



IBM System z Technical Conference Dresden – Germany – May 5-9



Systems Management on z/VM

Chuck Brazie

z/VM Development - IBM Endicott, NY

Session V61

Disclaimer

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "AS IS" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used instead.

Any performance data contained in this document was determined in a controlled environment and, therefore, the results which may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environments.

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming or services in your country.

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

CICS*	Language Environment*	S/370
DB2*	MQSeries*	S/390*
DB2 Connect	Multiprise*	S/390 Parallel Enterprise Server
DB2 Universal Database	MVS	VisualAge*
DFSMS/MVS*	NetRexx	VisualGen*
DFSMS/VM*	OpenEdition*	VM/ESA*
e business(logo)*	OpenExtensions	VTAM*
Enterprise Storage Server*	OS/390*	VSE/ESA
ESCON*	Parallel Sysplex*	WebSphere*
FICON	PR/SM	z/Architecture
GDDM*	QMF	z/OS*
HiperSockets	RACF*	zSeries*
IBM*	RAMAC*	z/VM*
IBM(logo)*	RISC	

* Registered Trademarks of the IBM Corporation

The following are trademarks or registered trademarks of other companies.

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation.

Tivoli is a trademark of Tivoli Systems Inc.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput 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 improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

IBM considers a product "Year 2000 ready" if the product, when used in accordance with its associated documentation, is capable of correctly processing, providing and/or receiving date data within and between the 20th and 21st centuries, provided that all products (for example, hardware, software and firmware) used with the product properly exchange accurate date data with it. Any statements concerning the Year 2000 readiness of any IBM products contained in this presentation are Year 2000 Readiness Disclosures, subject to the Year 2000 Information and Readiness Disclosure Act of 1998.

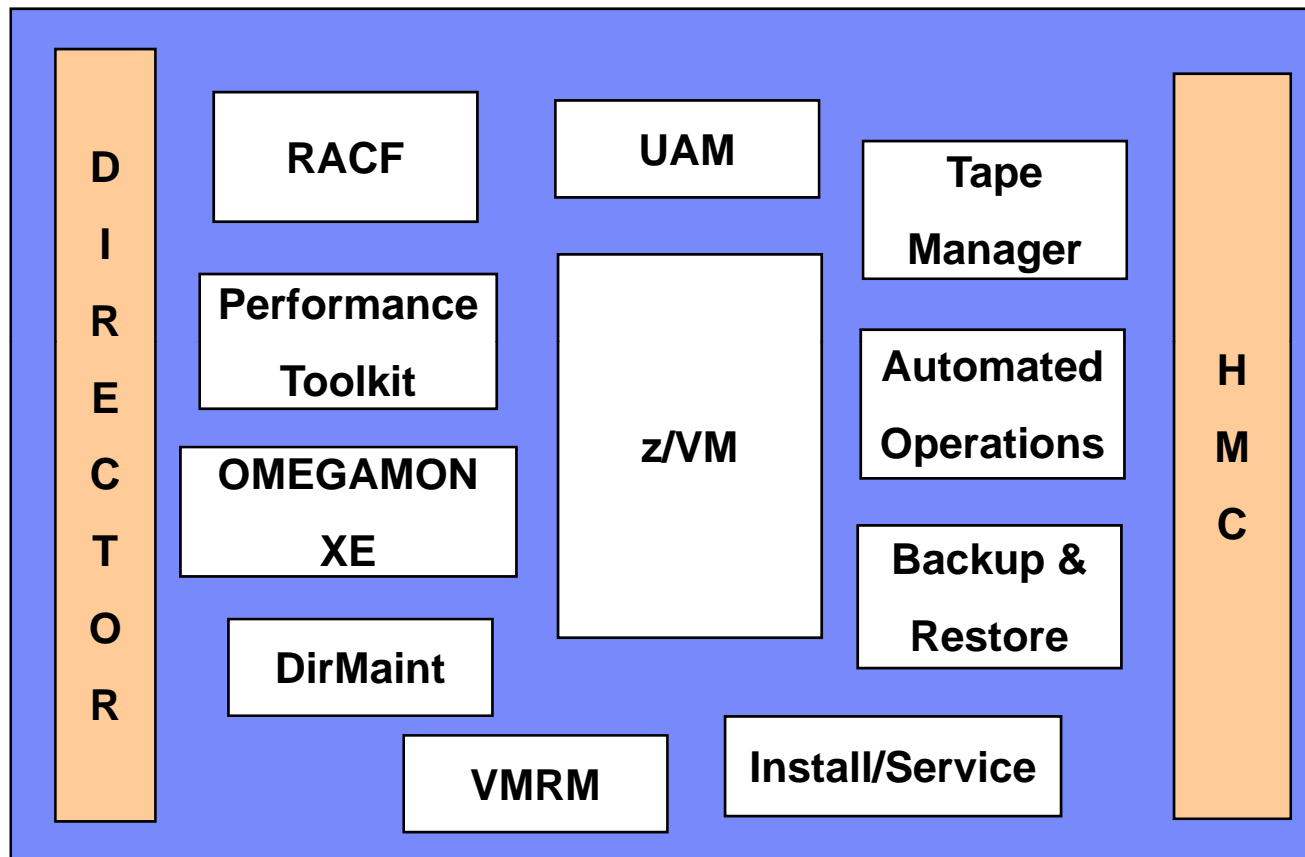
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Agenda

- **The Big Picture**
- **Virtual Systems Management API Enhancements**
- **Directory Management**
- **IBM Director for System Z**
- **Integrated Systems Management**
 - ▶ **CP System Services for VMTower, Proxy server**
- **Performance tools & monitoring**
- **Software Group and other products**
- **Summary**

