applications on one physical System z server

Reduce costs on a bigger scale

[1] Source: gomainframe.com Joe Clabby

System z Integrated Facility for Linux (IFL)

What you need to know about the specialty engine enabled for Linux and z/VM

- **IFL is a standard CP engine with different microcode**
- **Inherits same hardware RAS as GCP engines**
- **Always runs at "full speed" even on sub-capacity systems**
- **No impact on z/OS software stack cost**
- **Sold at a lower price than standard GCP engines**
- **Leverages incremental capacity on existing systems**
  - Exception: Enterprise Linux Server
- **Enables cost effective collocation benefits with z/OS applications and data**

Built-in IBM System z security for distributed workloads

Reducing risk – protecting businesses

- **Protect from intrusion**
  - z/VM Integrity Statement
- **Protect data**
  - Built in encryption accelerators in every server
  - FIPS-140-20 Level 4 certified encryption co-processors for highly secure encryption
- **Ensure privacy**
  - Access to all resources is controlled by an integrated central security manager
- **Respond to compliance regulations**
  - Up to 70% in security audit savings
- **Protect virtual servers; System z is the only commercial platform with EAL5 security classification**
  - The IBM zEnterprise 196 (z196) server achieved a Common Criteria certification at an EAL5+ level.
  - Certification is listed on BSI's website*

* https://www.bsi.bund.de/ContentBSI/EN/Topics/Certification/newcertificates.html

Top Drivers for Running Additional Workloads On Linux on the Mainframe

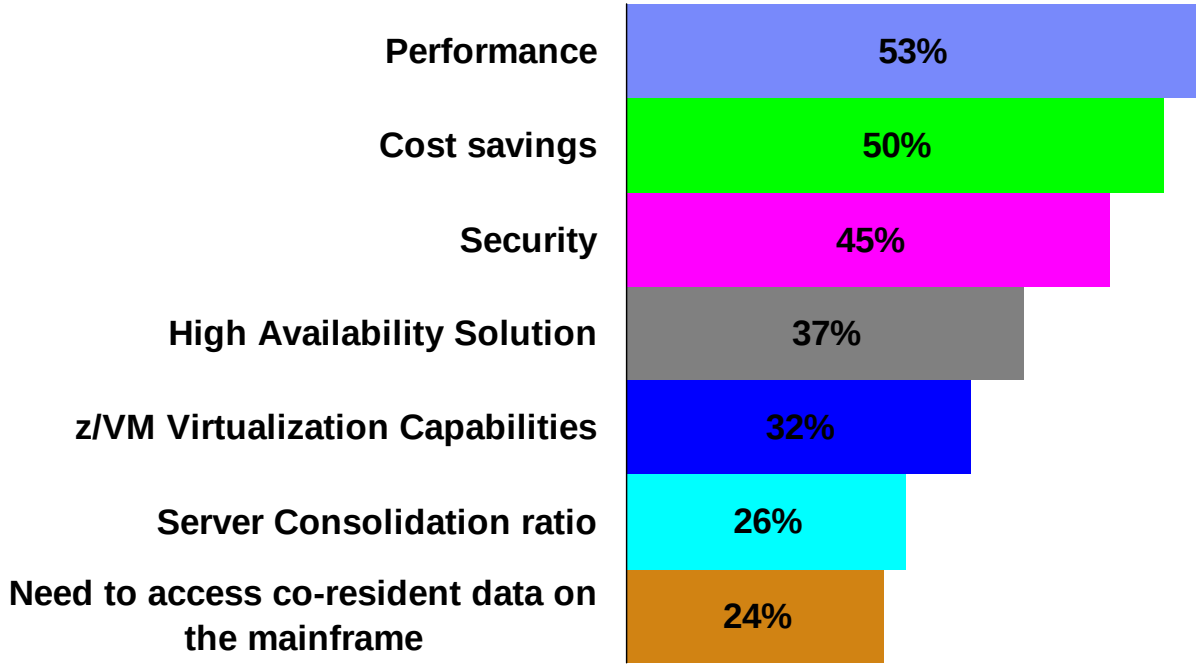

- **Performance, cost savings, and security** are the 3 most significant factors that influence decisions to run additional workloads on Linux on the Mainframe

