



IBM Software

# Virtualization Management on System z

Speaker Name: Mac Holloway  
([mhollowa@us.ibm.com](mailto:mhollowa@us.ibm.com))

# Trademarks and disclaimers

- Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries./ Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both. IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce. ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office. UNIX is a registered trademark of The Open Group in the United States and other countries. Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others. Information is provided "AS IS" without warranty of any kind.
- The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.
- Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.
- All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
- Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.
- Prices are suggested U.S. list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.
- Photographs shown may be engineering prototypes. Changes may be incorporated in production models.
- © IBM Corporation 1994-2009. All rights reserved.
- References in this document to IBM products or services do not imply that IBM intends to make them available in every country.
- Trademarks of International Business Machines Corporation in the United States, other countries, or both can be found on the World Wide Web at <http://www.ibm.com/legal/copytrade.shtml>.

## Acknowledgements


- **The following people contributed material (knowingly and unknowingly) to this presentation**
  - Bill Bitner
  - Bill Maddox
  - Mike Sine
  - Jon vonWolfersdorf
  - Kevin Yash
  - Bob Neill

# Agenda

- **Virtualization**
  - What is it ?
  - Why System z ?
- **System Management**
  - An incremental approach
  - Basic health – what are you using ?
- **Better Health**
  - What are you waiting on ?
- **Summary**

**Roles:**

- Consolidations
- Dynamic provisioning / hosting
- Workload management
- Workload isolation
- Software release migration
- Mixed production and test
- Mixed OS types/releases
- Reconfigurable clusters
- Low-cost backup servers

**Possible Benefits:**

- High resource utilization
- Great usage flexibility
- Enhanced workload QoS
- High availability / security
- Low cost of availability
- Low management costs
- Enhanced interoperability
- Legacy compatibility
- Investment protection

In the final analysis, the potential virtualization benefits take three forms:

- **Help reduce hardware costs**
  - Help increase physical resource utilization
  - Small footprints
- **Can improve flexibility and responsiveness**
  - Virtual resources can be adjusted dynamically to meet new or changing needs and to optimize service level achievement
  - Virtualization is a key enabler of on demand operating environments
- **Can reduce management costs**
  - Fewer physical servers to manage
  - Many common management tasks become much easier

**Can reduce Energy costs! Can reduce Energy costs! Can reduce Energy costs!**

# System z Virtualization Technology

## A Shared Everything Architecture

### Start Interpretive Execution

- Establish architecture for guest systems
- Maintain status
- Invoke SIE assists

PR/SM – SIE – EAL 5

### LPAR Zoning:

each partition has a zero-origin address space, allowing I/O access to memory without hypervisor intervention

Hardware support: 10% of circuits are used for virtualization

LPAR – Up to 60 Logical Partitions

### Most sophisticated and functionally complete hypervisors

- Able to **host** z/OS, Linux, z/VSE, z/TPF, and z/VM-on-z/VM
- Shared **everything** architecture
- Highly **granular** resource sharing (less than 1% utilization)
- Any virtual CPU can **access** any virtual I/O path within the attached logical channel subsystem
- z/VM can **simulate** devices not physically present
- Application **integration** with HiperSockets and VLANs
- Intelligent and **autonomic** workload management

### Shared resources per mainframe footprint

- Up to **64** OS-configurable CPUs
- Up to **11** SAP processors
- Up to **1.5** TB of memory
- Up to **1024** channel paths
- Up to **16** internal HiperSockets networks

z/VM – SIE – EAL 4+ – 100s of Virtual Machines – Shared Memory

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

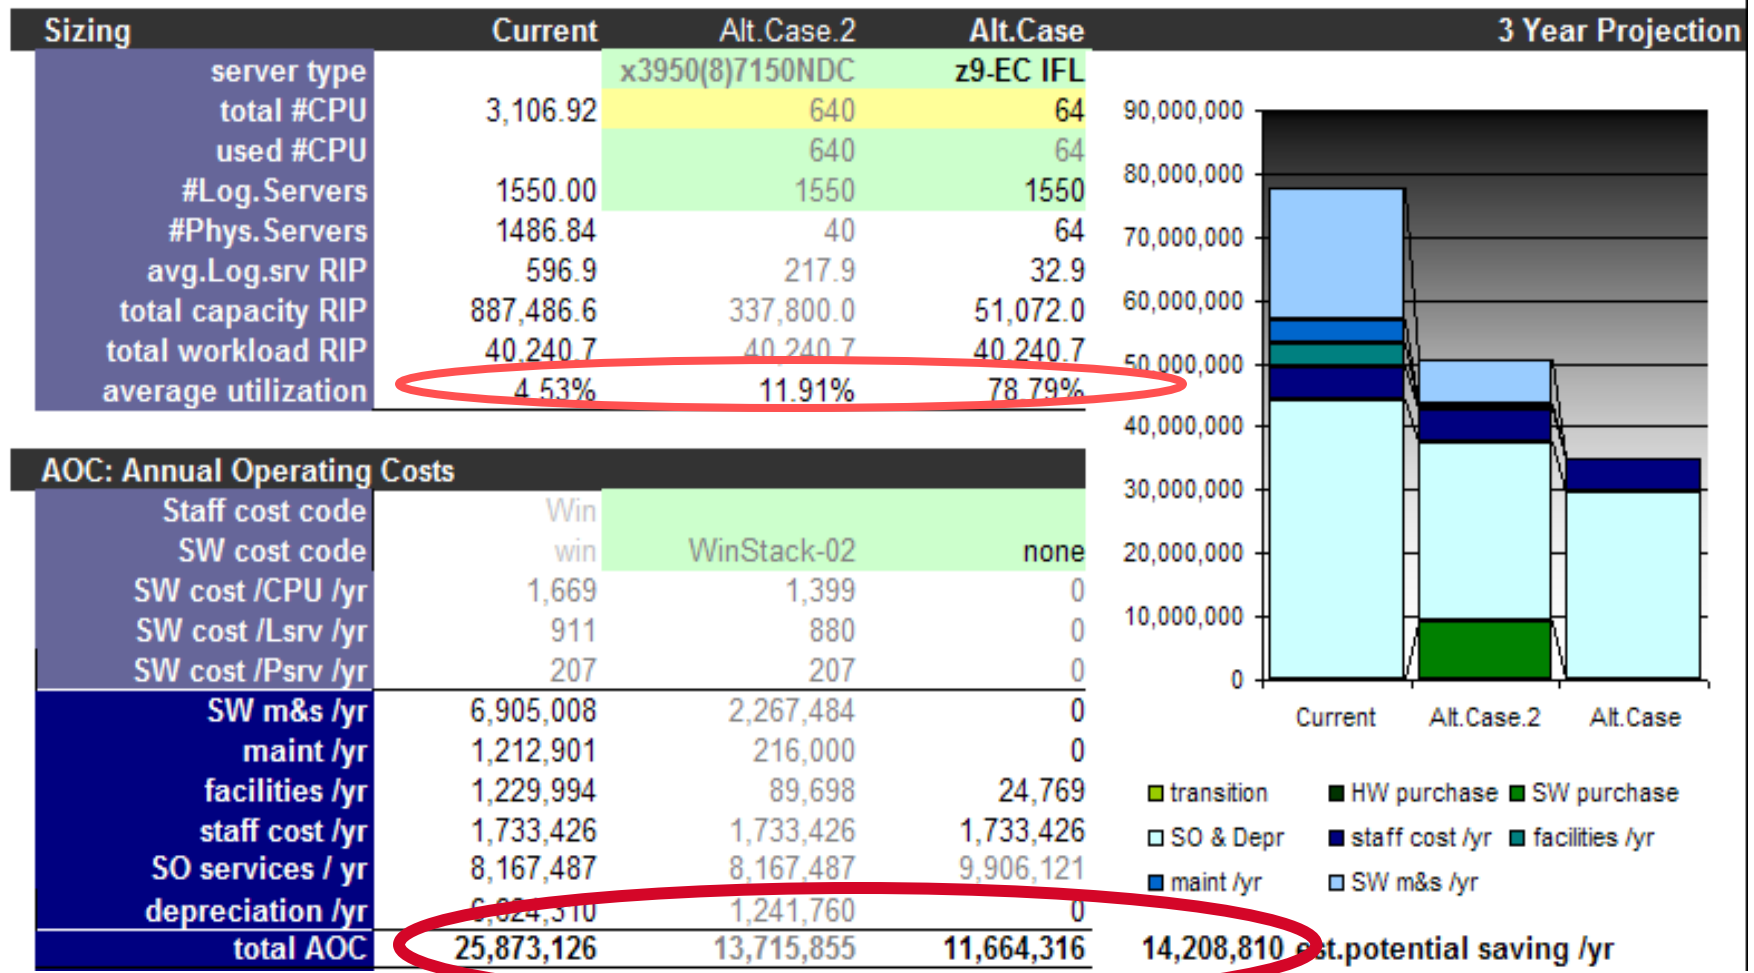
- Hardware support, SIE, microcode assist
- Virtualization is transparent for Op Sys execution
- Hardware-enforced isolation**