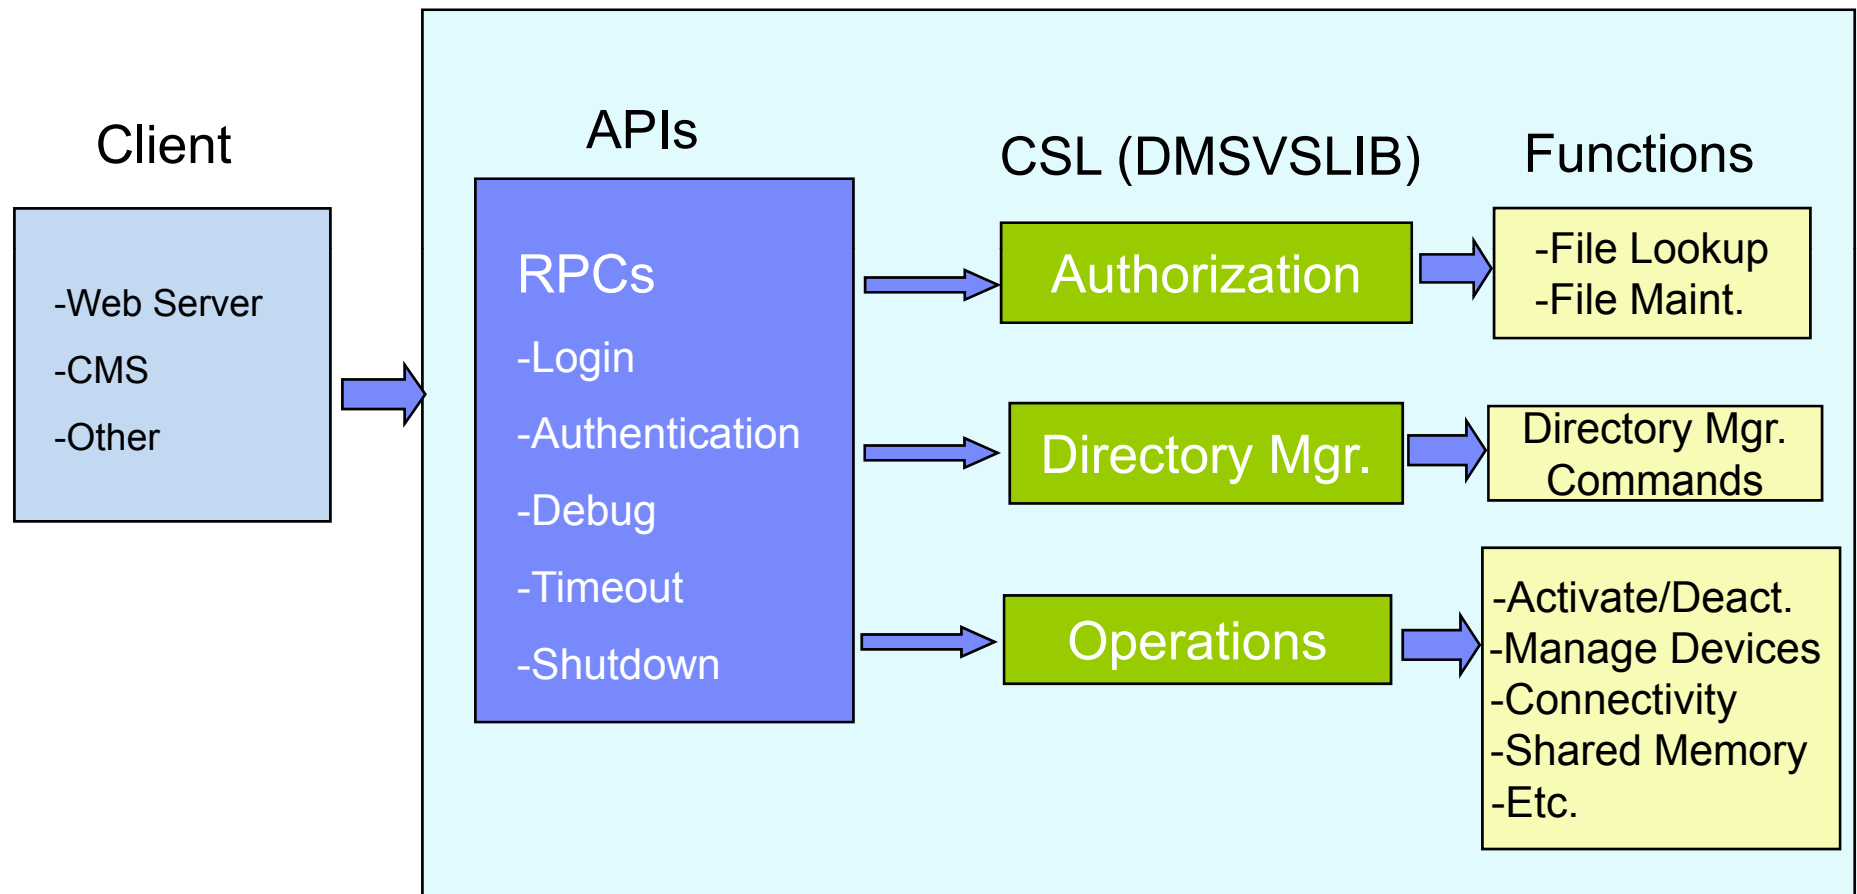
z/VM Systems Management Big Picture



- Security
- Health Metering
- Performance
- Capacity planning
- Business continuity
- Automation
- Media management
- User administration

Structure of Old Systems Management Server

VSMSEVERE

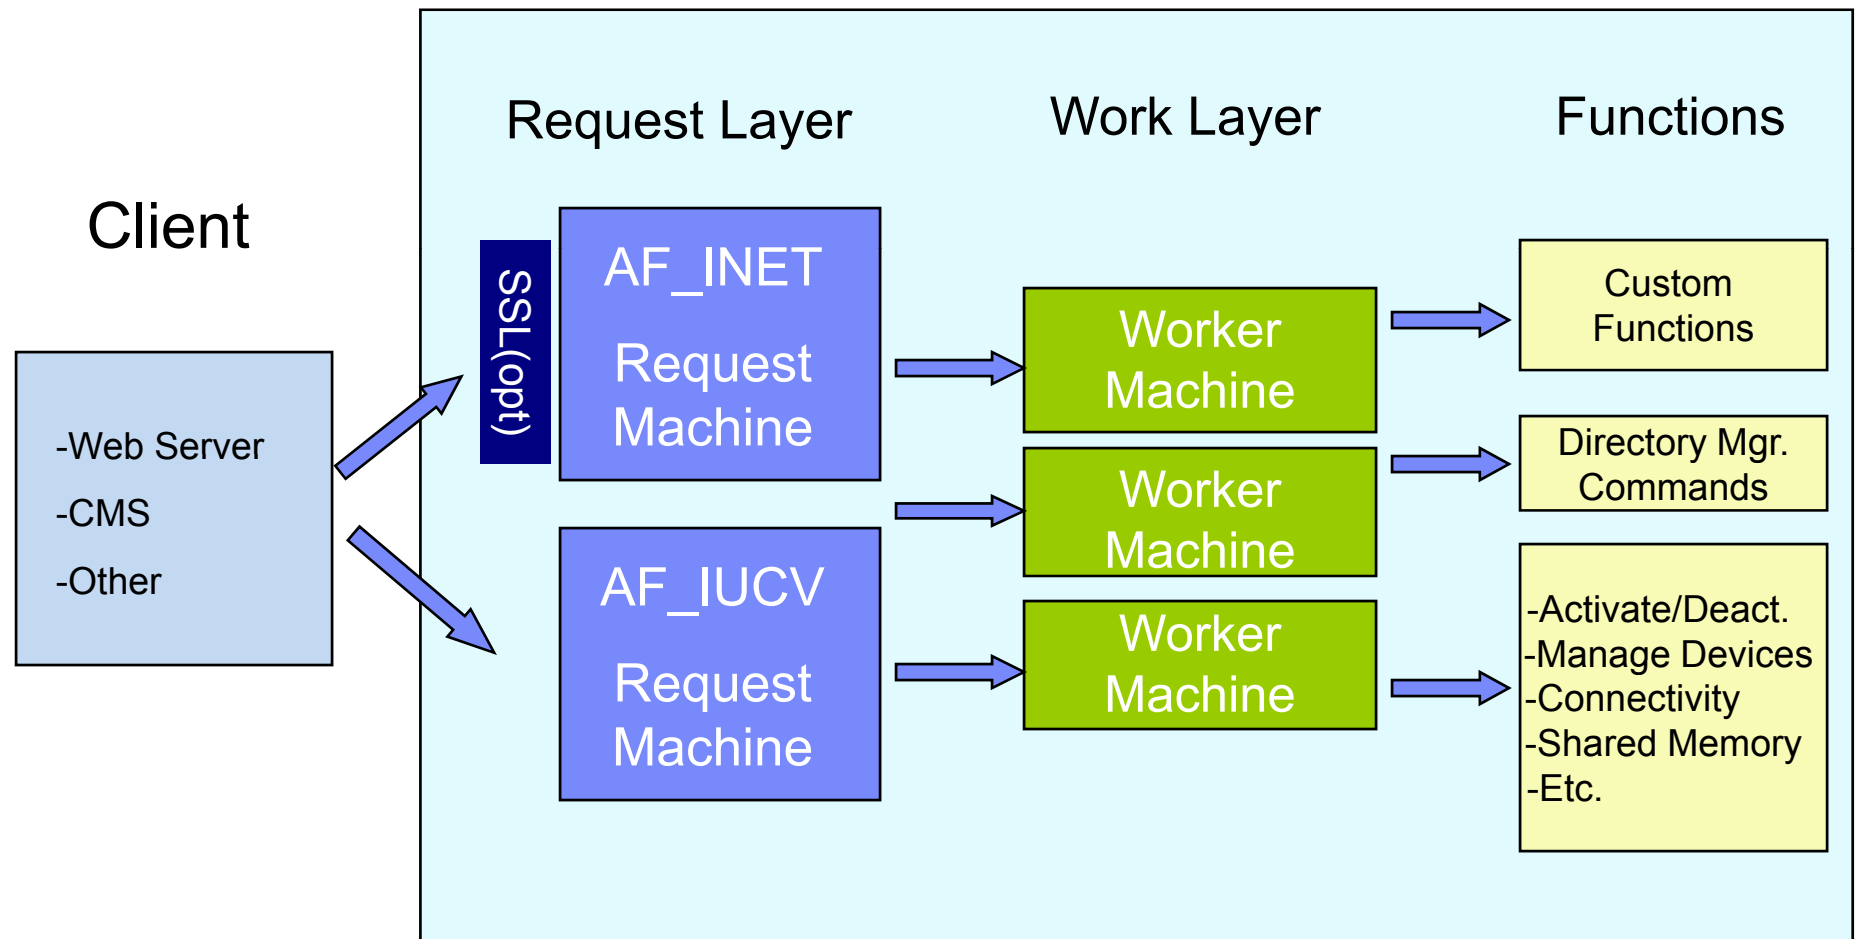


Virtual Systems Management APIs

- **New Sockets-based server in z/VM 5.3.0 to replace RPC and CSL interfaces**
 - ▶ **Multi-tasking capable**
 - ▶ **New APIs and enhancements to existing APIs in support of requirements from IBM Director**
 - ▶ **Architected output for some query APIs**
 - ▶ **Directory IPL and Directory Manager Task APIs**
 - ▶ **Directory Manager Exit Routine**
 - ▶ **Removal of Directory_Manager_Command_API**
 - ▶ **Enhancements to support passphrases**
 - ▶ **Obtain active configuration information for virtual images**
 - ▶ **Allows for user-written custom execs**
 - ▶ **Uses new CP FOR command and *ASYNCMD system service**

Structure of New Systems Management Server

VSM Server



Directory Management: IBM Directory Maintenance (DirMaint) Product

- **A CMS application that helps you manage your VM directory**
 - ▶ **Describes to CP the configuration and operating characteristics of each virtual machine**
 - ▶ **Maintains user passwords, authorizations, access to minidisks**
- **DirMaint provides a command interface for every directory statement, plus commands for managing DirMaint itself.**
- **Supports Systems Management APIs**
 - ▶ **_DM APIs can be replaced by other directory managers.**
- **Priced feature that comes installed with your VM system, but must be enabled to use it.**

Sample Directory Entry

USER VMRMSVM VMRMSVM 64M 128M AG
IPL CMS
OPTION APPLMON QUICKDSP
SHARE ABSOLUTE 3%
ACCOUNT MONITOR
MACHINE XA
IUCV *MONITOR MSGLIMIT 255
NAMESAVE MONDCSS
CONSOLE 0009 3215 T
SPOOL 000C 2540 READER *
SPOOL 000D 2540 PUNCH A
SPOOL 000E 1403 A
LINK MAINT 190 190 RR
LINK MAINT 193 193 RR
LINK MAINT 19D 19D RR
LINK MAINT 19E 19E RR
MDISK 191 3390 1783 040 510RES MR READ WRITE MULTIPLE

Directory Management...