Source: 2011 IBM Market Intelligence

Customer Perspective

> *"From every perspective, running applications under Linux on System z makes sense for our organization. Performance, reliability, disaster recovery, server provisioning and cost efficiency have all seen dramatic improvements – helping us deliver better service and value to our members across the state."*

Ted Mansk, Director of Infrastructure Engineering and Databases at Blue Cross Blue Shield of Minnesota

z/VM Version 6
Foundation for future virtualization growth

- **z/VM V6.1 is the base for all future z/VM enhancements**
  - This release implements a new Architecture Level Set available only on the IBM System z10 and zEnterprise servers, and future generations of System z servers
  - Includes several enhancements, plus support for the zEnterprise Unified Resource Manager
- **z/VM V6.2 accelerates the journey**
  - IBM's virtual machine facility has a very long track record of supporting production workloads
  - The major improvement offered by version 6.2 is support of clustering for up to four z/VM instances hosted on LPARs which can be hosted on a single machine or distributed over several machines
  - Linux guest environments can be relocated from one LPAR to another without disruption of the workloads being supported

z/VM Version 6 Release 2

Accelerate the journey to smarter computing with multi-system virtualization and virtual server mobility

Features:	Benefits:
▪ Multi-system virtualization allows up to 4 z/VM instances to be clustered, serviced, and administered as a Single System Image (SSI)	▪ Relief from the challenges associated with virtual machine sprawl on competitive systems
▪ Live Guest Relocation (LGR) moves running Linux virtual servers without disruption to the business	▪ Helps clients avoid planned outages for virtual servers when performing maintenance
▪ Provides a set of shared resources for the z/VM systems and their hosted virtual machines	▪ Provides a more manageable infrastructure for cloud computing
▪ Scales up to four systems horizontally, each with up to 32 CPUs and 256 GB memory	▪ Improved systems management to help manage the life cycle of the z/VM hypervisors and the virtual servers
▪ High server consolidation ratio with support for more virtual servers than any other platform in a single footprint	▪ Enhanced workload balancing with the added ability to move work to available resource in addition to long standing capability to move system resources to work
	▪ Rapid deployment of Linux virtual servers for under $1.45 per day on IBM zEnterprise System

Single System Image feature
Clustered hypervisor with Live Guest Relocation

- **Provided as an optional priced feature**
- **Connect up to four z/VM systems as members of a Single System Image (SSI) cluster**
- **Provides a set of shared resources for member systems and their hosted virtual machines**
- **Cluster members can be run on the same or different System z servers**
- **Simplifies systems management of a multi-z/VM environment**
  - Single user directory
  - 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
  - Built-in cross-member capabilities
  - Resource coordination and protection of network and disks