The potential performance impact of the Linux server farm is isolated from the other LPARs

These two alternatives (VMware and Mainframe) differ in level of risk and estimated transition cost, but both are favorable from a business perspective and address the majority of workloads at XYZ.

## 101: win-app

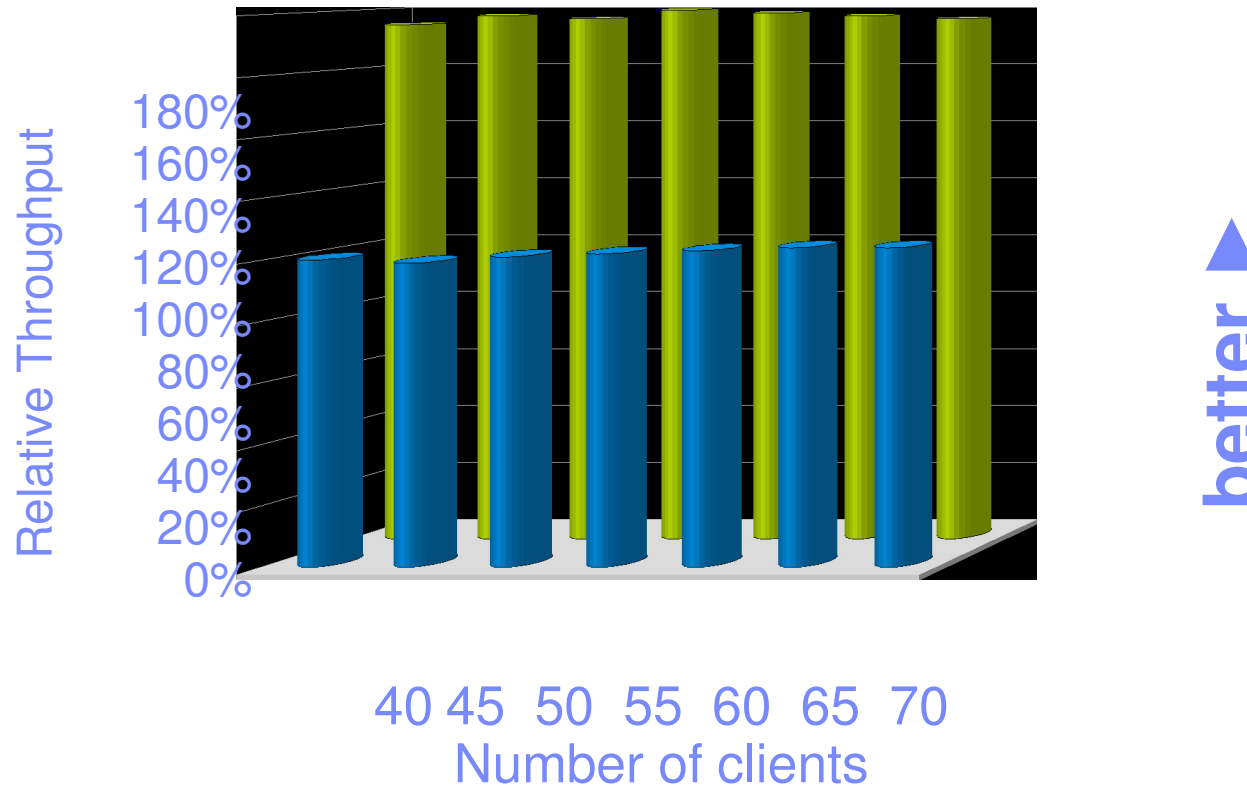
Actual



# WebSphere Application Server Cluster - Result

## Throughput - z10 versus z9

### Workload Scaling



**IBM System z10 provides constantly about 80% higher throughput!**



**WebSphere/AIM (18/13)+2**

- CICS Transaction Gateway
- IBM Java SE
- WebSphere Application Server
- WebSphere Application Server ND
- WAS Community Edition
- WebSphere Extended Deployment
- WS Business Events
- WS Business Modeler Publishing Server
- WS Business Monitor
- WS Business Services Fabric
- WS Commerce
- WS Enterprise Service Bus
- WS Message Broker
- WS MQ
- WS MQ File Transfer Edition
- WS Process Server
- WS Service Registry & Repository
- WS sMash

I35

**Information Management (13/10)+2**

- Alphablox
- Cognos 8 Business Intelligence
- DataQuant
- DB2 for LUW
- DB2 Connectit
- **Content Integrator**
- DB2 Content Manager
- DB2 Content Manager OnDemand
- Filenet P8
  - Content Manager
  - Records Manager
- Information Server
  - WS DataStage
  - WS Federation Server
  - WS QualityStage
- Informix Dynamic Server
- InfoSphere Master Data Management Server
- InfoSphere Warehouse for z/OS

I34

**STG & GTS (5/0)**

- IBM Active Energy Manager
- **IBM VMControl – Image Manger**
- IBM Systems Director
- IBM Integrated Removable Media Manager
- z/VM

**Lotus/Workplace (3/3)+3**

- Domino
- WS Portal
- **Lotus Forms**

**Rational (4/0)**

- Asset Manager
- BuildForge
- ClearCase
- Team Concert for System z

**Tivoli (32/27)+2**

**Service and Process Automation (SPA)**

- Application Dependency Discovery Manager
- Business Continuity Process Manager
- Change & Configuration Management Database
- Dynamic Workload Broker
- Provisioning Manager
- + Provisioning Manager for SW (TCM)
- **Provisioning Manager for OS Deployment**
- Release Process Manager
- Service Request Manager
- **Service Automation Manager**
- System Automation
- + System Automation Application Manager
- Workload Scheduler

**Service, Avail. and Perf Mgmt (SAPM)**

- Business Systems Manager
- Composite App Manager
  - for Appl Diagnostics
  - Web Resources
  - WebSphere
  - for SOA
- Monitoring
  - for Applications
  - for Databases
- NetView
- OMEGAMON XE for z/VM and Linux

