#SHAREorg

SHARE

Managing z/VM and Linux Performance Best Practices

Mike Sine
IBM

March 13, 2014
Session Number 14637

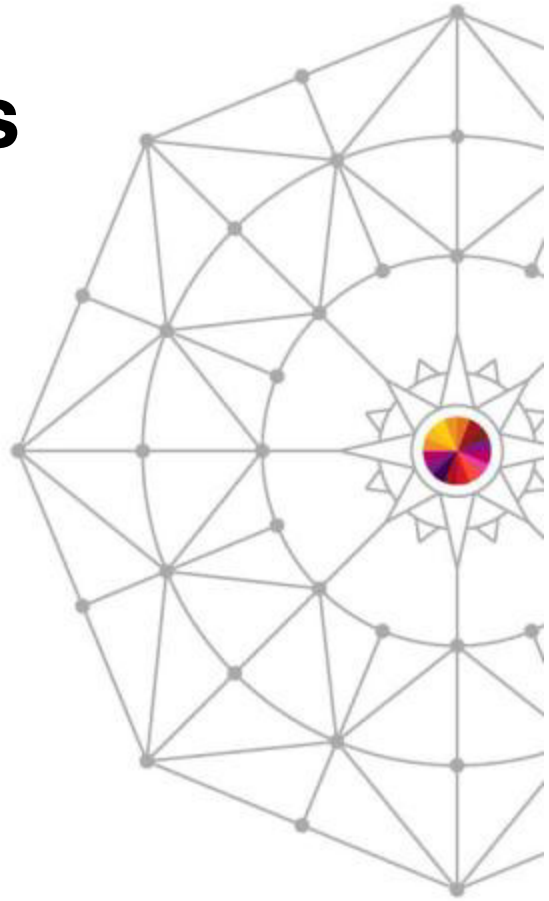

SHARE in Anaheim

Special Notices and Trademarks

Special Notices

This presentation reflects the IBM Advanced Technical Skills organizations' understanding of the technical topic. It was produced and reviewed by the members of the IBM Advanced Technical Skills organization. This document is presented "As-Is" and IBM does not assume responsibility for the statements expressed herein. It reflects the opinions of the IBM Advanced Technical Skills organization. These opinions are based on the author's experiences. If you have questions about the contents of this document, please contact the author at sine@us.ibm.com .

Trademarks

The following are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

IBM, the IBM logo, Candle, DB2, developerWorks, iSeries, Passport Advantage, pSeries, Redbooks, Tivoli Enterprise Console, WebSphere, z/OS, xSeries, zSeries, System z, z/VM.

A full list of U.S. trademarks owned by IBM may be found at http://www.ibm.com/legal/copytrade.shtml.

NetView, Tivoli and TME are registered trademarks and TME Enterprise is a trademark of Tivoli Systems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, Internet Explorer, and the Windows logo are registered trademarks of Microsoft Corporation in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Intel and Pentium are registered trademarks and MMX, Pentium II Xeon and Pentium III Xeon are trademarks of Intel Corporation in the United States and/or other countries.

Other company, product and service names may be trademarks or service marks of others.

AGENDA

- Introduction
- Monitoring Requirements
  - Virtual Linux and z/VM performance considerations
  - Don’t forget the hardware
  - Integration from hardware – systems – applications Persistent historical views
- Enterprise Management
- Operational Requirements
  - Centralized Control
  - Including all Enterprise Virtual Machines
- Integrating Monitoring and Operations
- Bringing it all together

Virtual Linux servers have unique challenges versus running on physical machines.

- z/VM System Programmers and Linux Administrators may not be in the same organization.
- We find that it is easy to over allocate resources; therefore, our monitoring examines resource usage of hardware, hypervisor, as well as the virtual machine. Real-time and historical metrics demonstrate peaks periods as well as average runtimes.
- Cross-platform virtualization increases these challenges

AGENDA

- Introduction
- Monitoring Requirements
  - Virtual Linux and z/VM performance considerations
  - Don’t forget the hardware
  - Integration from hardware – systems – applications Persistent historical views
- Enterprise Management
- Operational Requirements
  - Centralized Control
  - Including all Enterprise Virtual Machines
- Integrating Monitoring and Operations
- Bringing it all together