- **Enhanced Guest Configuration**
 - ▶ **User Directory COMMAND statement**
 - ▶ **Allows (almost) any CP command to be executed for a guest**
 - **After LOGON complete, immediately before IPL**
 - **Any privilege class**
 - **Subject to ESM auditing**
 - **Multiple statements allowed**
 - ▶ **Reduces need for future directory control statement enhancement**
 - **E.g., Specialty engines defined via COMMAND DEFINE CPU rather than by extensions to CPU statement**

IBM Director for Linux on System z

IBM Director for Linux on System z is a full member of the IBM Systems Director family

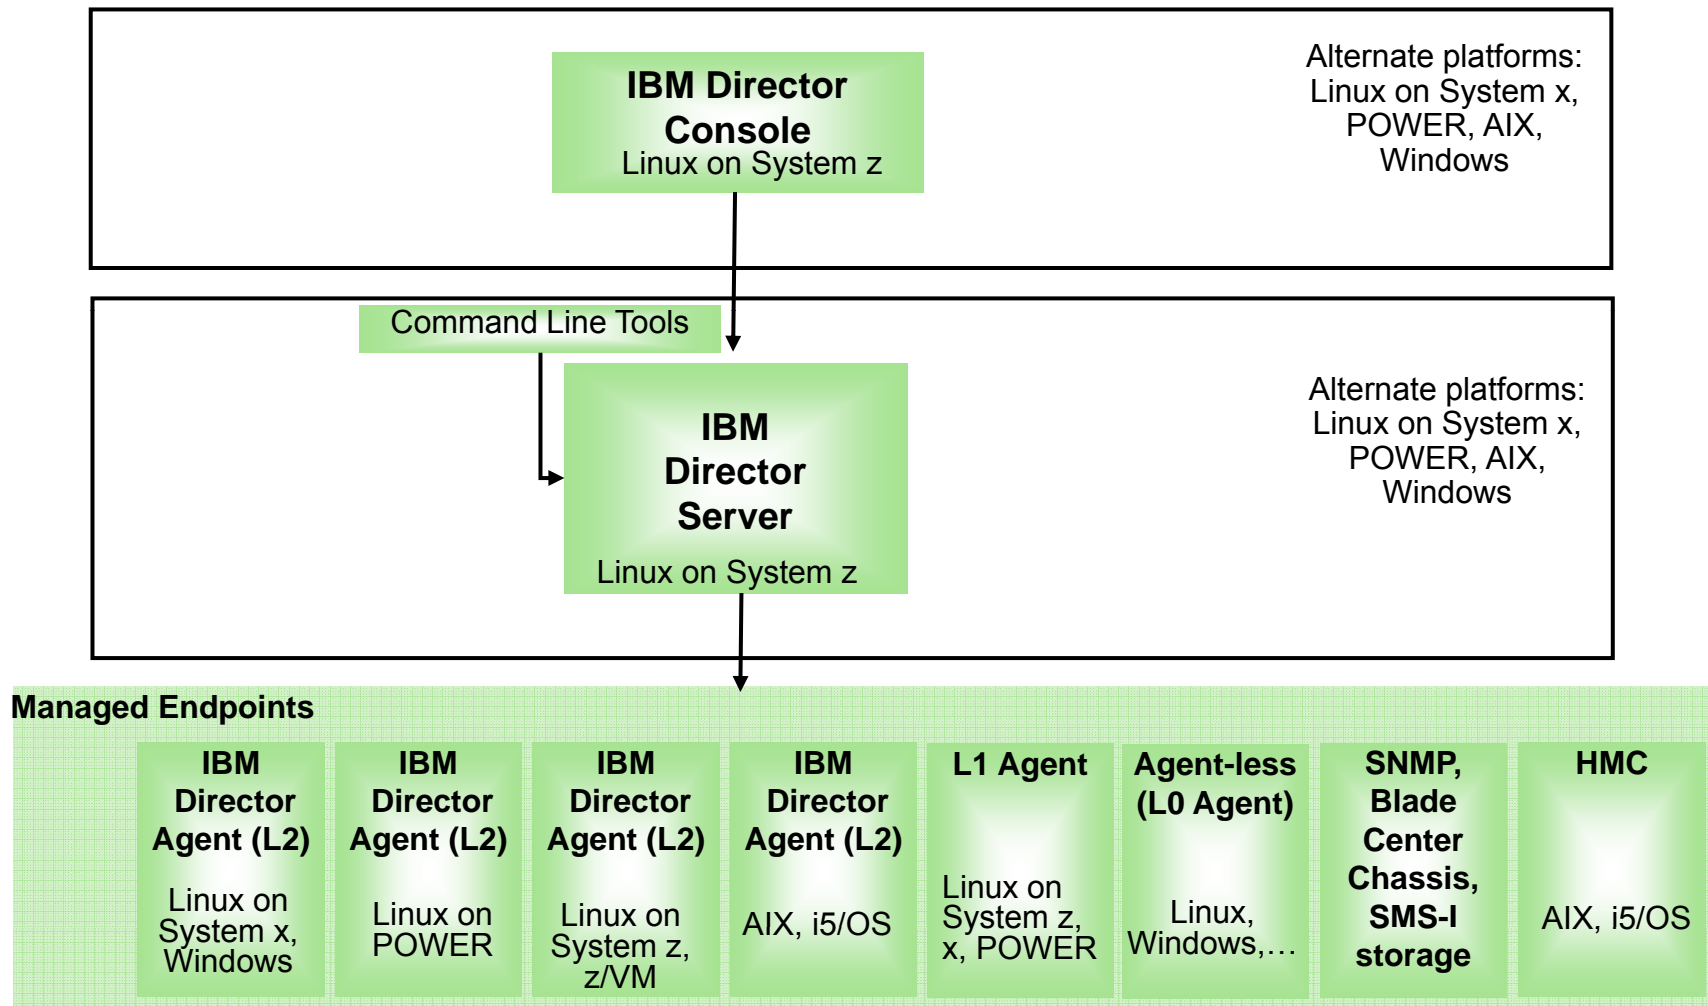
- **GUI interface using CIM (Common Information Model) Standards**
- **Includes Director Console, Server, Agents and two optional, priced extensions to run on Linux on System z:**
 - ▶ **IBM Director z/VM Center**
 - ▶ **IBM Director Software Distribution Premium Edition**
- **Free of charge core / base functions:** discovery, inventory, monitor, alert, etc.
- common to all IBM systems
- **Upgrade options to Service Management solutions from Tivoli**