**Security Mgmt**

- Access Manager
- Director Integrator
- Director Server
- Federated Identity Manager
- Identity Manager

**Network Mgmt**

- Enterprise Console
- NetCool/Omnibus
- Network Manager IP Edition

**Asset Management for IT**

- **Asset Management**
- Maximo Asset Management
- License Compliance Manager
- Usage and Accounting Manager

I36

**Storage Mgmt**

- Storage Manager
- TPC Agent

## Slide 9

---

- I34** Information Server also consists of:
- WS Business Glossary
  - WS Information Analyzer
  - WS Information Services
  - WS Metadata Server
  - WS Replication Server
  - Metadata Workbench

IBM\_USER, 12/3/2008

- I35** Compute Grid
- WS eXtreme Scale
  - WS Virtual Enterprise

IBM\_USER, 12/11/2008

- I36** Maximo Asset Management products consist of:
- Maximo Asset Management
  - Maximo Asset Management Essentials
  - Maximo Asset Navigator
  - Maximo Change and Corrective Action Manager
  - Maximo SLA Manager
  - Maximo Adapter for Microsoft Project
  - Maximo Enterprise Adapter
  - Maximo Asset Management for Service Providers
  - Maximo Mobile Work Manager
  - Maximo Mobile Inventory Manager
  - Maximo for Oil and Gas
  - Maximo for Government
  - Maximo Asset Configuration Manager

IBM\_USER, 10/1/2009

expected in 2H2009 or beyond

**WebSphere/AIM**

WAS Hypervisor (HV) Edition and WS CloudBurst Appliance  
z/VM support being pursued for 4Q09

**Information Management**

- InfoSphere Warehouse for z/OS
  - RHEL 5 support expected in July
- Cognos
  - 8.4 BI zNext now available in June "silent" refresh:
    - Go! Search
    - Virtual View Manager
    - DB2 z/OS Cubing Services
    - SQL Server
    - Additional data sources:
      - VSAM/IMS (Classic Fed)
      - Datacom/Adabas (Classic Fed)
      - Teradata (via ODBC)
    - IPv6 support
  - 64-bit product targeted for 2H2009

**STG & GTS**

- Systems Director VMControl Image Manager V2.1 (GA'd 7/28/09)
- GPFS (2Q-> 4Q)
- ISS Proventia Server (2Q-> 3Q)

**Lotus/Workplace**

Lotus Forms  
3.5.1 eGAs 8/21/09  
ST Server Standard 8.5.1 \*\*  
Target 4Q09 -> 2Q10  
Connections & Quickr  
Target 2Q10

**Rational**

- Rational Agent Controller V8.2
  - 64-bit & 31-bit support for WAS 7.0 eGA'd 7/27/09 for SLES 10 and RHEL 4
  - RHEL 5 and SLES 11 expected later in 3Q09.

**Tivoli**

**Service & Process Automation**

- TPMfOSD
  - 7.1.1 eGAs 8/21/09 with Server support for SLES 10 and RHEL 5
- TPM
  - 7.1.1 eGA 8/28/09 with Server support for SLES 10 and RHEL 5
  - RHEL 5 Agent support in 4Q09 \*\*

• TSAM

- Manage from support for SLES 10 & SLES11 plus Deploy of RHEL 5 Images targetted for 4Q09