OMEGAMON XE on z/VM and Linux agents with ITM example

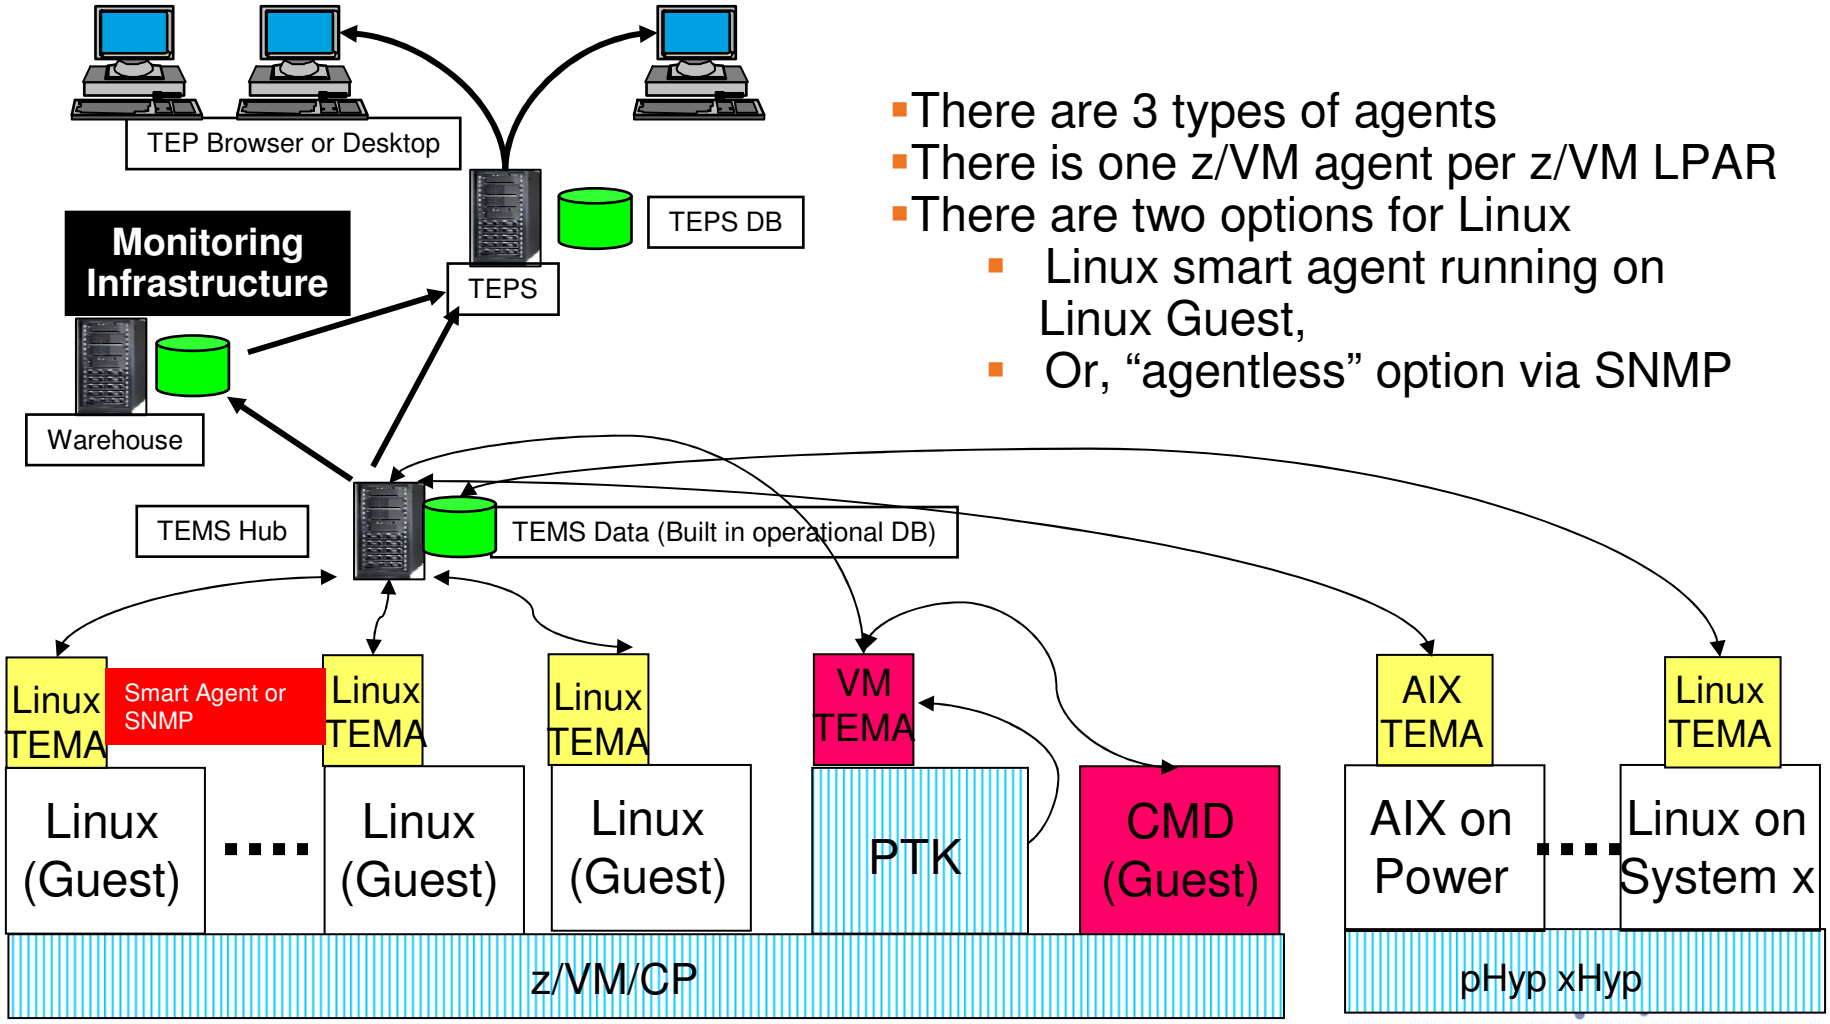

- There are 3 types of agents
- There is one z/VM agent per z/VM LPAR
- There are two options for Linux
  - Linux smart agent running on Linux Guest,
  - Or, "agentless" option via SNMP

Highly scalable across the entire Enterprise

An Integrated Monitoring Approach

- Provides performance monitoring for z/VM and Linux guests
- Executes automated actions in response to defined events or situations (monitoring without automation is overhead)
- Integrates and Scales well across Enterprise for central control and trending:
  - Specifically focused on z/VM and Linux guests
  - Able to integrate z/VM and Linux into Enterprise Solution
  - Data warehousing for trend analysis (interactive and batch)

Workspaces to Manage z/VM and Linux

z/VM

- **Processors**
- **SYSTEM Utilization, spinlocks**
- **Workload**
  - Linux Appldata
  - Scaled & total CPU values
- **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 – Users**
- **Resource Constraint (Wait states)**
- **System Health**

Linux

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

Have I allocated enough Virtual CPUs to my guest?

- Do not define more virtual CPUs for a Linux guest than are needed.
  - The use of more than one processor requires software locks so that data or control blocks are not updated by more than one processor at a time.
- Linux makes use of a global lock, and when that lock is held, if another processor requires that lock, it spins.
  - Set the number of virtual processors based on need and not simply match the number of real that are available.
  - Careful when cloning as some Linux guests require more Virtual CPUs (ex: Running Websphere, Oracle) than others.