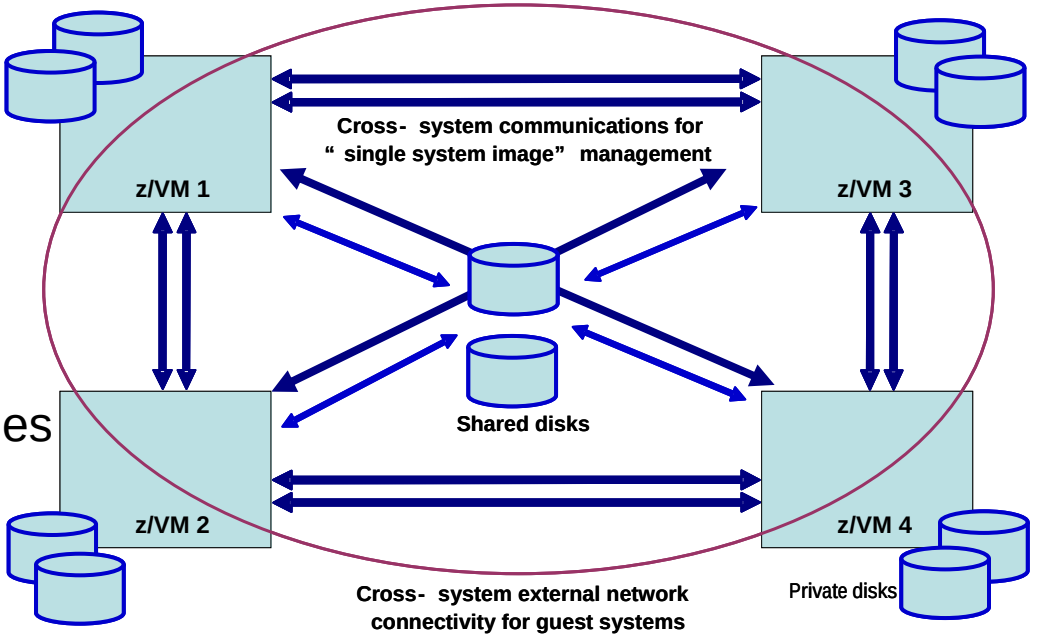

Single System Image feature
Clustered hypervisor with Live Guest Relocation

- **Dynamically move Linux guests from one member to another with Live Guest Relocation**
  - Reduce planned outages
  - Enhance workload management
  - Non-disruptively move work to available system resources and non-disruptively move system resources to work
- **When combined with Capacity Upgrade on Demand, Capacity Backup on Demand, and Dynamic Memory Upgrade, you will get the best of both worlds**
  - Bring additional resources to the workload!
  - Move the workload to the resources!

SSI Cluster Management – Features for Greater Reliability

- **Cross-checking of configuration details as members join cluster and as resources are used:**
  - SSI membership definition and identity
  - Consistent definition of shared spool volumes
  - Compatible virtual network configurations (MAC address ranges, VSwitch definitions)
- **Cluster-wide policing of resource access:**
  - Volume ownership marking to prevent dual use
  - Coordinated minidisk link checking
  - Autonomic minidisk cache management
  - Single logon enforcement
- **Communications failure “locks down” future resource allocations until resolved**
- **Comprehensive checking for resource and machine feature compatibility during relocation:**
  - Adjustment of “virtual architecture level” to support customer relocation policy

Early Client Feedback on z/VM 6.2[1]

BALDOR
A MEMBER OF THE ABB GROUP

- 26 manufacturing facilities WW
- High-quality industrial electric motors, mechanical power transmission products, drives and generators
- 15 percent growth per year over the last 20 years.
- IT costs have now been reduced to less than one percent of revenue
- Core SAP landscape runs on an z10 EC (6 CPs, 3 zIIPs, 16 IFLs). DB2® runs in z/OS partitions. 70 z/VM guests provide Linux environments for SAP application servers [2]
- Mark Shakelford, VP of Information Services: "*This implementation will provide many usability enhancements that increase our productivity and let us focus on customer projects, versus focusing on internal projects, with continued improvement in high availability. LGR is the very best z/VM software enhancement since 64-bit support was available.*"

Banking Industry
Major US financial firm

- One of the largest companies in the US, with operations around the world
- One of the largest banks by deposits, assets and by market capitalization
- Home mortgage servicing and debit cards in the US.
- The company has grown to 6,335 retail branches 12,000 automated teller machines, 280,000 employees and over 70 million customers
- The data center's operating systems engineer states that the ability to provide uninterrupted reliable service is critical.
- The bank's IT executive says, "z/VM allows ATM operations without interruption, which is critical to our maintaining customer satisfaction…. we find that the advantages of z/VM 6.2 features will provide a solid basis for the future."