**Service, Avail, & Perf Mgmt**

**Security Mgmt**

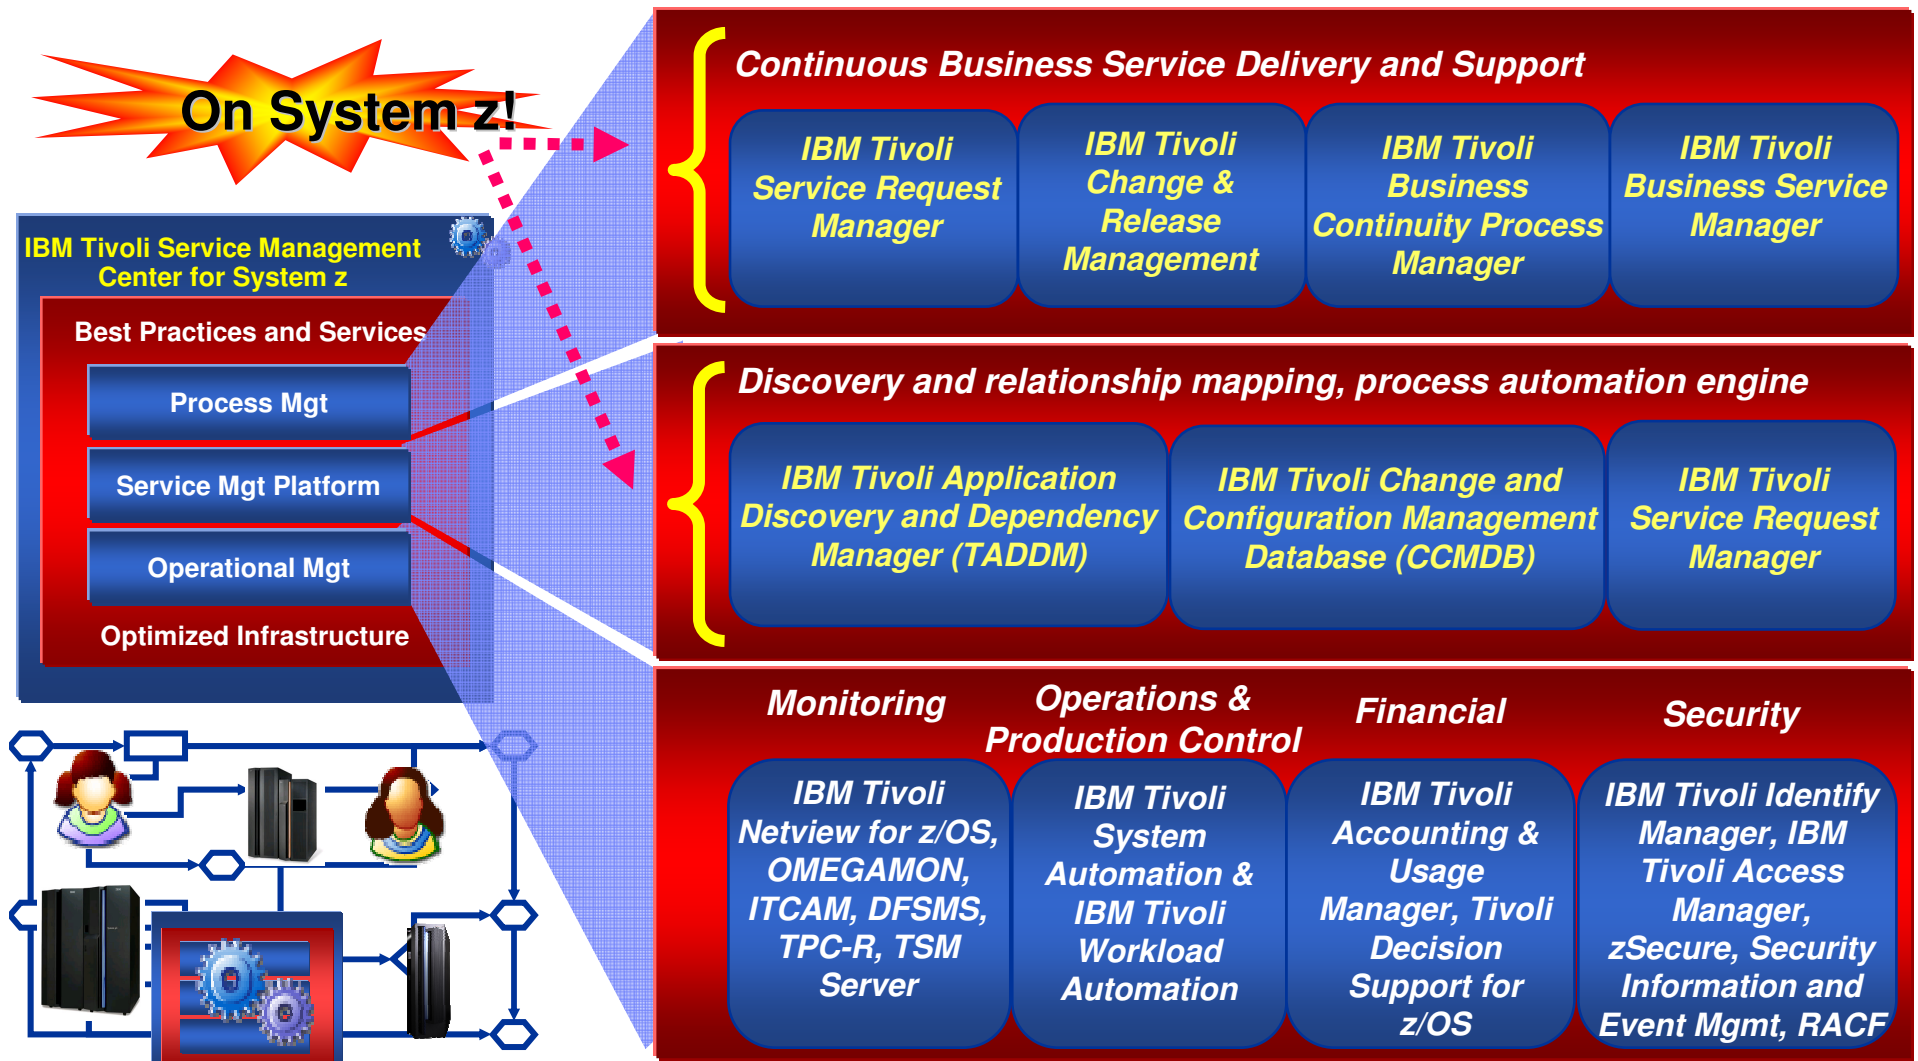
**Asset Management**

**Network Mgmt**

**Storage Mgmt**

# IBM Tivoli Service Management Center for System z


*Enabling clients to strategically use their System z as an integrated, enterprise-wide, hub for the efficient management of business and IT services*



# Agenda

- **Virtualization**
  - What is it ?
  - Why System z ?
- **System Management**
  - An incremental approach
  - Basic health/TCO - performance
- **Better Health**
  - What are you waiting on ?
  - And ...
- **Summary**

# IBM System z Virtualization Architecture



- Multi-dimensional virtualization technology**

- System z provides logical (LPAR) and software (z/VM) partitioning
  - PR/SM enables highly scalable virtual server hosting for LPAR *and* z/VM virtual machine environments
  - IRD coordinates allocation of CPU and I/O resources among z/OS and non-z/OS LPARs\*

\* Excluding non-shared resources like Integrated Facility for Linux processors

## Incremental management value – Crawl, walk, run

### ■ **Manage the environment**

- Base monitoring feature
- Storage & Data Protection feature
- Security & compliance feature

### ■ **Manage the workload**

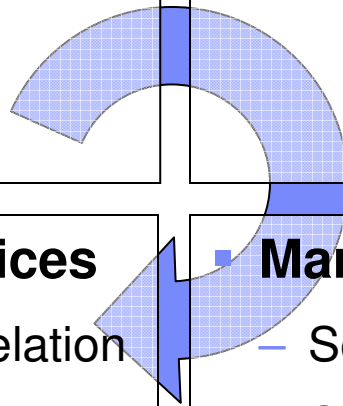
- Automation feature
- Application management feature
- Capacity planning feature

### ■ **Manage the business services**

- Event consolidation and correlation feature
- Financial management feature
- Business Service Management feature
- High availability feature

### ■ **Manage service delivery**

- Service operations feature
- Service transition feature



# Project Green Operational Goal

## Minimize change while leveraging System z

**The distributed and mainframe support teams collaborated with IBM's Design Center to develop the reference architecture and the basis for the operational approach**

### **Approach:**

- Adapt existing UNIX® team processes
- Engage z team to operate System z and z/OS®
- Broaden the VM/mainframe knowledge of the mid-range team through training, to assist in support of VM Hypervisor (z/VM®) and Linux
- Use existing monitoring and operational tools, i.e.
  - VM Resource Manager
  - Tivoli Monitoring and Enterprise Portal
  - Monitor and Performance Toolkit
  - Administer Capacity, OS Provisioning and Software Distribution Tools