Aggregate monitoring of Virtual CPUs

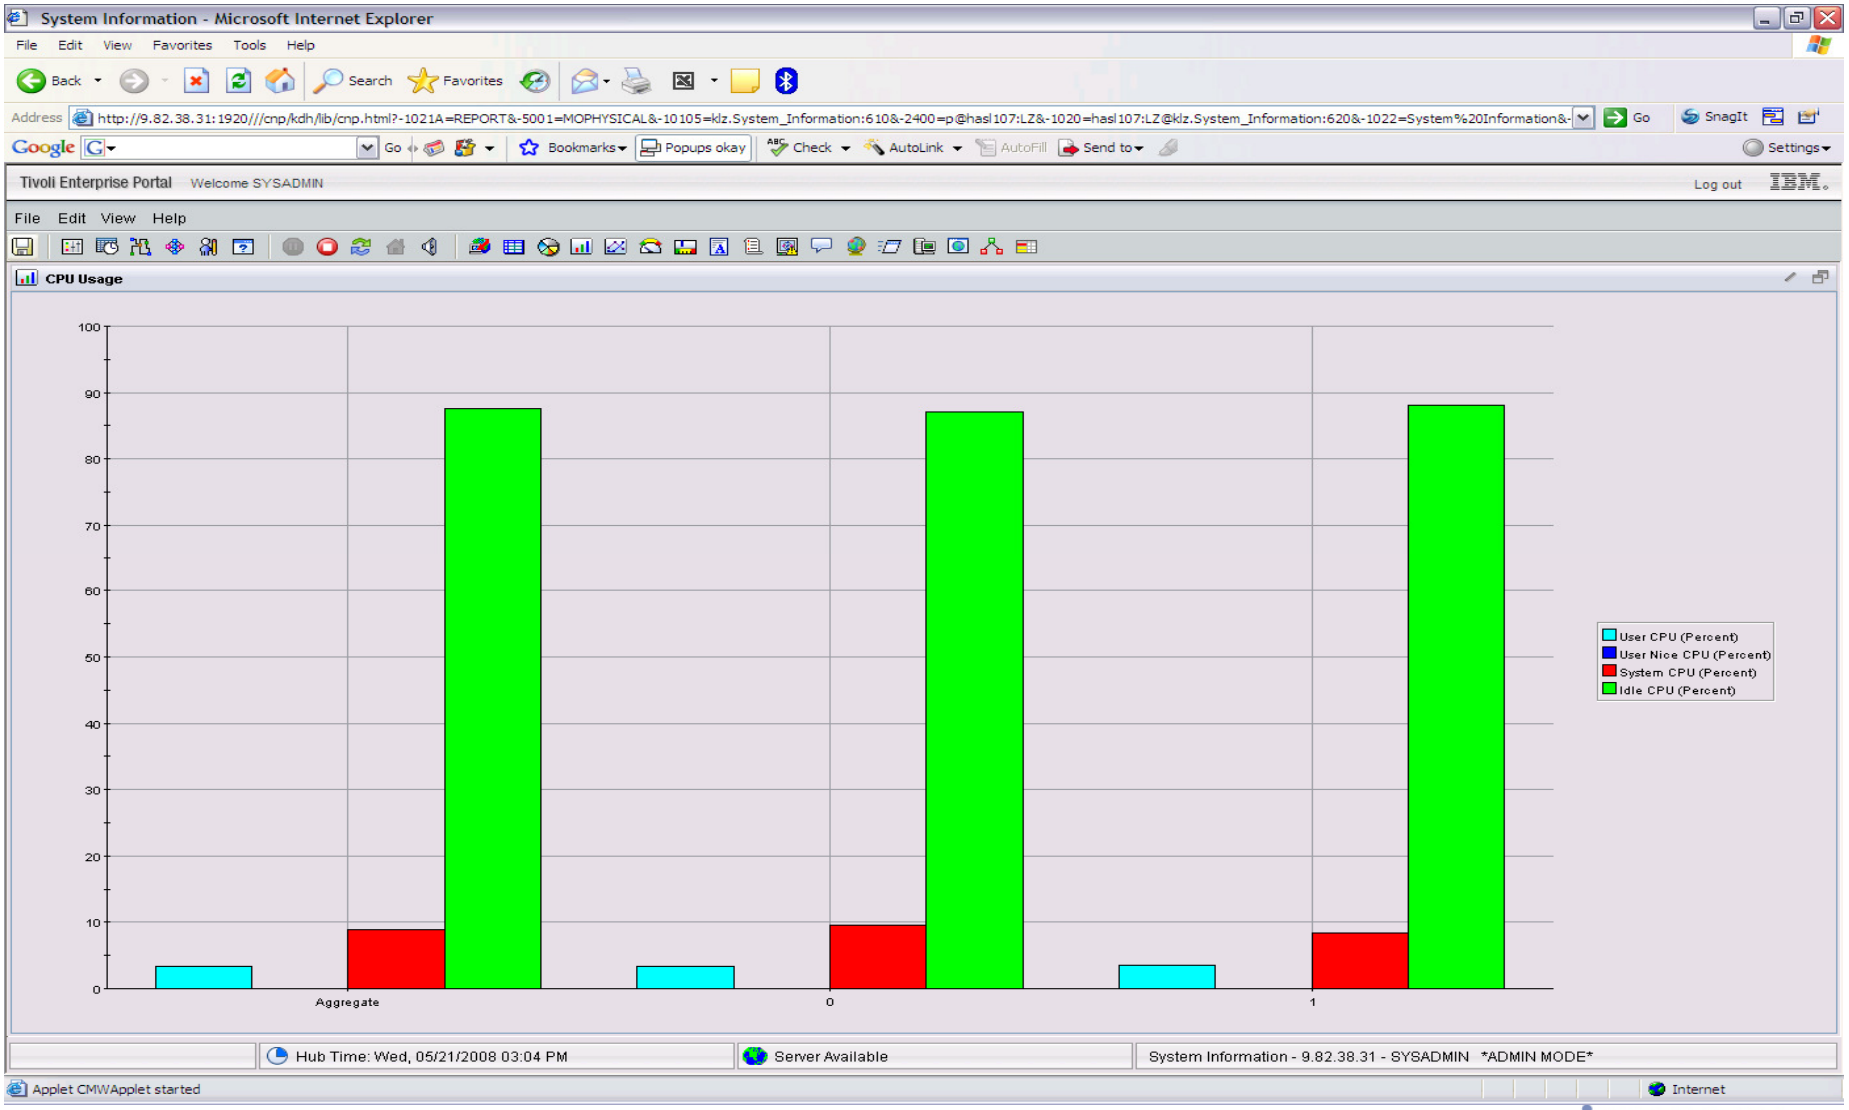

z/VM Processor Utilization

- **Total Processor Utilization** This is the processor utilization from the VM perspective and includes CP, VM System, and Virtual CPU time.
- **System Time:** This is the processor time used by the VM control program for system functions that are not directly related to any one virtual machine. This should be less than 10% of the total.
- **CP Processor Time:** This is the processor time used by the VM control program in support of individual virtual machines.
- **Virtual Processor Time: (Emulation Time):** This is processor time consumed by the virtual machine and the applications within it.
- **Total to Virtual Ratio** The ratio of total processor time to virtual processor time is often used as an indicator of z/VM efficiency or overhead. The closer to 1.0, the better the z/VM efficiency. RoT: Should explore causes of a ratio over 1.30.