[1] Edison Whitepaper: Using Linux on IBM z/VM

[2] Baldor consolidates hundreds of servers and cuts IT and energy cost

z/VM Pricing in the US

- **z/VM pricing is an OTC price per value unit and an S&S annual support charge price per value unit:**
  - Price per value unit charge metric is used on z/VM V5 and z/VM V6
- **The z/VM price schedule below applies to V5 and V6 (including V6.2)**
  - The SSI feature includes LGR and it is priced in line with the RACF feature
- **The z/VM value unit schedule below helps illustrate how the number of engines translates into VUs to be used in pricing:**

z/VM Price Schedule per Value Unit		
Feature	**OTC**	**S&S**
Base	**$2,250**	**$563**
RSCS	**$1,000**	**$250**
Performance Toolkit	**$175**	**$44**
RACF	**$300**	**$75**
Directory Maintenance	**$38**	**$10**
Total Suite	**$3,763**	**$942**
New Feature in V6.2 – Announce 10/12/11		
SSI (incl. LGR)	**$300**	**$75**

z/VM Value Unit Schedule	
Number of Engines	**VUs per Engine**
1 to 3	**10**
4 to 6	**9**
7 to 9	**8**
10 to 12	**7**
13 to 16	**6**
17 to 20	**5**
21 to 25	**4**
26 and above	**3**