## Planned incremental transformation



# What are you using ?

## z/VM

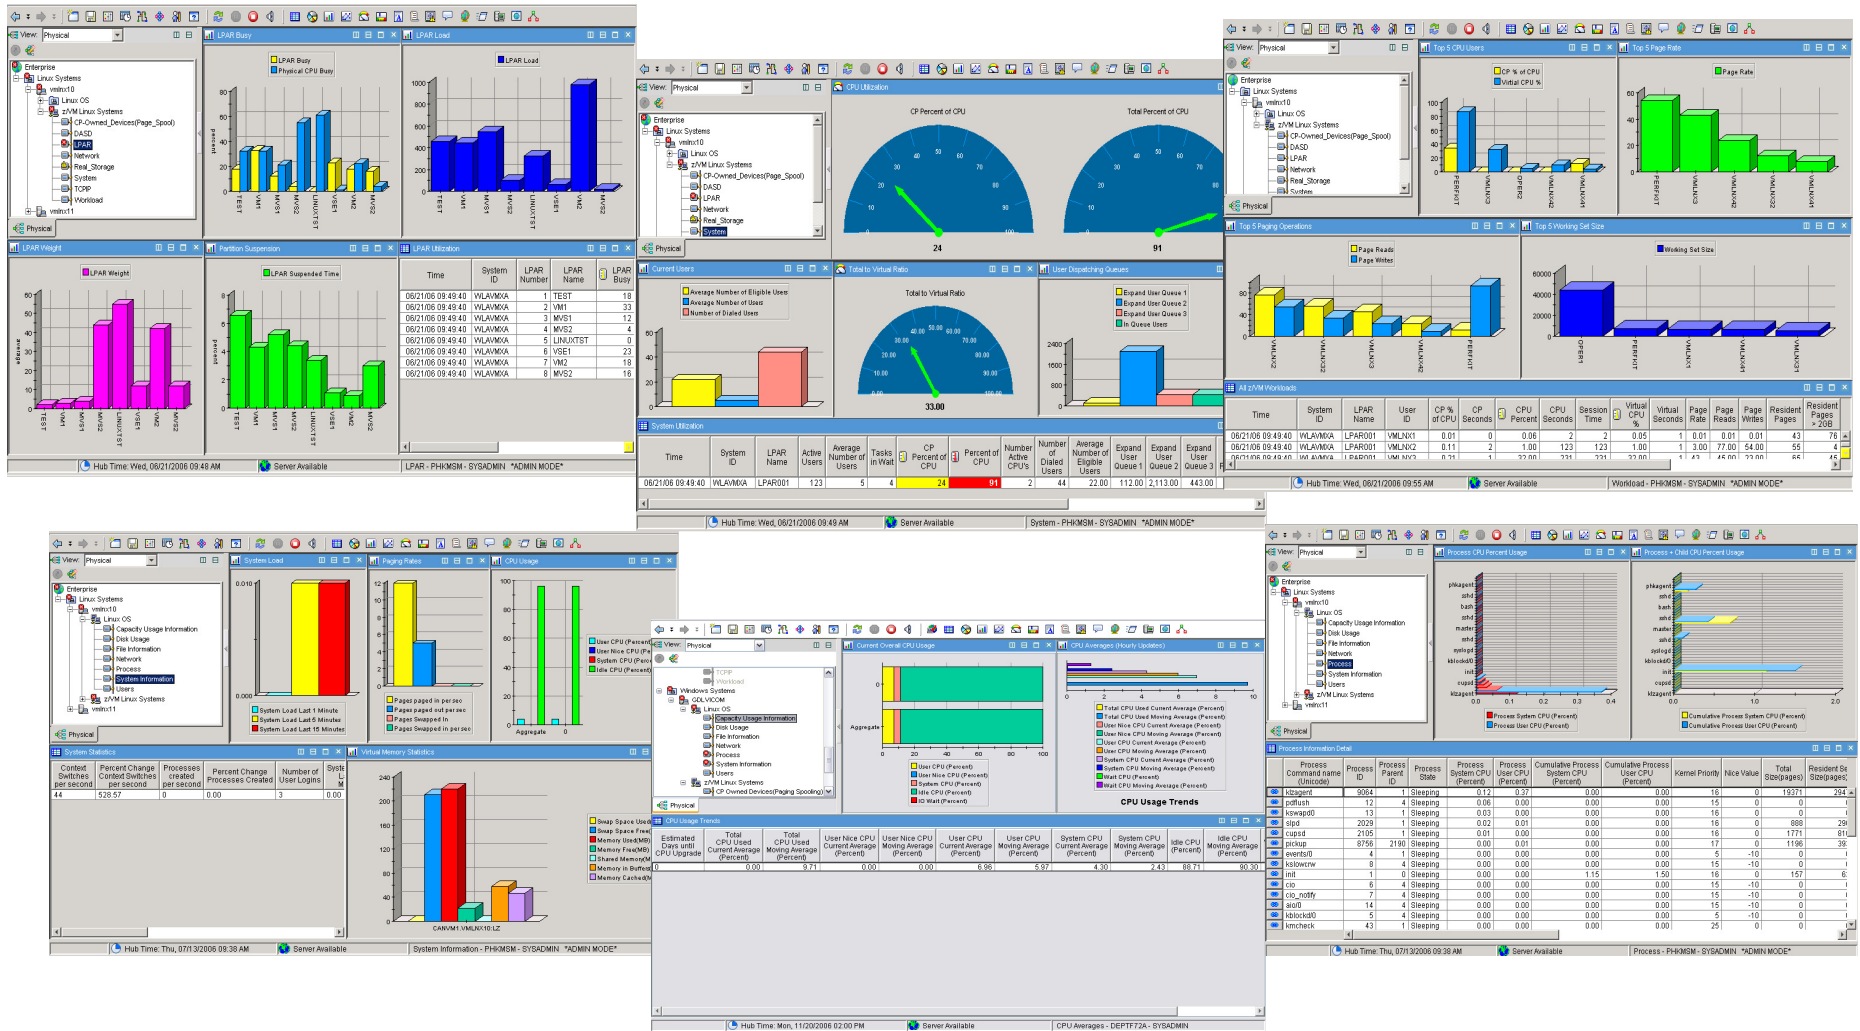
- Processors
- SYSTEM Utilization, spinlocks
- Workload
  - Linux Appldata
  - Scaled & total CPU values (almost new)
  - **Resource Constraint Analysis (new)**
- LPAR Utilization
- PAGING and SPOOLING Utilization
- DASD
- Minidisk Cache
- Virtual Disks
- Channels
- CCW Translation
- REAL STORAGE Utilization
- NETWORK Utilization (Hiper Socket and Virtual Switch)
- TCPIP Utilization – Server
- TCPIP Utilization – Use

## Linux

- Linux OS
- System Information
  - CPU aggregation
  - Virtual Memory Statistics
- Process
- Users
- Disk Usage
- File Information
- Network

# OMEGAMON XE on z/VM and Linux – CPU data example

## z/VM - CEC/LPAR/Guest Linux – System/CPU/Process



## Overall Health of Your System – Expert Advice

### ■ Things to look for

- CPU usage
  - Is any one system using more CPU than expected
  - Is any one system using less CPU than expected—you may have an underutilized processor and be wasting capacity
  - Remember, a DEDICATED processor will show 100%
- Users waiting for resources
  - Number of users at the end of the monitoring interval who are either in:
    - Eligible list—**waiting** to enter the dispatch list
    - Nondispatchable
      - > **Waiting** for paging
      - > **Waiting** for I/O completion
    - Dispatchable
      - > **Waiting** for a processor

## Overall Health of Your System – Expert Advice

### ■ Things to look for

#### – System paging rate

- Number of page reads per second
- Not a complete indicator of your paging effectiveness, but a good first glance
  - If the rate is low, and you don't have many users **waiting** for paging to complete (dispatch list), then you don't have a problem
  - If rate is low and you DO have many users in dispatch list, it may be an indication of a paging problem.
    - > High dispatch list number could be for other reasons such as **I/O contention**. You need to check.
- If the rate is high, then you may need to tune your paging subsystem.
- High and low are relative. You need to keep historical data so you can tell when your rate has changed dramatically.

# Agenda

- **Virtualization**
  - What is it ?
  - Why System z ?
- **System Management**
  - An incremental approach
  - Basic health – what are you using ?
- **Better Health**
  - What are you waiting on ?
  - And ....
- **Summary**

## What are you waiting on ?

- **Active Percent**
- **Running Percent**
- **Waiting Percent**
  
- **Eligible List Percent.**
- **Q0 List Percent.**
- **Q1 List Percent**
- **Q2 List Percent**
- **Q3 List Percent**
  
- **Test Idle Wait Percent**
- **Limit List Wait**
  
- **Page Active Wait Percent**
- **Page Wait Percent**
  
- **I/O Wait Percent**
- **Wait While I/O Active Percent**
  
- **Instruction Simulation Wait Percent**
  
- **Loading Percent**
  
- **SVM and Dormant Wait List Percent**
- **SVM and Eligible List Wait Percent**
- **SVM and Test Idle Wait Percent**
- **Console Function Wait Percent**
  
- **Other State Percent**

## Resource Constraints

### ■ CPU Wait

- Constraint in processor resources
  - May need to re-allocate workload
  - May need to add processors to user
  - If most processors are running above 90%, may need to add real processors (another IFL)

### ■ I/O Wait

- Waiting for I/O completion. Virtual machine can not run
- Check your minidisk cache, vdisks and I/O configuration

### ■ Page Wait

- Waiting to read a page into storage.
  - Check paging system
  - If a key server virtual machine, consider reserving pages

## Agenda

- **Virtualization**
  - What is it ?
  - Why System z ?
- **System Management**
  - An incremental approach
  - Basic health – what are you using ?
- **Better Health**
  - What are you waiting on ?
  - And ....
- **Summary**