System Processor Utilization Workspace

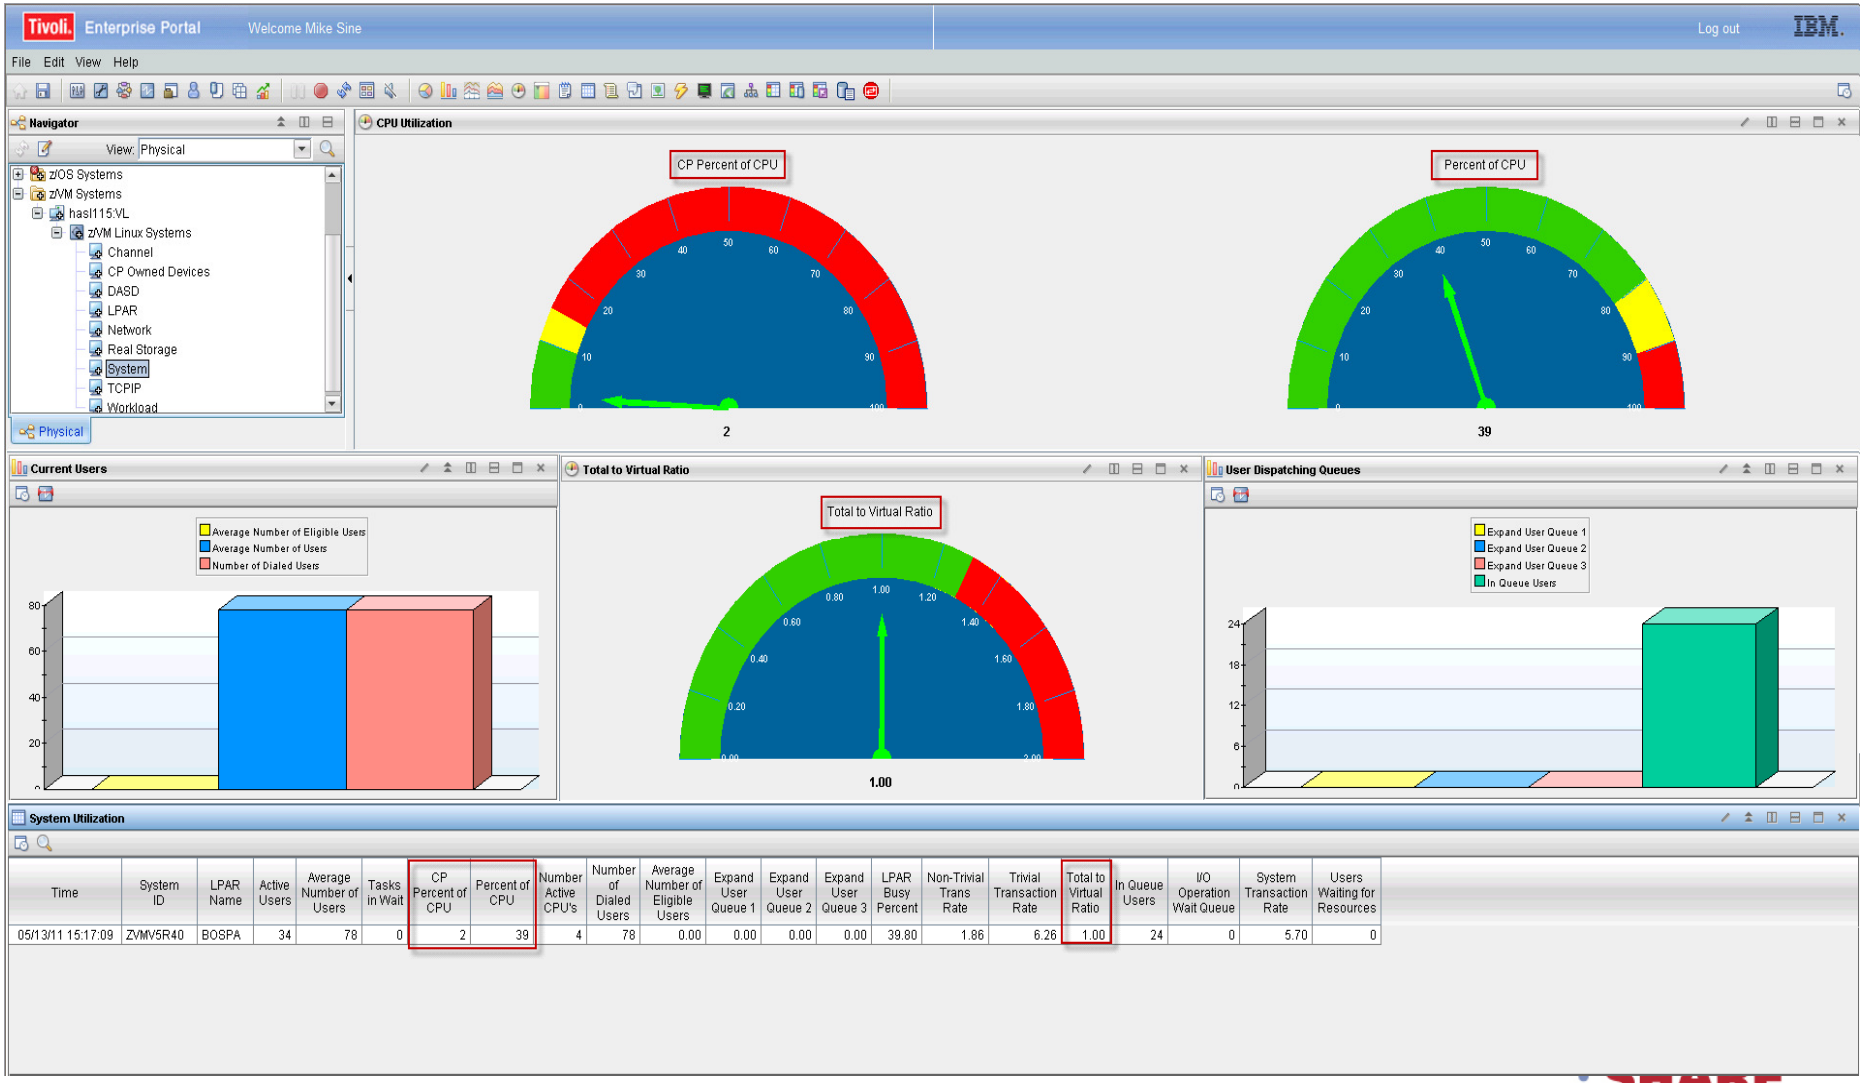

Time	System ID	LPAR Name	Active Users	Average Number of Users	Tasks in Wait	CP Percent of CPU	Percent of CPU	Number Active CPU's	Number of Dialed Users	Average Number of Eligible Users	Expand User Queue 1	Expand User Queue 2	Expand User Queue 3	LPAR Busy Percent	Non-Trivial Trans Rate	Trivial Transaction Rate	Total to Virtual Ratio	In Queue Users	I/O Operation Wait Queue	System Transaction Rate	Users Waiting for Resources
05/13/11 15:17:09	ZVMV5R40	BOSPA	34	78	0	2	39	4	78	0.00	0.00	0.00	0.00	39.80	1.86	6.26	1.00	24	0	5.70	0