Putting zEnterprise System to the task

Use the smarter solution to improve your application design

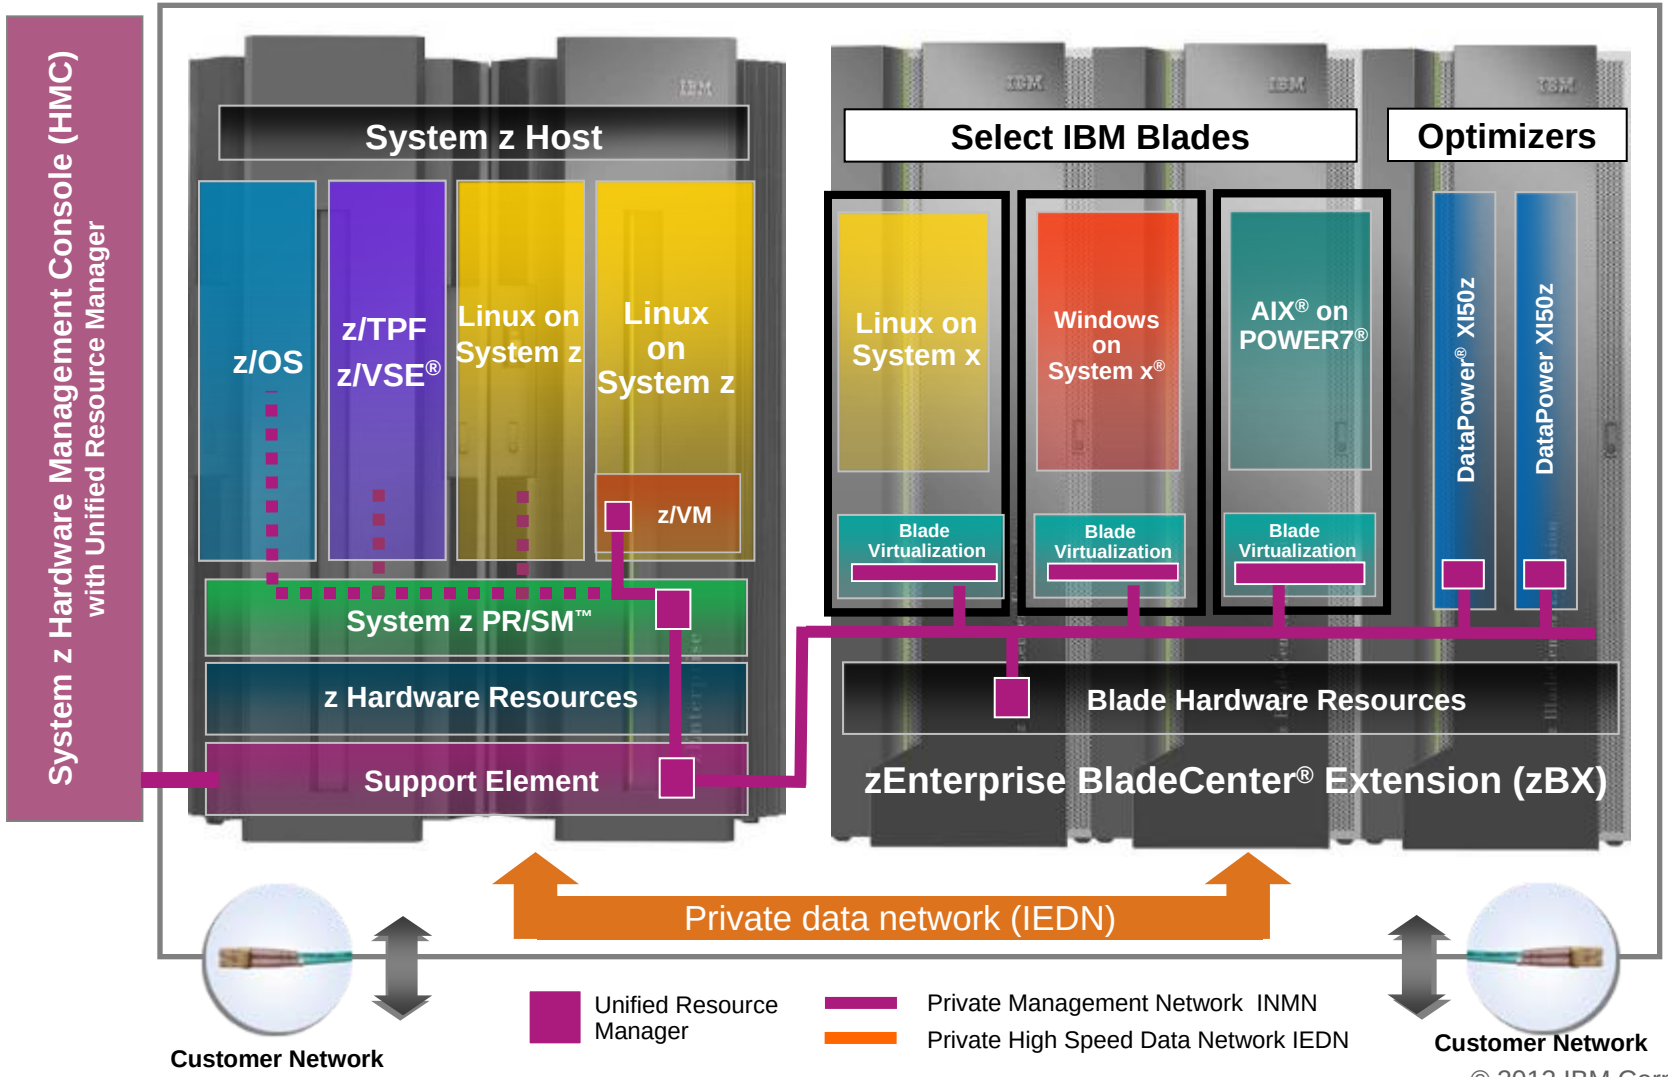

Single System Image feature

Clustered hypervisor with Live Guest Relocation

- **Unified Resource Manager does not support SSI and LGR**
- **IBM Systems Director does not support SSI and LGR**
- **Suggested best practice is to not combine SSI and LGR with the above offerings**
  - Work with your IBM Sales Team, Business Partner, IBM Lab Services, or z/VM Development Lab to determine which technologies are most critical to your environment and business