## ACTIVE Per Cent

- Percentage of samples where the user was active, either consuming CPU or waiting for some service.
- If %Active value is low (less than 25%), then all wait states for that workload are only reflecting a small percentage of the data interval. May not be meaningful. Should always look at multiple intervals, but if the %ACTIVE is low then multiple intervals is a must. Gives you a perspective on the rest of the data

Time	System ID	LPAR Name	User ID	Active Percent	Running Percent	CPU Wait Percent	Loading Percent	Page Wait Percent	I/O Wait Percent	Instruction Simulation Wait Percent	Test Idle Wait Percent	Console Function Wait Percent	SVM and Test Idle Wait Percent
09/02/09 15:21:55	WLAVM0A	CANVM1	VMLNX11	100	27	3	0	0	0	0	70	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VMLNX5	100	0	20	0	0	0	0	80	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VMSERVS	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VSMPROXY	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VSMREQIN	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VMLNX03	83	4	12	0	0	0	0	84	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VMLNX01	57	0	0	0	0	0	0	100	0	0

## CPU wait

- **Percentage of active samples where the user was waiting to run on the real CPU.**
- **This means workload is in the dispatch list and waiting for a processor.. Almost always a lack of CPU resources. The System workspace shows the CPU utilization.**
  - May well need to add more processors to LPAR
  - Check to see if you have a multi-threaded application. May need to add virtual processors to a particular workload
  - It would be a very odd circumstance to see a high %CPU Wait and not be CPU constrained.

Time	System ID	LPAR Name	User ID	Active Percent	Running Percent	CPU Wait Percent	Loading Percent	Page Wait Percent	I/O Wait Percent	Instruction Simulation Wait Percent	Test Idle Wait Percent	Console Function Wait Percent	SVM and Test Idle Wait Percent
09/02/09 15:21:55	WLAVM0A	CANVM1	VMLNX11	100	27	3	0	0	0	0	70	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VMLNX5	100	0	20	0	0	0	0	80	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VMSERVS	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VSMPROXY	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VSMREQIN	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VMLNX03	83	4	12	0	0	0	0	84	0	0
09/02/09 15:21:55	WLAVM0A	CANVM1	VMLNX01	57	0	0	0	0	0	0	100	0	0

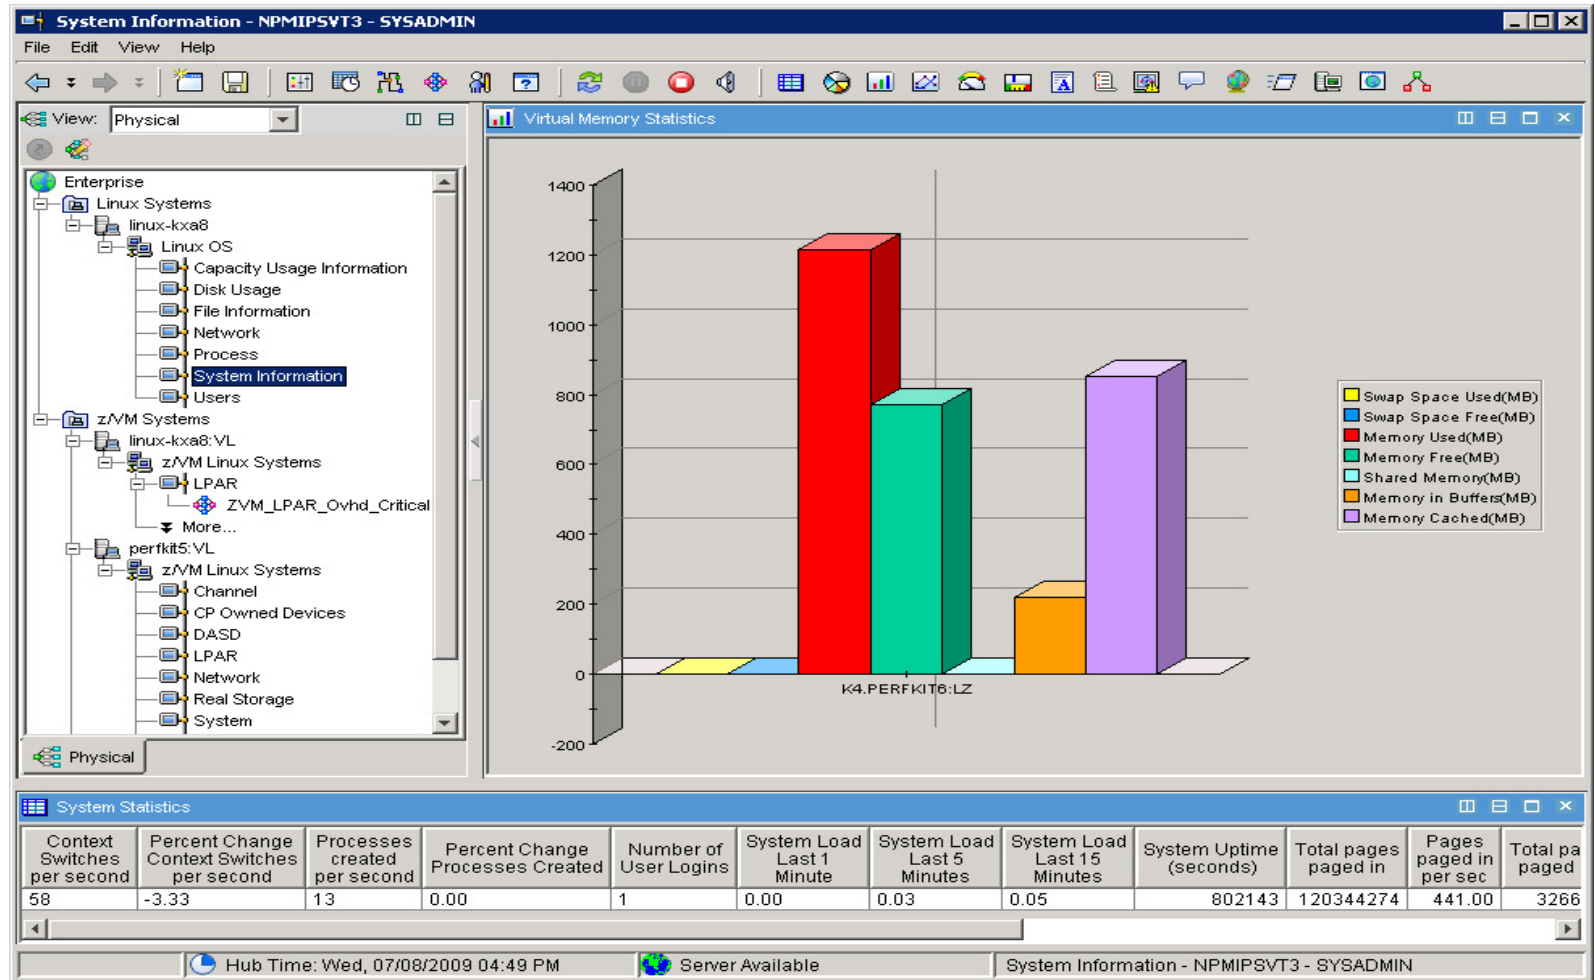
## Limit List

- Percentage of active measurements where the user was in the limit list, that is in the dispatch list waiting for CPU, but prevented from running by a maximum share setting.
- This is where a workload goes if its used up its SHARE settings. This tells you the effectiveness of how you've set your SHARES. It may be good or bad, depending on what you want to accomplish. Bad if work you want to run is there, good if you are trying to keep someone from using too much resource.