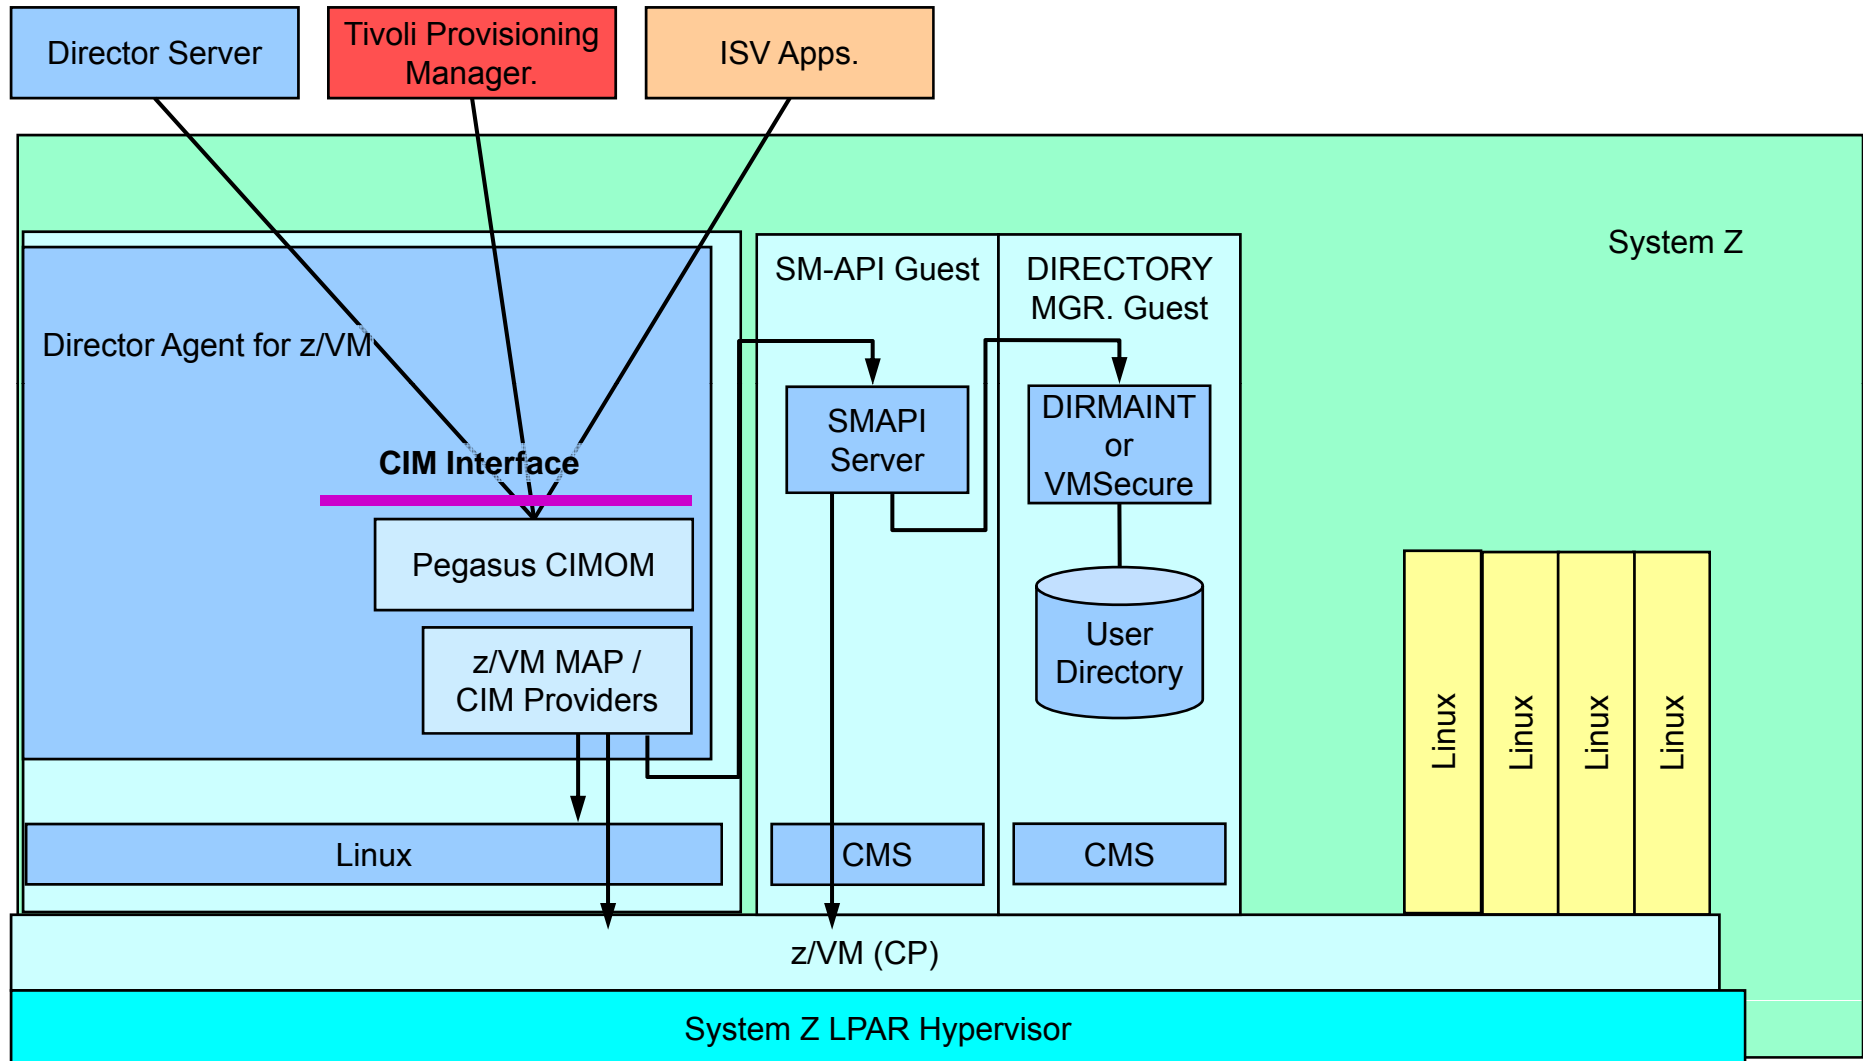
IBM Director for Linux on System z...

- **Managed Access Point (MAP) Server for z/VM**
 - ▶ **Uses Systems Management APIs to invoke CP function**
 - ▶ **z/VM Center Extensions (5.20.n)**
 - **Creation of virtual servers under z/VM**
 - **Deployment of Linux systems into those virtual servers**
 - **Allows for autonomic adjustment of users in workloads for CPU and DASD I/O management**
- **IBM Director V5.20 replaces the IBM Director V5.10 base function in the *IBM Virtualization Engine and Infrastructure Services for Linux on System z9 and zSeries* product**


IBM Director - Topology



Systems Management for z/VM with Director




IBM Director z/VM Center – System Set-up



IBM Director V5.20 z/VM Center

- **Virtual Server Deployment** – easy deployment of Linux servers on z/VM
 - ▶ Create z/VM virtual machines using the IBM Director graphical user interface
 - ▶ Template-based provisioning of Linux virtual machines
- **Server Complexes** – one-step provisioning of multiple Linux virtual servers
 - ▶ Exploits Virtual Server Deployment task
 - ▶ Manages configuration settings of Linux virtual servers in order to reduce customization steps needed by Virtual Server Deployment
 - Definition of virtual networking, minidisk attachments
 - Supports configuration scripts
 - Controls virtual server resource assignment via z/VM Virtual Machine Resource Manager (VMRM)
- **Utility Service Configuration Manager**
 - ▶ New task that provides ease-of-use enhancements for installing and customizing Linux utility services
 - ▶ A utility service combines applications, configurations, and a system environment (e.g., network setup, storage) needed for a specific end-user function
 - ▶ Learn more about Linux utilities at: ibm.com/systems/z/os/linux/utilities
- **z/VM Center is priced on a per-managed processor basis**

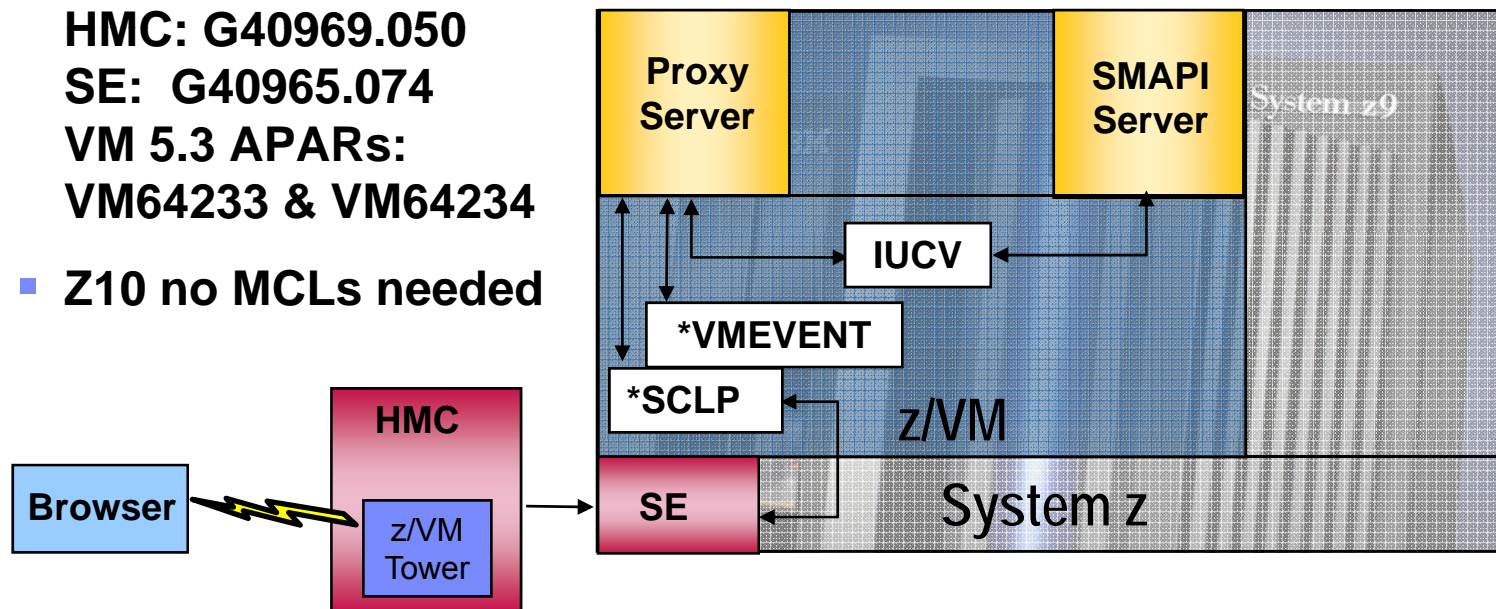


z/VM Integrated Systems Management

- **z/VM V5.3 provides an interface to allow basic z/VM systems management functions to be performed from the Hardware Management Console (HMC)**
- **Enables administration of z/VM guests without having to establish network connections, reducing complex system configuration**
- **HMC recognizes z/VM logical partition and extends management to virtual machine level**
 - Select guests to manage
 - Display guest configuration and status
 - Display Guest LAN and Virtual Switch configurations
 - Activate/Deactivate guest
 - Grouping
 - Event monitoring

z/VM Integrated Systems Management

- Web-browser access to the HMC interfaces to z/VM services including the Systems Management API to activate and deactivate guests and display guest status
- Z9 Requires MCLs:
HMC: G40969.050
SE: G40965.074
VM 5.3 APARs:
VM64233 & VM64234
- Z10 no MCLs needed