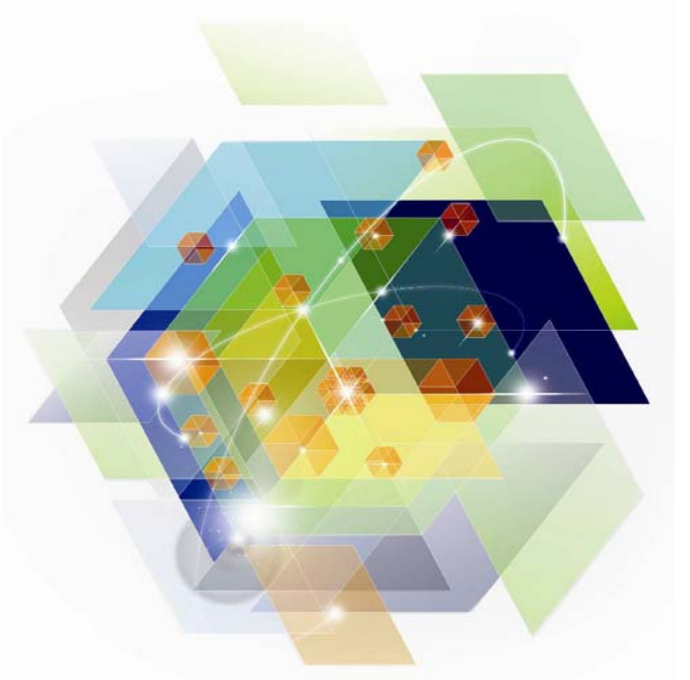

zEnterprise Unified Resource Manager
VM64822, VM64904, VM64917, VM64956, VM64957

- **z/VM V6 only**
  - Turnkey installation option to enable virtual server management via zEnterprise Unified Resource Manager (z/VM V6.2 only)
- **Enables Unified Resource Manager to perform system and virtual server management tasks**
  - Virtual server configuration
  - Disk storage management
  - Virtual network management
  - Performance monitoring
- **CP, CMS, LE, TCP/IP, DirMaint™, Performance Toolkit for VM, HCD**
- **http://www.vm.ibm.com/service/vmrequrm.html**

APIs: Enable Management Tools

- **API allows programmatic access to the same underlying zManager function as is accessed via the HMC UI**
  - Same resource types, instances and policy
  - API functions corresponding to views and tasks in the UI
    - Listing resource instances
    - Creating, changing, deleting resource instances
    - Operational control of resource instances
    - etc.
  - Requests and responses structured as RESTful Web services
  - Status and property changes delivered via JMS
- **Enables ensemble management from external (to HMC) tools**
  - Tivoli® Director
  - 3rd party tools
- **Web-services orientation facilitates access from scripting environments (e.g. Python) that can be used to develop user-written tools and automation**
  - New opportunity for automating z/VM management using modern scripting languages

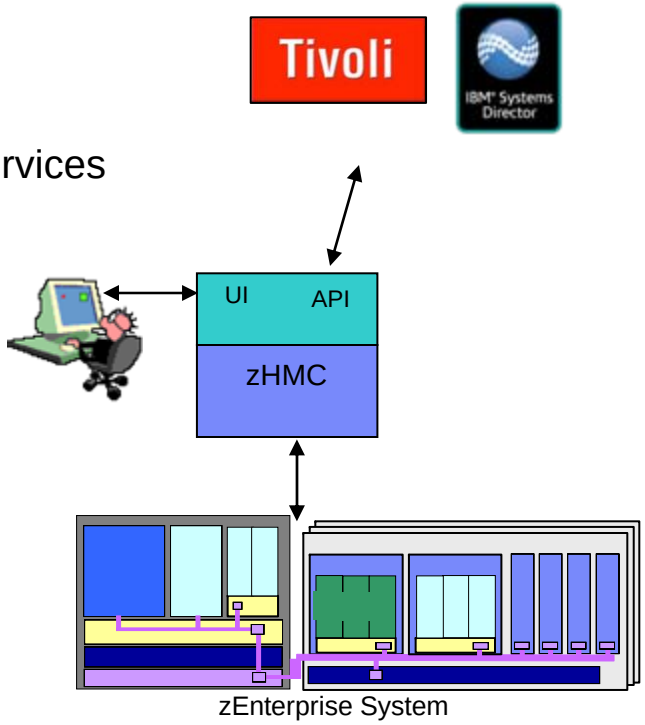

HiperSockets VSWITCH integration with zEnterprise IEDN

z/VM Statement of Direction: Delivered April 13, 2012!