SVM and Test Idle Wait Percent	SVM and Eligible List Wait Percent	SVM and Dormant List Wait Percent	Wait While I/O Active Percent	Page Active Wait Percent	Limit List Wait Percent	Other State Percent	Q0 List Percent	Q1 List Percent	Q2 List Percent	Q3 List Percent	Eligible List Percent	Sum of Top 5 Wait States
0	0	0	0	0	0	0	0	0	0	100	0	0
0	0	0	0	0	0	0	0	0	0	100	0	0
0	0	100	0	0	0	0	0	0	0	0	0	0
0	0	100	0	0	0	0	0	0	0	0	0	0
0	0	100	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	28	72	0	0	0
0	0	0	0	0	0	0	0	59	35	6	0	0

# Sizing Linux Guests

**Memory usage of a particular Linux virtual machine**



## Sizing Linux Guests

- **More memory is not always better**
  - Linux will use all available memory
    - Any space it doesn't need will be used as file buffer cache. Notice the large amount of cache used in example—indicates that guest may be sized too large
    - Larger Linux guests means that z/VM has to swap out larger virtual machines when running other guests
    - Check on Page Wait
  - One method—use monitor to watch for swapping. Shrink guest size until it starts swapping.
  - Another method. Look at the Working Set Size for the Virtual Machine. This shows what z/VM is using for the guest. May still be too large if storage is used for cache
  - To handle some swapping, define a Vdisk. This is much faster than swapping to a real device
- **Don't define more virtual processors than you have logical processors**

# Check that Paging waiting

All z/VM Workloads

Time	System ID	LPAR Name	User ID	Active Percent	Running Percent	CPU Wait Percent	Loading Percent	Page Wait Percent	I/O Wait Percent	Instruction Simulation Wait Percent	Test Idle Wait Percent	Console Function Wait Percent	SVM and Test Idle Wait Percent
09/02/09 15:21:55	WLAVMXA	CANVM1	VMLNX11	100	27	3	0	0	0	0	70	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VMLNX5	100	0	20	0	0	0	0	80	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VMSERVS	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VSMPROXY	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VSMREQIN	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VMLNX03	83	4	12	0	0	0	0	84	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VMLNX01	57	0	0	0	0	0	0	100	0	0

SVM and Test Idle Wait Percent	SVM and Eligible List Wait Percent	SVM and Dormant List Wait Percent	Wait While I/O Active Percent	Page Active Wait Percent	Limit List Wait Percent	Other State Percent	Q0 List Percent	Q1 List Percent	Q2 List Percent	Q3 List Percent	Eligible List Percent	Sum of Top 5 Wait States
0	0	0	0	0	0	0	0	0	0	100	0	0
0	0	0	0	0	0	0	0	0	0	100	0	0
0	0	100	0	0	0	0	0	0	0	0	0	0
0	0	100	0	0	0	0	0	0	0	0	0	0
0	0	100	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	28	72	0	0	0
0	0	0	0	0	0	0	0	59	35	6	0	0

## Agenda

- **Virtualization**
  - What is it ?
  - Why System z ?
- **System Management**
  - An incremental approach
  - Basic health – what are you using ?
- **Better Health**
  - What are you waiting on ?
- **Summary**

## Finally

- **Virtualization on System z must be managed**
- **Key to success is performance management**
- **What you are using is important but ...**

**What you are waiting on may be more important !!!**



# Questions

## SWG Product Matrix

- For Internal Matrix:

[IBM Linux Portal - IBM Middleware Available on Linux](#)

- For External Matrix:

[Linux and IBM - IBM Middleware Available on Linux](#)

- SWG Federated Integration Test (FIT) Stacks descriptions:

[SWG Federated Integration Test \(FIT\) Stacks](#)

## OMEGAMON XE on z/VM and Linux

- **Provides performance monitoring for z/VM and Linux guests**
- **Linux agents gather detailed performance data from Linux guests**
- **z/VM agent gathers performance data from z/VM and Linux**
  - Including z/VM view of guests
  - Uses IBM Performance Toolkit for VM as its data source
  - Linux provides APPLDATA to CP monitor.
- **Executes automated actions in response to defined events or situations**
- **Part of the Tivoli Management Services infrastructure and OMEGAMON family of products**
  - Specifically focused on z/VM and Linux guests

# Overall Health of Your System

At a quick glance you can see the %CPU usage, number of users in a wait state, and paging rates of all your z/VM systems

The screenshot displays the z/VM Systems - NPMIPSVT3 - SYSADMIN interface. It features a tree view on the left showing the system hierarchy. The main area contains several charts and a table:

- z/VM Systems: Top 5 CPU Busy Percent:** A 3D bar chart showing CPU usage for systems WLAVMKA, GDLVICOM, and GDLVMK4.
- z/VM Systems: Top 5 Waiting Users Count:** A 3D bar chart showing the number of users waiting for resources for the same three systems.
- z/VM Systems: Top 5 Paging Rate:** A 3D bar chart showing the system paging rate per second for the same three systems.
- z/VM Systems: Active z/VM Systems Registered:** A table listing active systems with their respective metrics.

Time	System ID	LPAR Name	Active Users	Average Number of Users	Tasks in Wait	CP Percent of CPU	Percent of CPU	Number Active CPU's	System Paging Rate
07/07/09 15:06:10	GDLV MK4	K4	54	118	0	0	0	18	0
07/07/09 15:06:23	WLAVMKA	CANVM1	22	41	0	1	6	2	1
07/07/09 15:06:16	GDLV ICOM	VIC	17	26	0	0	1	9	0