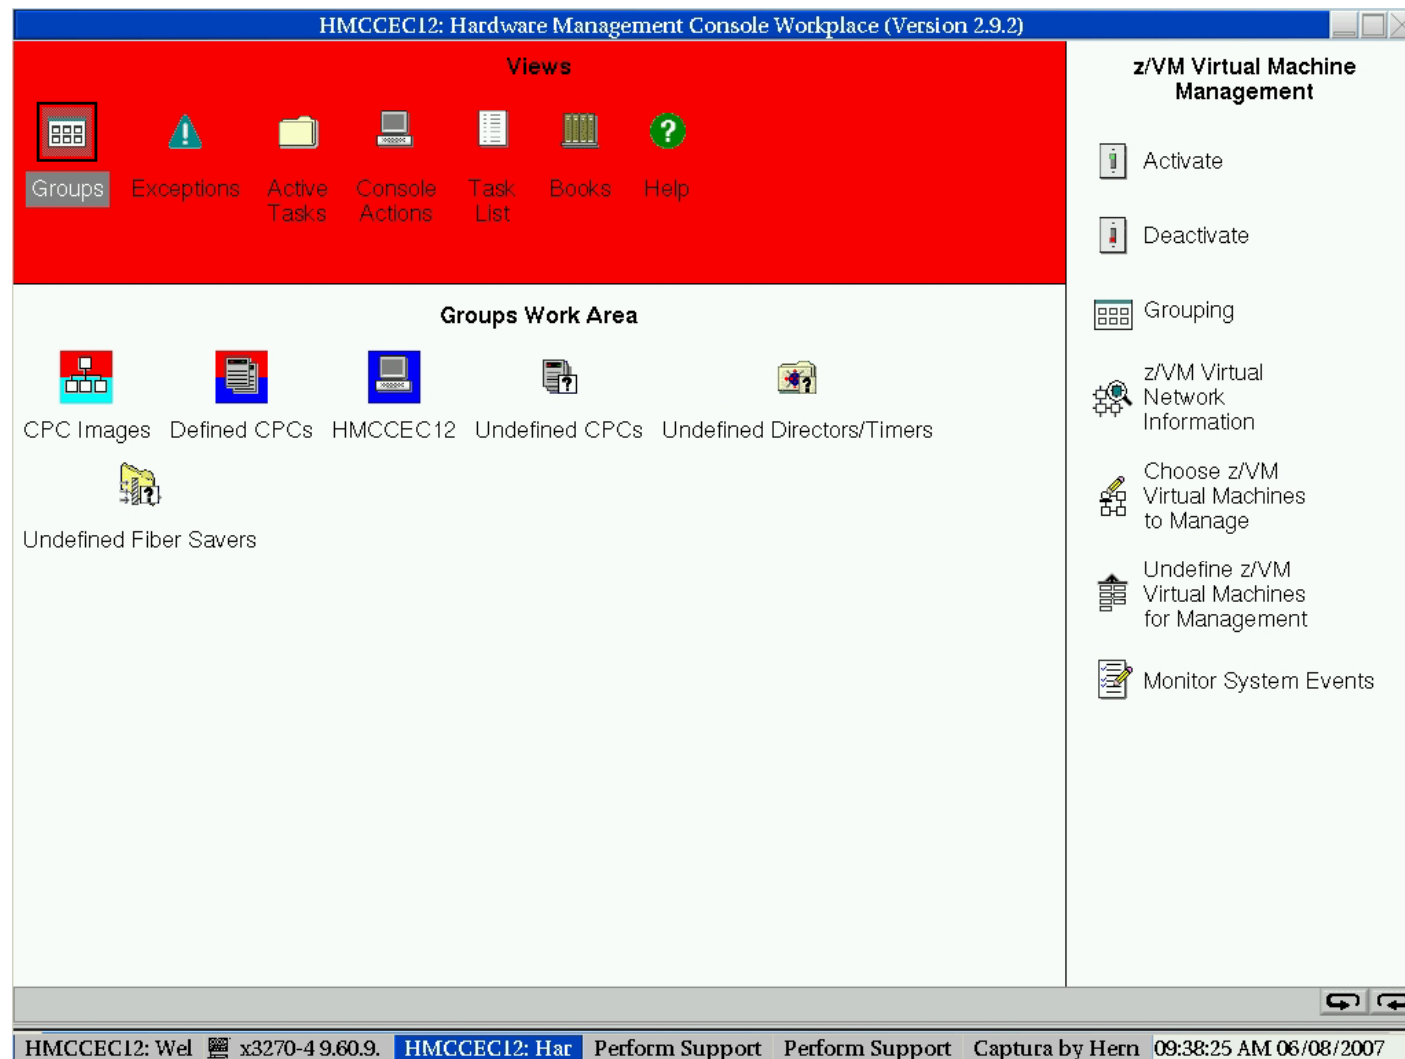
New CP System Services & Command

- ***SCLP System Service**
 - ▶ Receives and transmits HMC events
 - ▶ Exploited by a new proxy server that directs HMC requests to the SMAPI server to perform the desired function
- ***VMEVENT IUCV System Service**
 - ▶ Gives notification about certain events such as virtual machine status changes
 - User LOGON, LOGOFF, Failure conditions, LOGOFF timeout initiated, Forced Sleep initiated, Previous exception condition cleared
 - ▶ Proxy server receives the event notification and reports these back to the HMC via *SCLP
- **Query IUCV Command**
 - ▶ shows users connected, path identifiers, and associated event classes