z/VM Workload Workspace

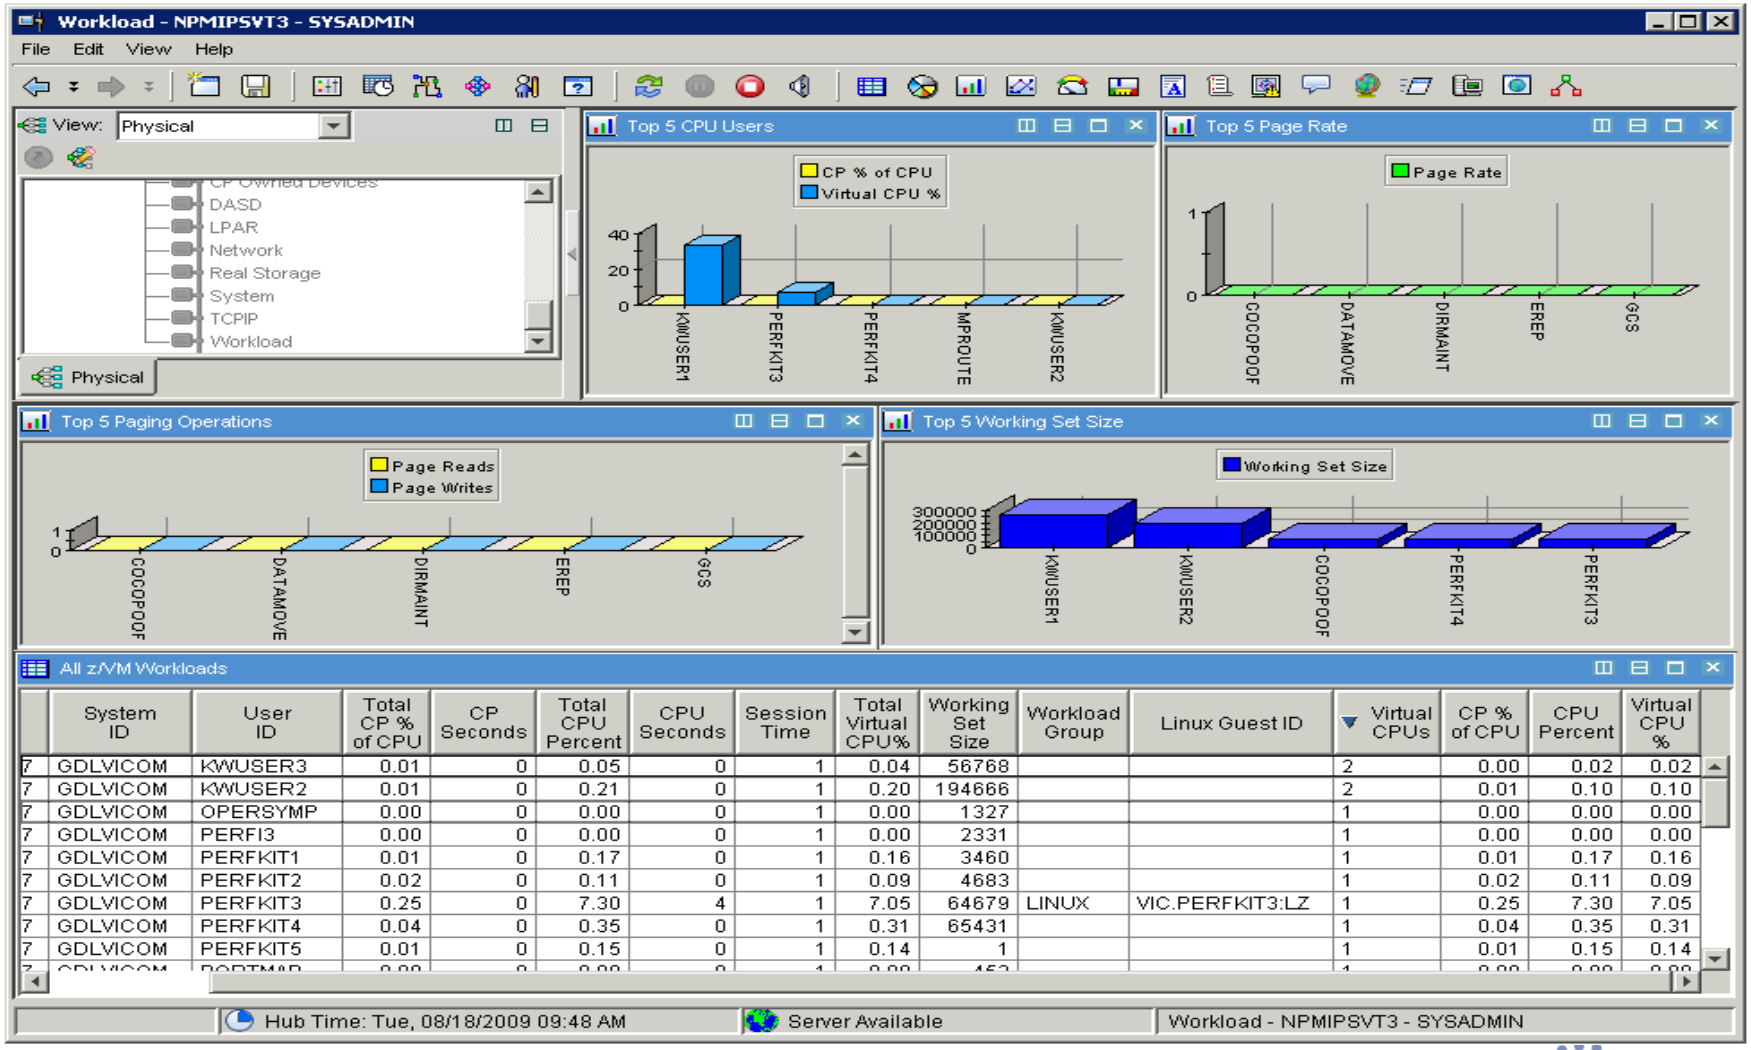

	System ID	User ID	Total CP % of CPU	CP Seconds	Total CPU Percent	CPU Seconds	Session Time	Total Virtual CPU%	Working Set Size	Workload Group	Linux Guest ID	Virtual CPUs	CP % of CPU	CPU Percent	Virtual CPU %
7	GDLVICOM	KWUSER3	0.01	0	0.05	0	1	0.04	56768			2	0.00	0.02	0.02
7	GDLVICOM	KWUSER2	0.01	0	0.21	0	1	0.20	194666			2	0.01	0.10	0.10
7	GDLVICOM	OPERSYMP	0.00	0	0.00	0	1	0.00	1327			1	0.00	0.00	0.00
7	GDLVICOM	PERFI3	0.00	0	0.00	0	1	0.00	2331			1	0.00	0.00	0.00
7	GDLVICOM	PERFKIT1	0.01	0	0.17	0	1	0.16	3460			1	0.01	0.17	0.16
7	GDLVICOM	PERFKIT2	0.02	0	0.11	0	1	0.09	4683			1	0.02	0.11	0.09
7	GDLVICOM	PERFKIT3	0.25	0	7.30	4	1	7.05	64679	LINUX	VIC.PERFKIT3:LZ	1	0.25	7.30	7.05
7	GDLVICOM	PERFKIT4	0.04	0	0.35	0	1	0.31	65431			1	0.04	0.35	0.31
7	GDLVICOM	PERFKIT5	0.01	0	0.15	0	1	0.14	1			1	0.01	0.15	0.14

Spin Lock Wait

- **Time Spinning on Locks Percent:**
  - The percentage of time processors spend spinning on formal spin locks. RoT: Should be less than 10%.
  - Increases as number of logical processors increases.