# Tivoli OMEGAMON XE on z/VM and Linux a Scenario

## Problem

- Uneven Linux guest CPU consumption

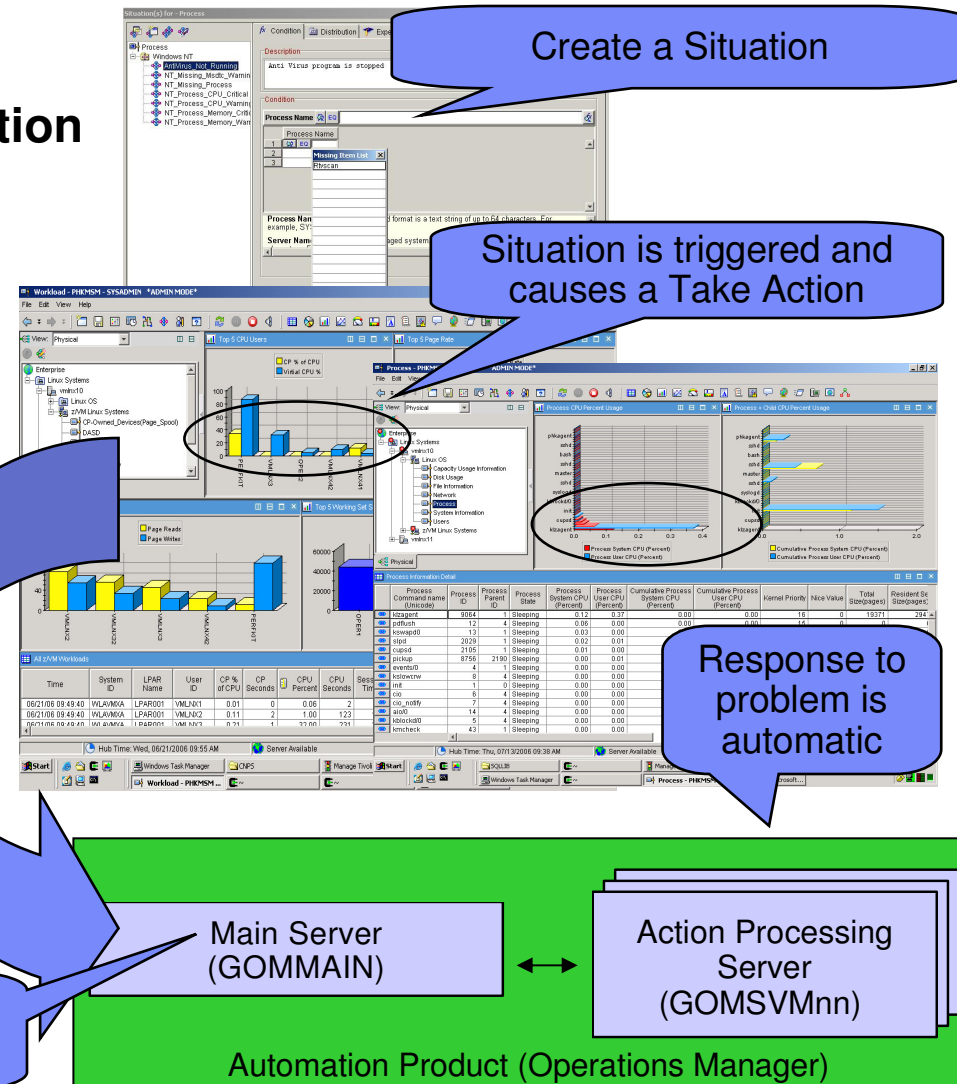
## Solution

- Use situation to recognize high swapping with high CPU and working set size
- Send message to Operations Manager
- Operations Manager invokes a rule to execute a CP tuning command to allocate more resource to the Linux Guest

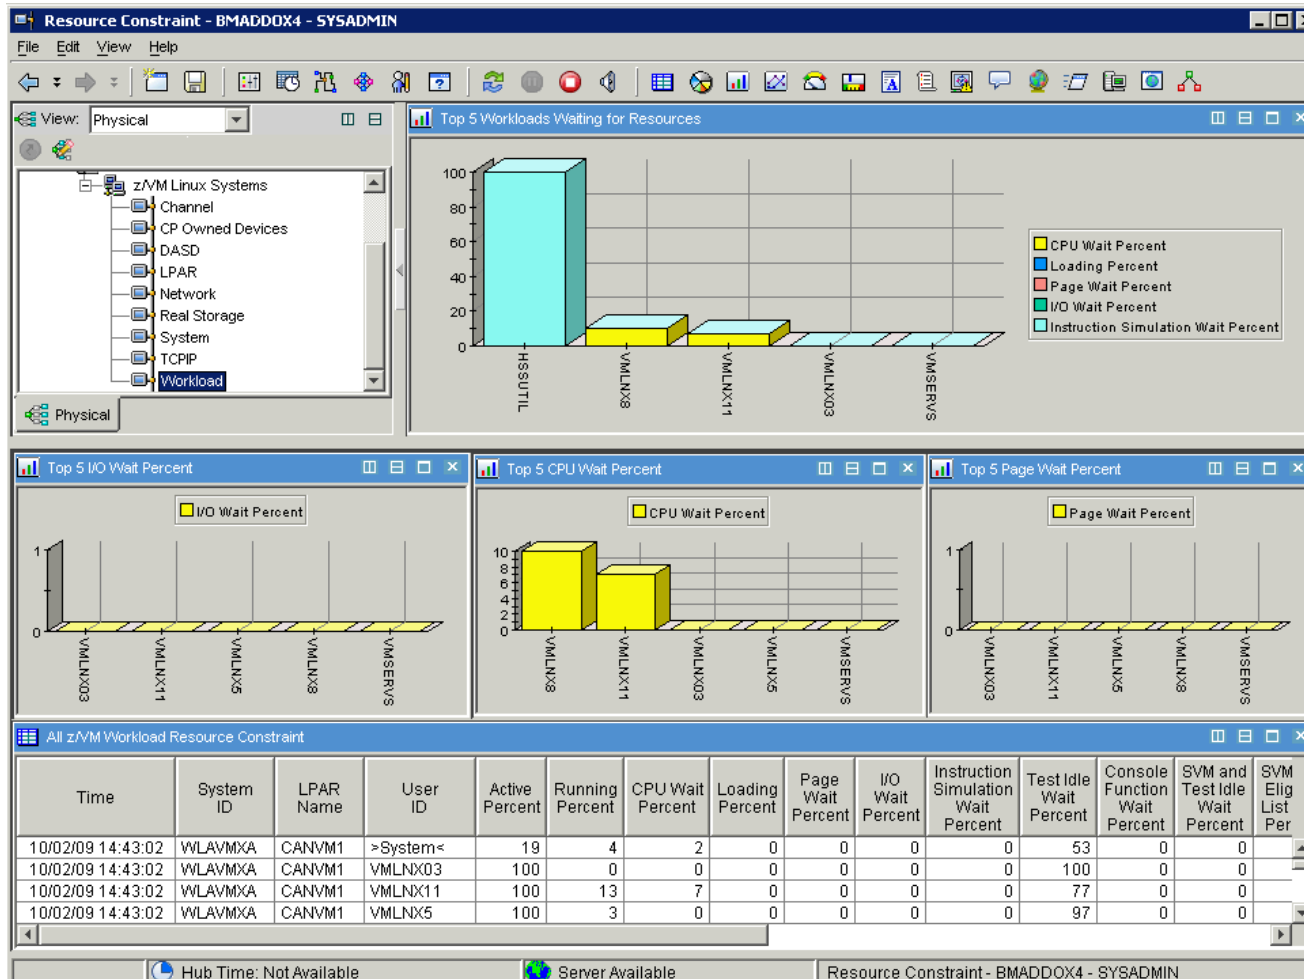
## Potential Benefits

- Automated problem resolution
- Integrated so

Message is sent and triggers z/VM automation



# What am I waiting on ?



# Let's talk about waiting

All z/VM Workloads													
Time	System ID	LPAR Name	User ID	Active Percent	Running Percent	CPU Wait Percent	Loading Percent	Page Wait Percent	I/O Wait Percent	Instruction Simulation Wait Percent	Test Idle Wait Percent	Console Function Wait Percent	SVM and Test Idle Wait Percent
09/02/09 15:21:55	WLAVMXA	CANVM1	VMLNX11	100	27	3	0	0	0	0	70	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VMLNX5	100	0	20	0	0	0	0	80	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VMSERVS	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VSMPROXY	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VSMREQIN	100	0	0	0	0	0	0	0	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VMLNX03	83	4	12	0	0	0	0	84	0	0
09/02/09 15:21:55	WLAVMXA	CANVM1	VMLNX01	57	0	0	0	0	0	0	100	0	0

SVM and Test Idle Wait Percent	SVM and Eligible List Wait Percent	SVM and Dormant List Wait Percent	Wait While I/O Active Percent	Page Active Wait Percent	Limit List Wait Percent	Other State Percent	Q0 List Percent	Q1 List Percent	Q2 List Percent	Q3 List Percent	Eligible List Percent	Sum of Top 5 Wait States
0	0	0	0	0	0	0	0	0	0	100	0	0
0	0	0	0	0	0	0	0	0	0	100	0	0
0	0	100	0	0	0	0	0	0	0	0	0	0
0	0	100	0	0	0	0	0	0	0	0	0	0
0	0	100	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	28	72	0	0	0
0	0	0	0	0	0	0	0	59	35	6	0	0