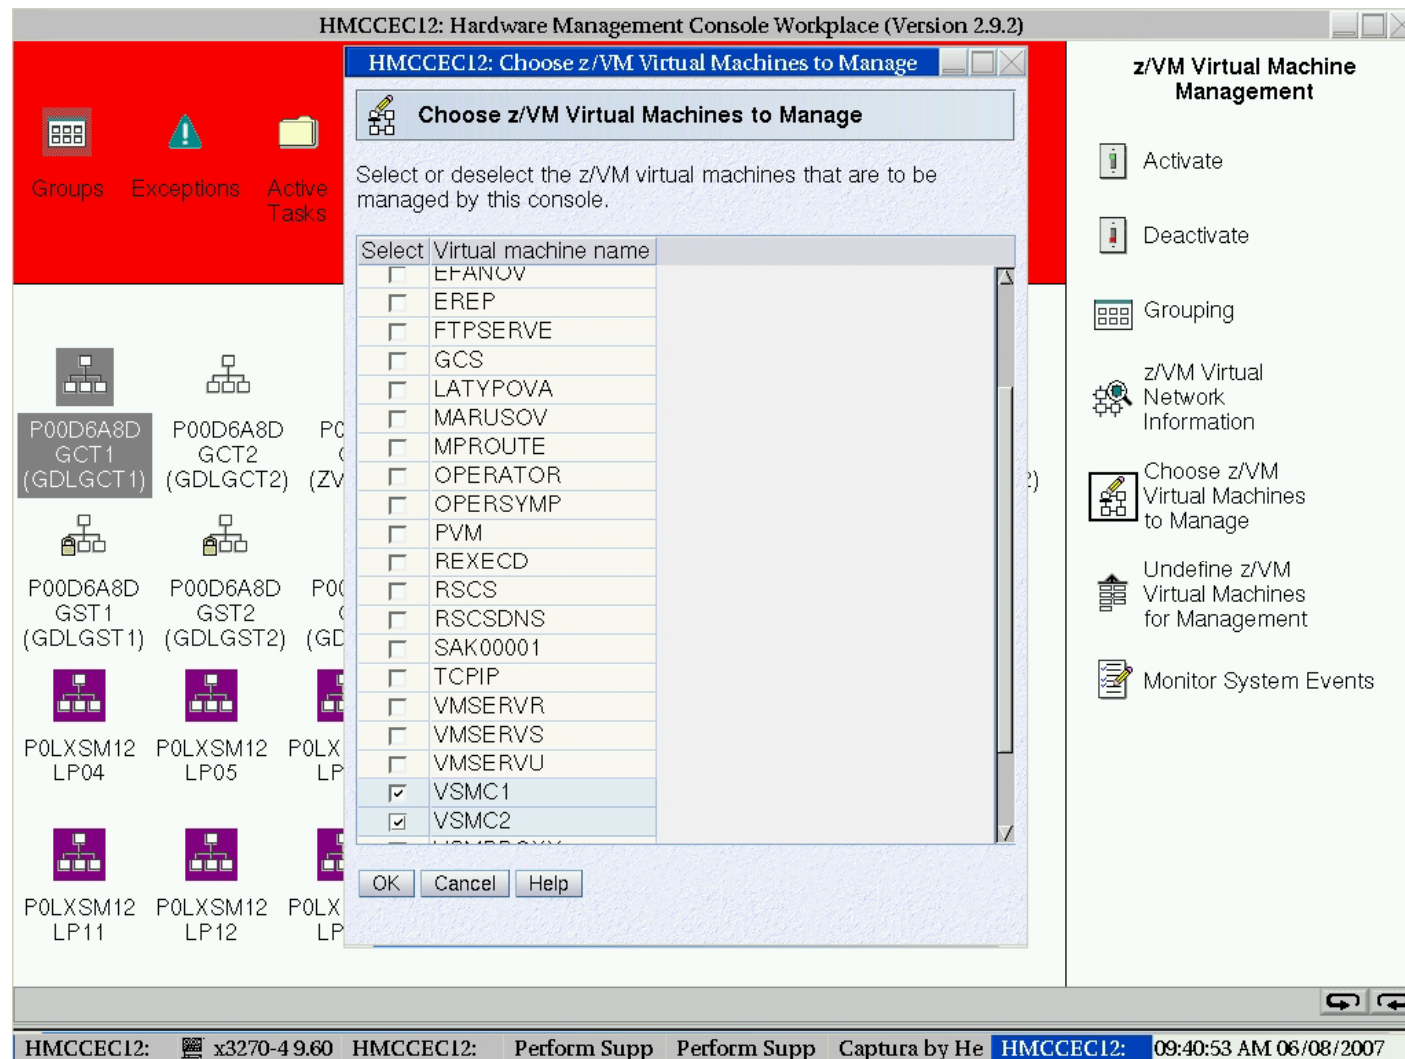
SM Proxy Server - VSMPROXY

- **Connects to *SCLP and waits for requests**
 - ▶ **Each request is converted to a *SCLP event and handled on its own thread.**
 - **Request thread collates its request and issues SMAPI request to existing AF_IUCV request server.**
 - **Request thread takes response from SMAPI server and sends it back over *SCLP**
 - ▶ **If request is to subscribe for virtual machine state changes, connect to *VMEVENT and wait for status data.**
 - **Status data sent back over *SCLP**
 - **Thread continues until an “unsubscribe” request**
 - **Only 1 thread waiting for *VMEVENT exists at a time.**

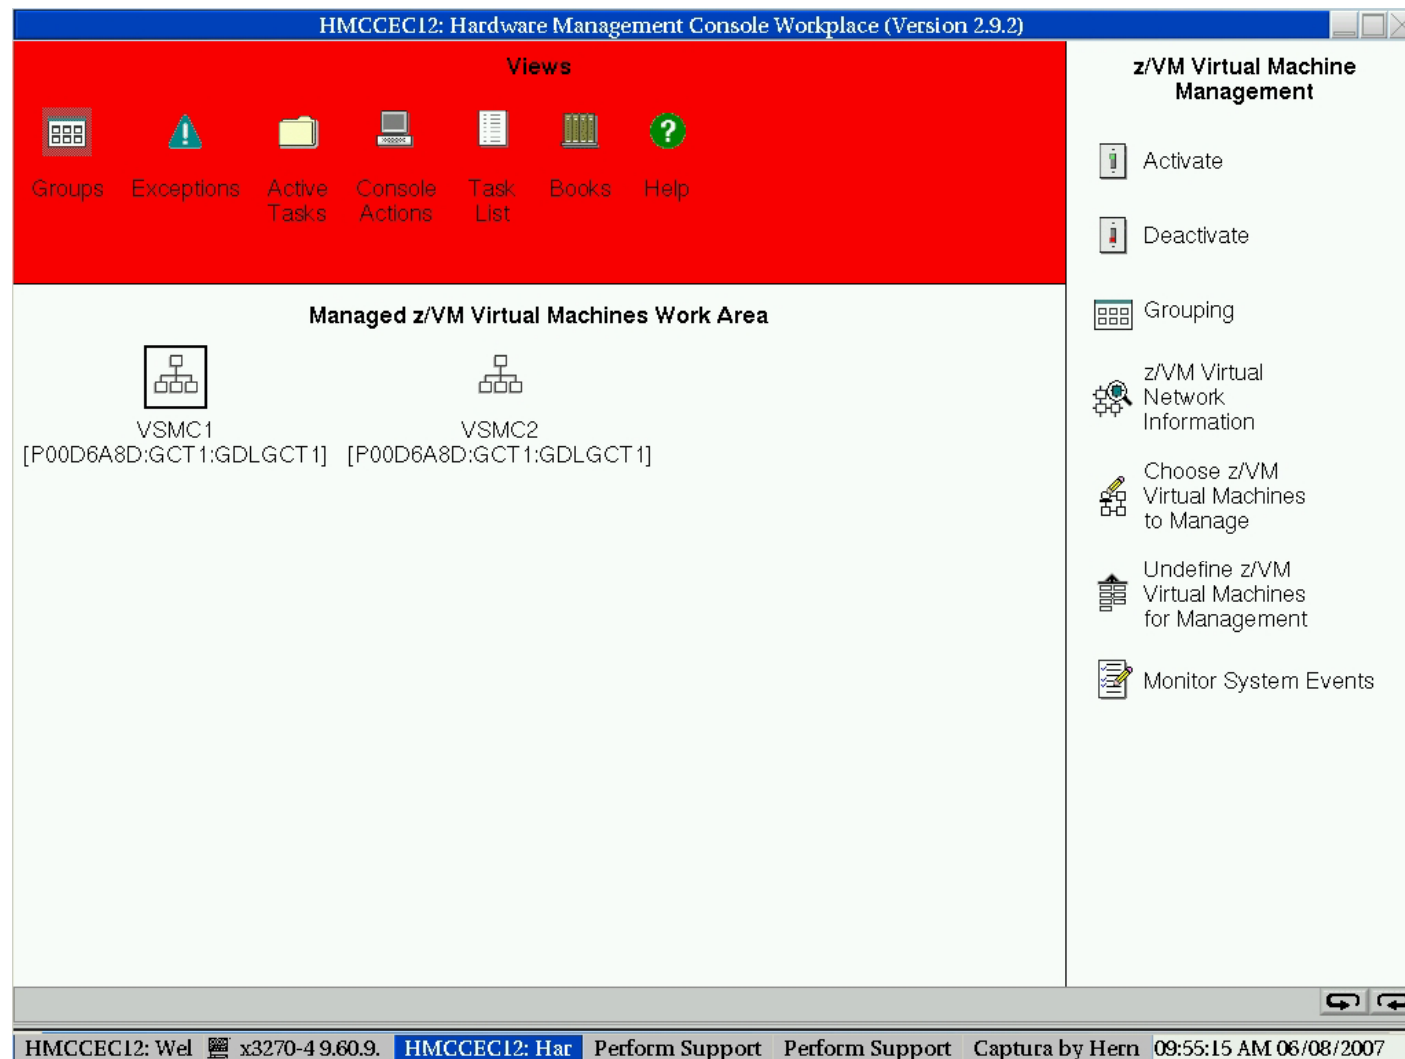
z/VM Integrated Systems Management ...