- **Virtual Switch bridge between Ethernet LAN and HiperSockets**
  - zEnterprise IEDN (OSX) connections
  - Guests can use simulated OSA or dedicated HiperSockets
  - VLAN aware
  - One HiperSockets CHPID only
- **Full redundancy**
  - Up to 5 bridges per CEC
  - One bridge per LPAR
  - Automatic takeover
  - Optionally designate one “primary”
    - Primary will perform “takeback” when it comes up
  - Each bridge can have more than one OSA uplink

High Performance FICON

z/VM Statement of Direction: Delivered April 13, 2012!

- **Enable guests to use High Performance FICON for System z (zHPF)**
  - Different I/O model
  - Single and multiple track I/O
- **Requires host and control unit compatibility**
  - Consult a storage specialist for details
- **z/OS and Linux provide exploitation**

Support for GDPS/PPRC 3.8
z/VM Statement of Direction: New function

- **Disk subsystem preemptive HyperSwap™**
  - Storage controllers will notify host when failure is predicted
  - HyperSwap before I/O errors are generated
- **HyperSwap scalability**
  - Summary “PPRC Suspend” event notification by storage controller
  - Avoid separate notification for each disk
- **Future z/VM release support for an alternate subchannel set to place PPRC secondary devices**

System z and z/VM Tuned to the Task

- **Consolidate 3x or more* servers per core than virtualized x86 offerings: spend less on software, energy, floor space and disaster recovery**
  - Reduce costs on a bigger scale
  - Consolidation on a “scale up” machine like Linux on zEnterprise means fewer cables, fewer components to impede growth
  - Delivers on Service Level Agreements
- **Do more with less**
  - Consolidate more servers, more networks, more applications, and more data
  - 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
- **Manage more server images with fewer people**
  - Exploit extensive z/VM facilities for life cycle management: provisioning, monitoring, security, workload mgmt, capacity planning, charge back, patching, backup, recovery,
  - Single System Image
- **Deploy new servers and applications faster – in seconds instead of hours or days**
  - Add hardware resources to an already-running system without disruption
  - Consolidating to a single zEnterprise server offers significant advantages in terms of flexibility
  - Rapid provisioning reduces lead time for new IT projects, helping to increase business agility

Questions?

Back Up

Overview of z/VM Releases

z/VM Level		GA	End of Service	End of Marketing	Minimum Processor Level	Security Plan
Ver 6	Rel 2	12 / 2011	4 / 2015	TBD	z10	-
	Rel 1	10 / 2009	4 / 2013	12 / 2011	z10	EAL 4+ *1
Ver 5	Rel 4	9 / 2008	9 / 2013	3 / 2012	z900 / z800	-
	Rel 3	6 / 2007	9 / 2010	9 / 2010	z900 / z800	EAL 4+

*1 Under evaluation

Scalability and performance enhancements
Available by PTF to prior releases where shown

- **Reduction of memory and CPU resources required to manage larger memory sizes**
- **Control of the guest page re-ordering process, improving the performance characteristics of guests with large memory footprints (VM64774)**
- **Reduced system overhead of guest page release function, thereby helping to increase guest throughput (VM64715)**
- **Improved contiguous frame coalescing algorithms help to increase system throughput (VM64795)**
- **More accurate scheduling algorithm for guests that have LIMITHARD shares (VM64721)**
- **Reduce LPAR suspend time by reducing the number of DIAGNOSE 0x9C and 0x44 instructions issued when obtaining system locks (VM64927 for z/VM 6.1 only)**
- **Improve workload dispatch algorithm to eliminate erratic virtual machine pause in busy systems with more than 14:1 total virtual to logical CPU over-commitment (VM64887)**