Spinlock Workspace

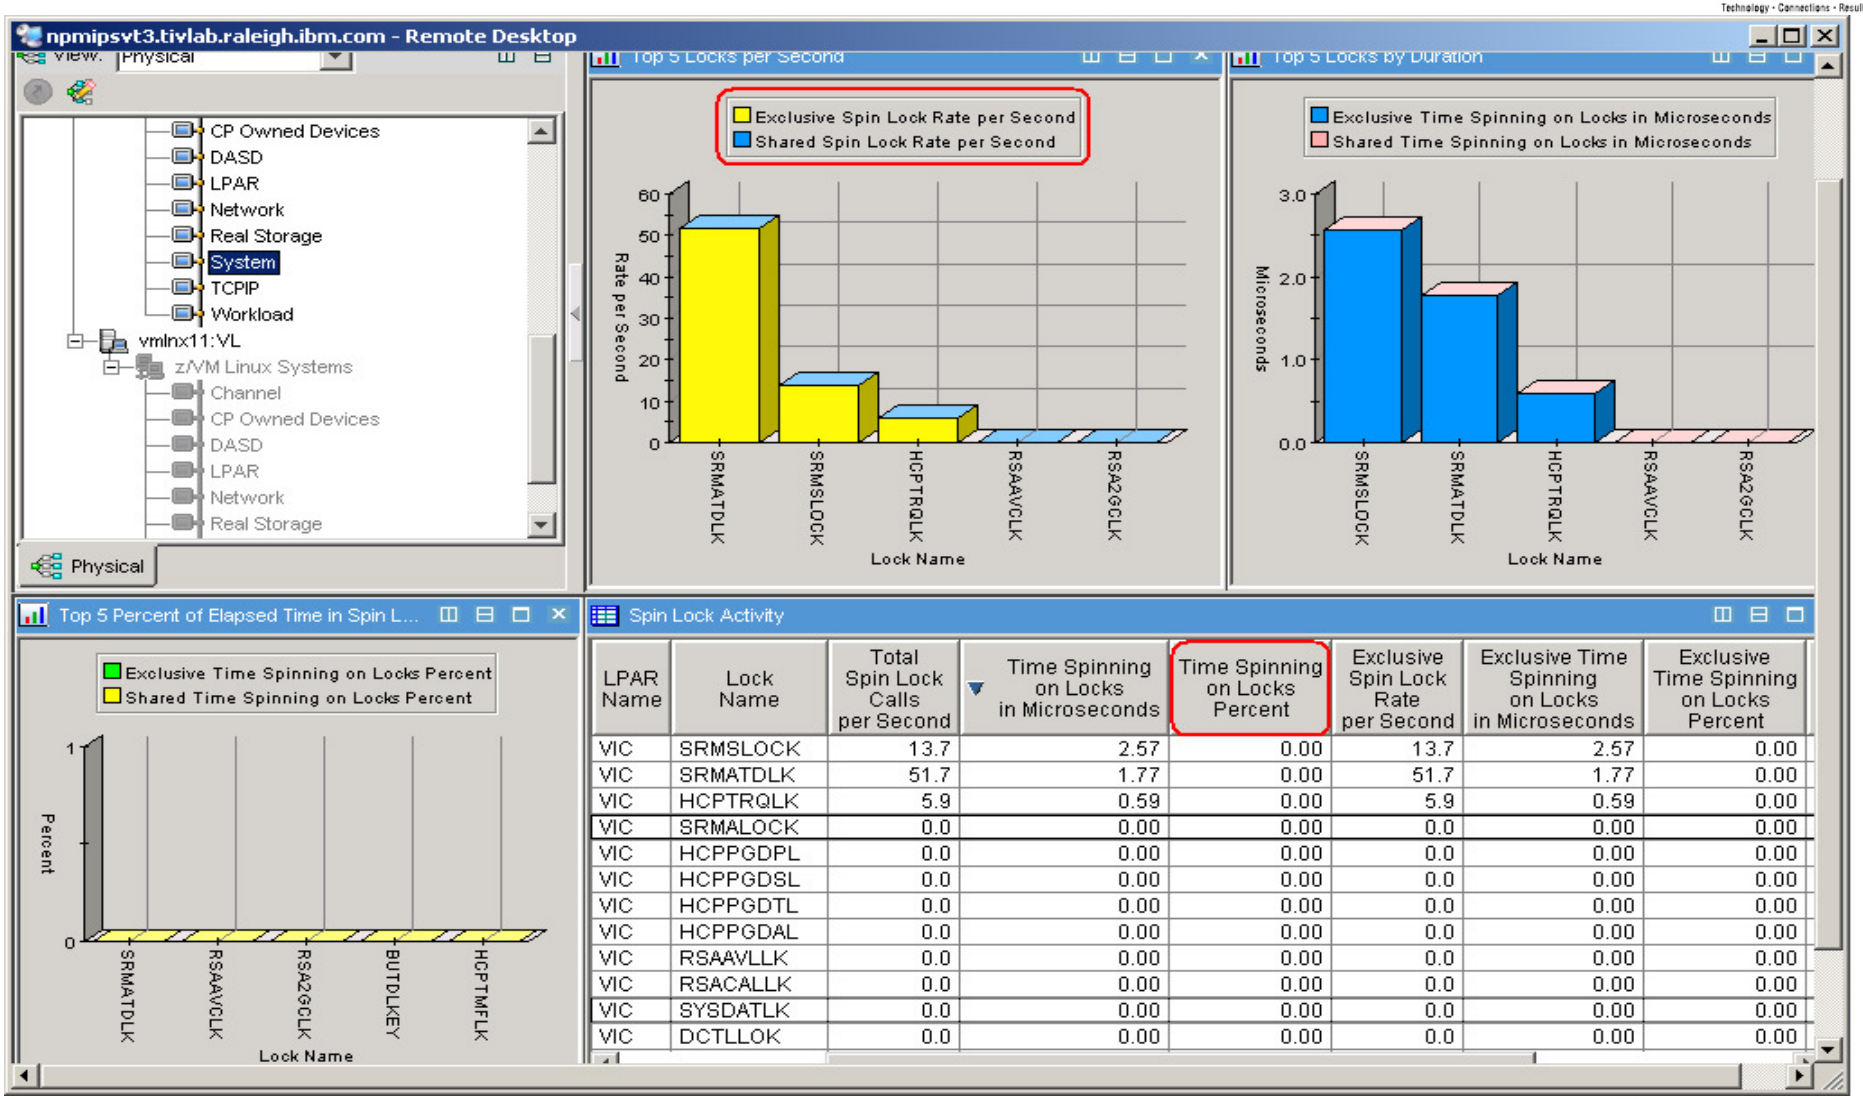

LPAR Name	Lock Name	Total Spin Lock Calls per Second	Time Spinning on Locks in Microseconds	Time Spinning on Locks Percent	Exclusive Spin Lock Rate per Second	Exclusive Time Spinning on Locks in Microseconds	Exclusive Time Spinning on Locks Percent
VIC	SRMSLOCK	13.7	2.57	0.00	13.7	2.57	0.00
VIC	SRMATDLK	51.7	1.77	0.00	51.7	1.77	0.00
VIC	HCPTRQLK	5.9	0.59	0.00	5.9	0.59	0.00
VIC	SRMALOCK	0.0	0.00	0.00	0.0	0.00	0.00
VIC	HCPPGDPL	0.0	0.00	0.00	0.0	0.00	0.00
VIC	HCPPGDSL	0.0	0.00	0.00	0.0	0.00	0.00
VIC	HCPPGDTL	0.0	0.00	0.00	0.0	0.00	0.00
VIC	HCPPGDAL	0.0	0.00	0.00	0.0	0.00	0.00
VIC	RSAAVLLK	0.0	0.00	0.00	0.0	0.00	0.00
VIC	RSACALLK	0.0	0.00	0.00	0.0	0.00	0.00
VIC	SYSDATLK	0.0	0.00	0.00	0.0	0.00	0.00
VIC	DCTLLOK	0.0	0.00	0.00	0.0	0.00	0.00

Is my Linux guest sized correctly?

- In general, do not define the Linux virtual machine larger than you need.
  - Excessive virtual machine sizes negatively impact performance.
  - Linux uses any extra storage for caching of data. For shared resources, this is an impact.
  - Reduce the size of the Linux guest until it starts to swap (use VDISK for swap).
  - Monitor swap usage.