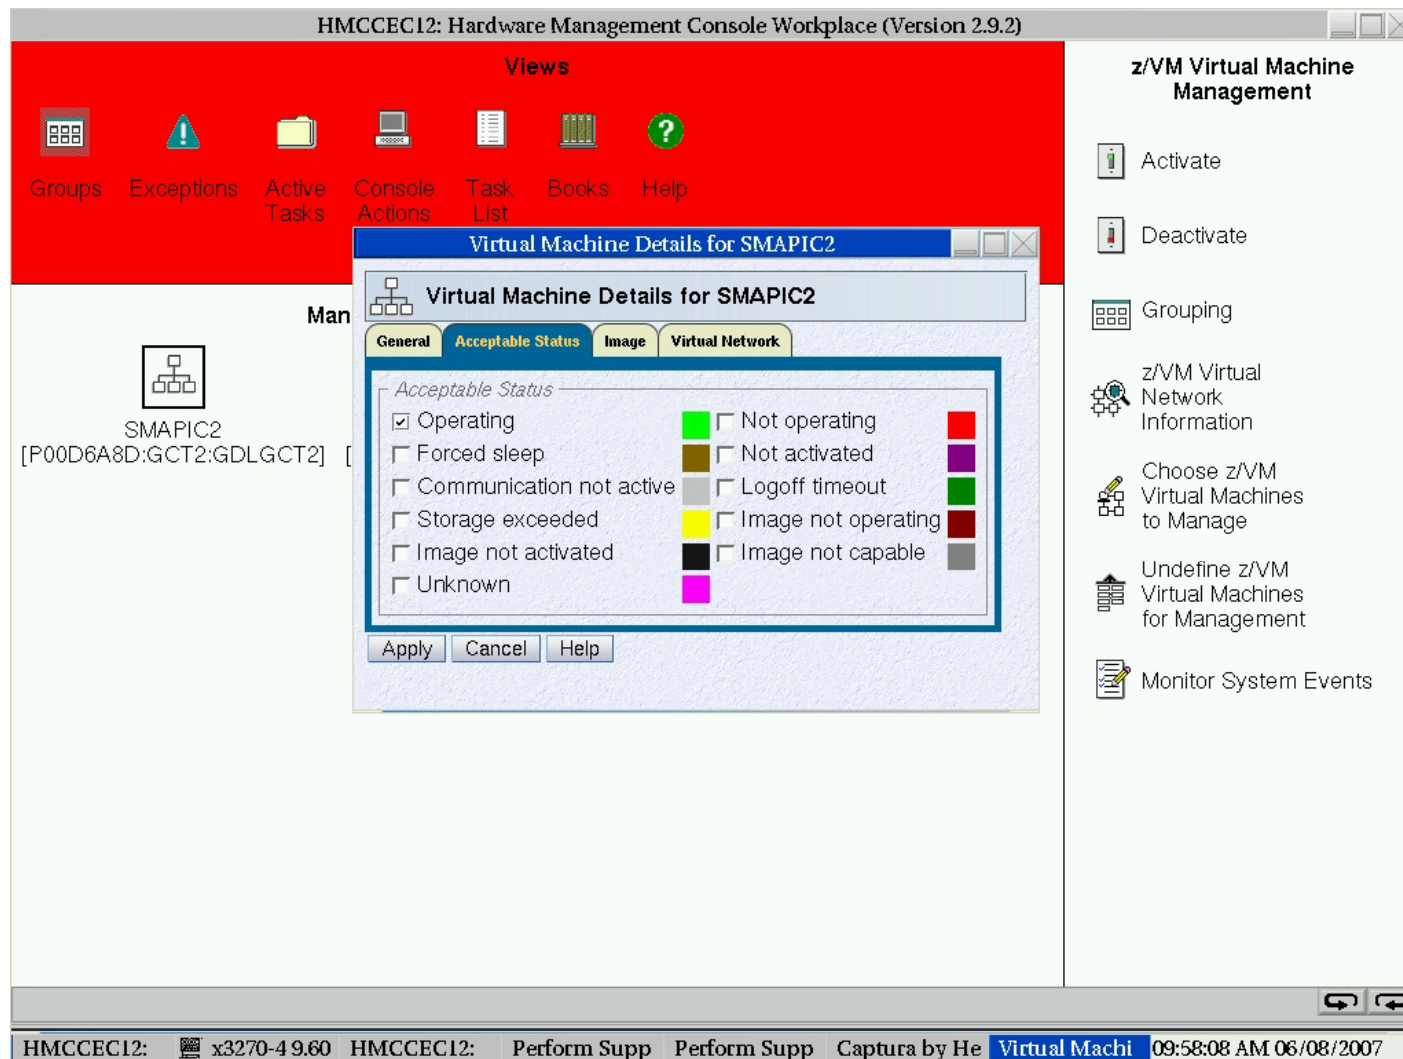
z/VM Integrated Systems Management ...



z/VM Integrated Systems Management ...



z/VM Integrated Systems Management ...



z/VM Integrated Systems Management ...

HMCCEC12: Hardware Management Console Workplace (Version 2.9.2)

Views

Groups Exceptions Active Tasks Console Actions Task List Books Help

Virtual Machine Details for SMAPIC2

Virtual Machine Details for SMAPIC2

General Acceptable Status Image Virtual Network

Memory Size: 64 MB
 Share Type: Relative
 Share Value: 100
 Number of CPUs: 1

CPU Information

ID	Number	Status
FF026A8D20948000	0	Base

Device Information

Address	Type
0009	Console
000C	Reader
000D	Punch
000E	Print
0100	DASD

Apply Cancel Help

z/VM Virtual Machine Management

Activate
 Deactivate
 Grouping
 z/VM Virtual Network Information
 Choose z/VM Virtual Machines to Manage
 Undefine z/VM Virtual Machines for Management
 Monitor System Events

HMCCEC12: x3270-4 9.60 HMCCEC12: Perform Supp Perform Supp Captura by He Virtual Machi 09:58:32 AM 06/08/2007

z/VM Integrated Systems Management ...

- **No configuration or setup required**
- **Supported on any System z**
- **HMC must be at or above level 2.9.2**
- **Support Element (SE) must be at or above**
 - ▶ **2.9.2 for z9 EC and z9 BC**
 - ▶ **1.8.2 for z990 and z890**
 - ▶ **1.7.3 for z900 and z800**
- **Requires MCLs: HMC - G40969.050, SE: G40965.074, and PTFs for APARs VM64233 and VM64234**
- **Z10 – no MCLs needed**

Virtual Machine Resource Manager (VMRM)

- **Manages performance of selected virtual machines**
 - ▶ based on customer-defined goals in configuration file
- **Part of CMS component; executables on MAINT's 193 disk**
- **Implemented as a Service Virtual Machine: VMRMSVM**
 - ▶ PROFILE EXEC starts the server by calling the IRMSERV EXEC (may also be invoked from the command line)
 - ▶ IRMSERV reads statements in the VMRM configuration file set up by a System Administrator
 - Workload definitions (may include multiple virtual machines)
 - Goal specifications for CPU and DASD I/O velocity goals
 - Association between Workloads and Goals and importance
 - New configuration file to put into production at a later time
 - Cooperative Memory Management notifications

Virtual Machine Resource Manager (VMRM)...

- **Uses VM monitor data**
 - ▶ Obtains 1-minute interval measurements of virtual machine resource consumption
 - ▶ Computes the achievement levels of interest (actuals) for each workload
- **VMRM adjusts user CPU shares or I/O priority based on:**
 - ▶ Velocity goals set for the user's workload class
 - ▶ Virtual machine CPU and/or I/O actual achievement levels
 - ▶ Uses CP SET SHARE and SET IOPRIORITY commands for adjustments
- **Provides Application Monitor Data (APPLDATA)**
- **Systems Management APIs provided for VMRM configuration file management**


VMRM-CMM (Cooperative Memory Management)

- A collaboration between VM and Linux to optimize memory management when memory constraint is detected
- Linux guests to be notified are identified in the VMRM configuration file, treated with equal priority
- VMRM tracks system memory utilization/demand and computes target "resident footprint" for each guest
- VMRM sends SMSG to guests to adjust footprint (SHRINK message)
- Guest device driver receives messages and uses existing guest logic to return the least valuable pages
 - ▶ Subsequent SHRINK messages may also indicate how much memory to reclaim (when SHRINK value is less than previous value)
- VMRM-CMM support in base z/VM 5.3.0, and enablement APAR VM64085 for z/VM 5.2.0 (see Reference page for Linux updates)

Linux and z/VM Technology Exploitation

Cooperative Memory Management

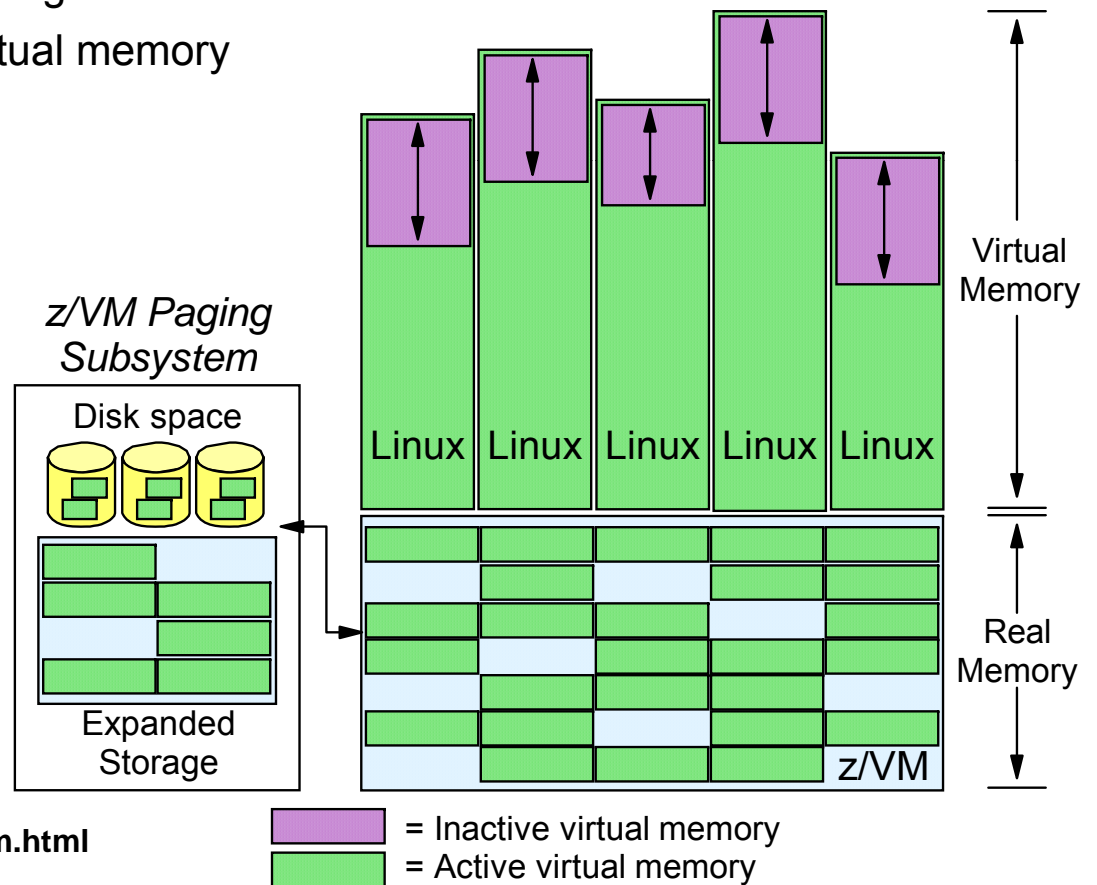
- Problem scenario: virtual memory utilization far exceeds real memory availability
- z/VM Control Program paging operations become excessive
- Overall system performance and guest throughput suffers