TCP/IP enhancements
SSL server upgrade and FIPS compliance

- **Upgraded to z/OS R12 System SSL technology**
- **Can now be configured for FIPS 140-2 mode of operation**
  - Validation and certification activities are underway
  - AES validation complete

    http://csrc.nist.gov/groups/STM/cavp/documents/aes/aesval.html#1712
- **Digital signature applied to the key database**
- **System SSL library is verified prior to loading**

RACF Security Server

- **Single System Image**
  - Automatic propagation of most RACF® commands
  - Also works with multiple RACF servers on same z/VM system
- **Protected Users**
  - User without a password or password phrase will not be revoked due to too many invalid password attempts or inactivity
- **High Level Assembler no longer required for most common customizations**

LDAP server upgrade
Tivoli Directory Server

- **z/OS R12 level**
- **Management and change logging of general resources**
- **Password management policy support to improve LDAP authentication from open systems such as Linux**
  - Expiry warnings
  - Interactive password change when password has expired
  - Password rule validation

Additional z/OS R12 equivalency upgrades

- **Language Environment® (LE) runtime libraries**
- **Program Management Binder**
  - COMPAT supports ZOSV1R10, ZOSV1R11, ZOSV1R12
  - New sub options on RMODE
  - Compiler parameters can be read from IEWPARMS DDNAME
  - New C/C++ API
- **Support for IBM XL C/C++ Compiler for z/VM, V1.3 (5654-A22)**
  - Details can be found in IBM US Announcement 211-369 or IBM Canada Announcement A11-0787
- **MPROUTE**

z/CMS

- **Previously shipped with z/VM as a sample program, now supported as an optional CMS**
  - IPL ZCMS
- **Enables CMS programs to use z/Architecture® instructions and 64-bit registers**
- **Existing ESA/390 architecture programs continue to run unchanged**
  - CMS not exploit memory above 2 GB
  - CMS does not provide memory management API for memory above 2 GB
- **Programs that examine or change architecture-sensitive memory locations (NUCON) must be updated in order to use z/CMS**
- **No architectural support for XC mode**
  - VM Data Spaces not available

Installation improvements

- **Significant changes to system layout to support Single System Image**
- **Choose a non-SSI system or a complete 1- to 4-member SSI cluster**
  - First or second level
- **All installation information is gathered at one time**
- **All DASD volumes can be labeled at installation time, including the system residence volume**
- **Turnkey support for zEnterprise ensembles**
  - Enable clients new to z/VM to get started with Unified Resource Manager
  - Those who purchase DIRMAINT or another directory manager, or who require an external security manager, need to perform manual enablement
    - Decline this option during installation

Previously shipped enhancements included in z/VM V6.2

- XRC timestamps via *VM64814 and VM64816*
- HyperSwap improvements via *VM64815 and VM64816*
- SSL Server Reliability and Scalability via *PK97437, PK97438, PK75662*
- CPU Measurement Facility counter host support
- zEnterprise Unified Resource Manager
- APAR numbers shown apply to z/VM 6.1 and z/VM 5.4 unless otherwise stated

CPU Measurement Facility (MF) counters
VM64961

- **CPU MF counters are a System z hardware facility that characterizes the performance of the CPU and nest**
  - Instructions, cycles, cache misses, and other processor related information
- **Available on z10 EC/BC, z196, and z114**
- **The CPU MF counter values:**
  - Help IBM to understand how your workload stresses a CEC for future design
  - Help IBM to map your workload into the LSPR curves for better sizing results
  - Help IBM better understand your system when there is a processor performance related problem.
- **z/VM 6.2, 6.1, and 5.4 can all collect the CPU MF counters from the hardware**
  - Counters are put in new z/VM monitor record