Linux and z/VM Technology Exploitation

Cooperative Memory Management

- Solution: real memory constraint detected by z/VM Virtual Machine Resource Manager
- Linux images signaled to reduce virtual memory consumption
- Linux memory pages are released
- Demand on real memory and z/VM paging subsystem is reduced
- Helps improve overall system performance and guest image throughput
- z/VM V5.2 support available with PTF for APAR VM64085
- Linux support currently available with SLES 9 SP3 and 10 RHEL 5.1



Learn more at:
<http://www.vm.ibm.com/sysman/vmr/vmrmmcm.html>

Linux and z/VM Technology Exploitation

CMMA - Collaborative Memory Management Assist

- Linux uses new ESSA instruction to mark page state changes
- z9 EC and z9 BC processors, z10
- CP knows when Linux releases storage and can select those pages for removal
- CP recognizes clean disk cache pages, Linux can reconstruct them. These do not have to be paged out by CP.

Learn more at:

<http://www.ibm.com/common/ssi/fcgi-bin/ssialias?infotype=an&subtype=ca&appname=GPA&htmlfid=897/ENUS207-019>

Monitoring

- **Periodic collection of performance data for early detection of problems and their cause**
- **System performance measurement commands available**
 - ▶ **MONITOR commands**
 - **controls selection, collection, and reporting of data from the host system**
 - **Data communicated to applications using the *Monitor system service**
 - ▶ **INDICATE command**
 - **gives information on the use of and contention for system resources**
- **Performance Toolkit for VM, IBM Operations Manager, OMEGAMON XE, and VMRM have monitoring capabilities**

Performance Toolkit for VM

- **z/VM priced feature based on the FCON/ESA product**
 - ▶ Priced on a one-time charge, per-processor basis
 - ▶ Replaces VMPRF and RTM
- **Functional highlights**
 - ▶ Provides an immediate view of system performance
 - ▶ Post processes its own history files or VM Monitor data
 - ▶ Threshold monitoring, User loop detection
 - ▶ Can monitor remote systems
 - ▶ View results using web browsers or PC-based 3270 emulator graphics
 - ▶ Processes Linux APPLDATA and Linux data obtained from the Resource Management Facility (RMF)
 - ▶ Report for SCSI devices

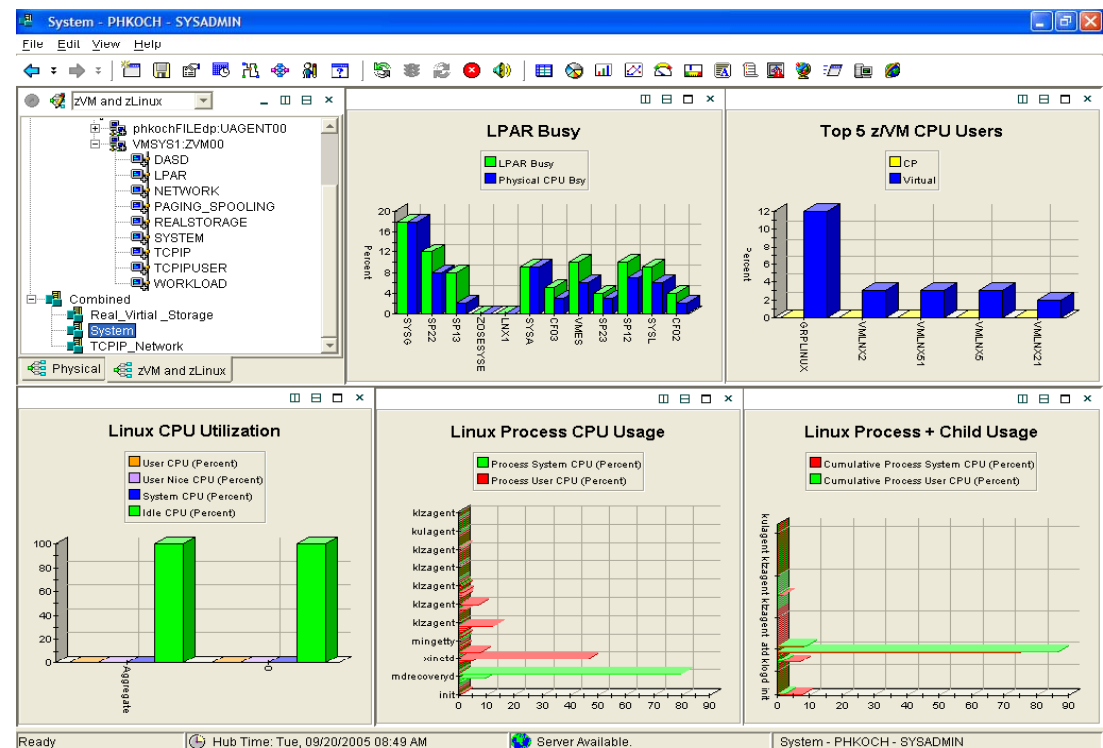
IBM Tivoli OMEGAMON XE on z/VM and Linux

- **Priced Product that collects system metrics and performance data about z/VM and Linux**
 - ▶ **z/VM agent is part of IBM Tivoli Management Services suite of products**
- **Can view data collected from multiple systems**
 - ▶ **Monitors all systems from a single, integrated browser with capability of customized filters**
 - ▶ **Real time performance data: memory usage, I/O, paging, spool usage, TCP/IP network statistics, CPU data.**
 - ▶ **Can view z/VM data obtained from Performance Toolkit for VM product**
 - ▶ **Data collected is displayed by java-based Tivoli Enterprise Portal (TEP) application**

IBM Tivoli OMEGAMON XE for z/VM and Linux

- **Combined product offering**
- **Provides work spaces that**
 - ▶ **Overall system health**
 - ▶ **Workload metrics for logged-in users**
 - ▶ **Individual device metrics**
 - ▶ **LPAR Data**
- **Provides composite views of Linux running on z/VM**
- **Supported on z/VM V5.2 and z/VM V5.3**
 - ▶ **Requires the Performance Toolkit for data collection**
- **Available: 13 October 2006**

*Refer to IBM Software
Announcement 206-251



z/VM Systems Management Products from IBM

- **Software Group priced products for z/VM**
- **New z/VM product release from IBM (announced 15 Aug 2006)**
 - ▶ **IBM Backup and Restore Manager for z/VM V1.2 – August 2006**
- **Other z/VM Systems Management products from IBM**
 - ▶ **IBM Tape Manager for z/VM V1.2 – February 2006**
 - ▶ **IBM Operations Manager for z/VM V1.2 – February 2006**
 - ▶ **IBM Archive Manager for z/VM V1.1 – August 2005**
- **Same pricing model as z/VM V5**
 - ▶ **Value Units based on processors**

Other Systems Management Products

- **ISV's, related products, licensed products websites:**

- ▶ <http://www.ibm.com/vm/vendor>

- ▶ <http://www.ibm.com/vm/related>

- ▶ <http://www.ibm.com/vm/techinfo/lpmigr>

- ▶ <http://www.ibm.com/software/tivoli>

Future Areas of Emphasis

- **Ease-of-Use**
 - ▶ Hypervisor Management
 - ▶ Systems Management APIs, IBM Director enhancements
 - ▶ Installation, service, release-to-release upgrades
 - ▶ HMC media install of Linux z/VM guest
- **Performance Monitoring and Reporting**
 - ▶ Performance Toolkit enhancements
 - ▶ Virtual Machine Resource Manager
 - ▶ Memory management
- **Autonomic computing**
 - ▶ Based on user-supplied definitions
 - ▶ Self-Managing, Self-Healing, High Availability

Summary

- **Systems Management covers many areas of the system and a wide range of tasks**
- **Many system functions and tools available**
 - ▶ both IBM and vendor products
- **Continued commitment to address customer needs**
 - ▶ enhancements each release
- **What are your requirements ?**
 - ▶ **Contact info: brazie@us.ibm.com**

Questions ?

Contact Info: email: brazie@us.ibm.com
phone: (607)429-3399

Webpage: <http://www.vm.ibm.com/sysman>

Linux updates for CMM support:

<http://www-128.ibm.com/developerworks/linux/linux390/linux-2.6.5-s390-34-april2004.html>