Need breakdown of memory use

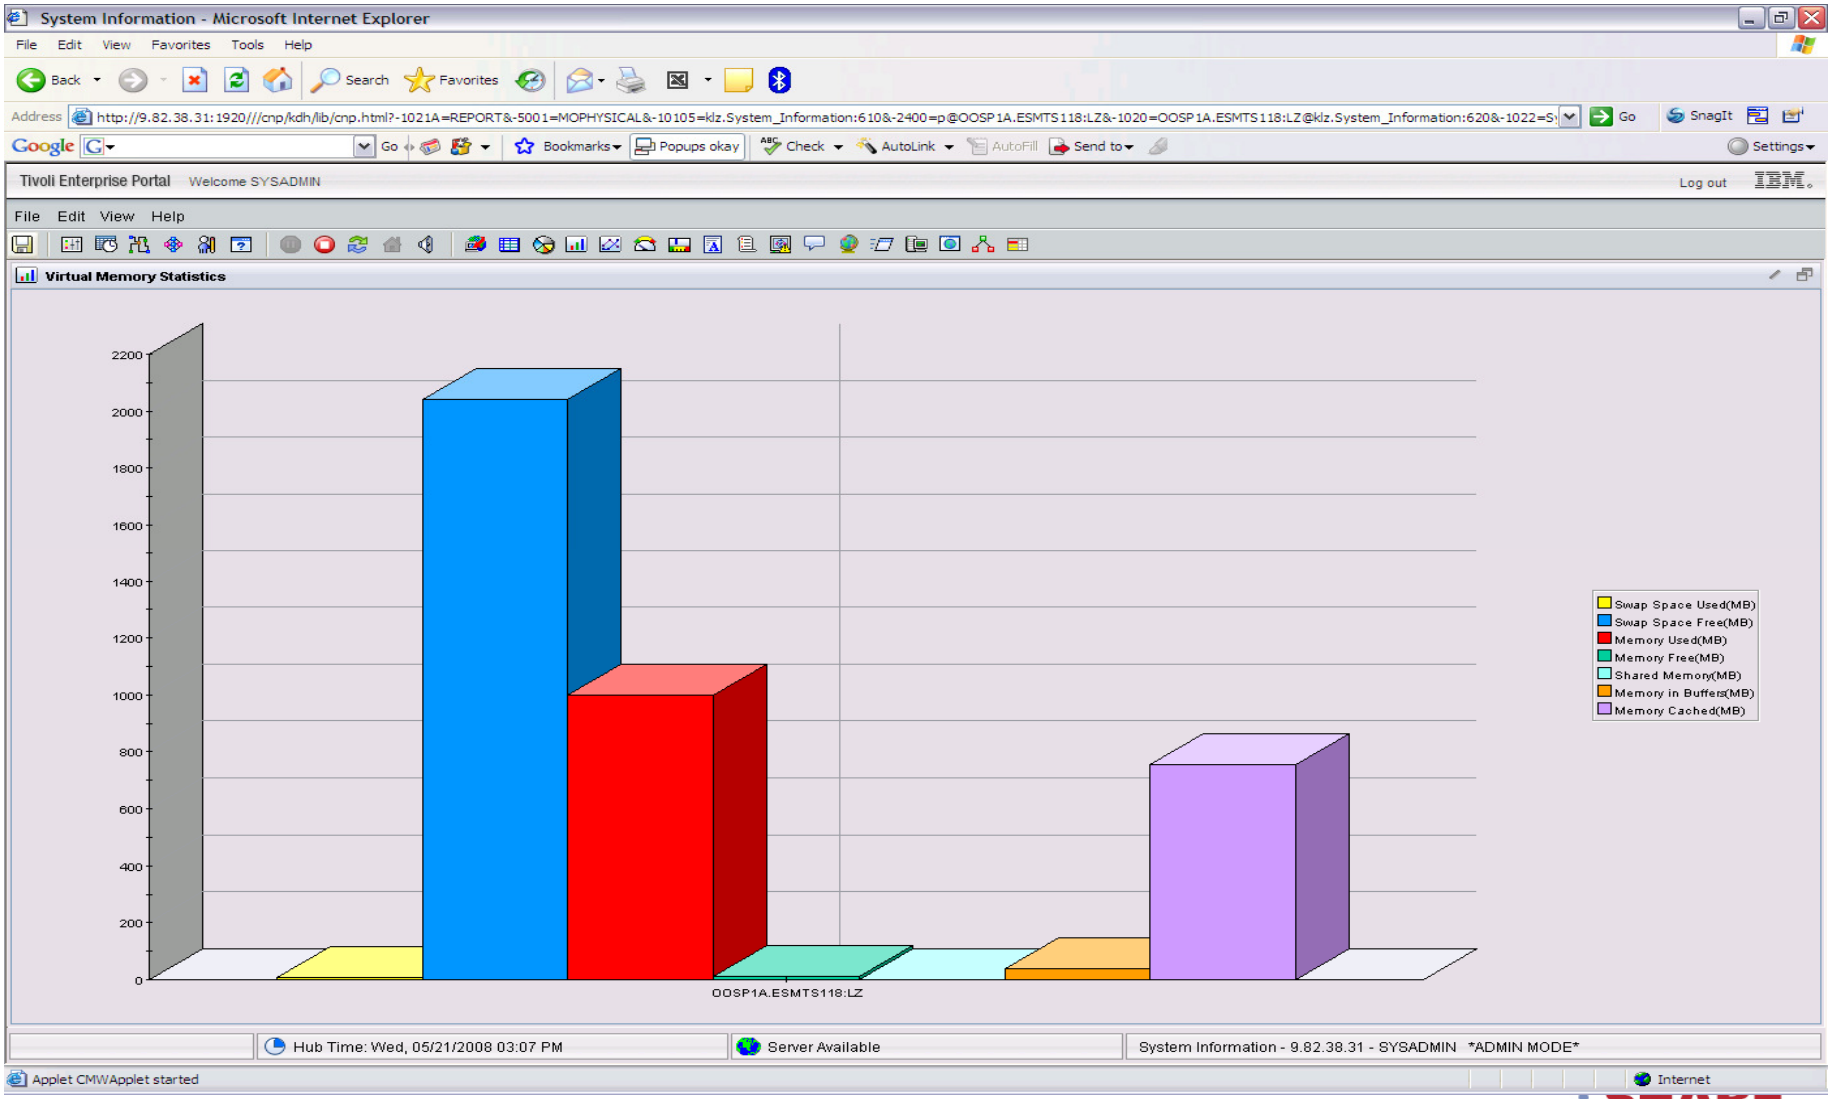

Page/Swap Attributes

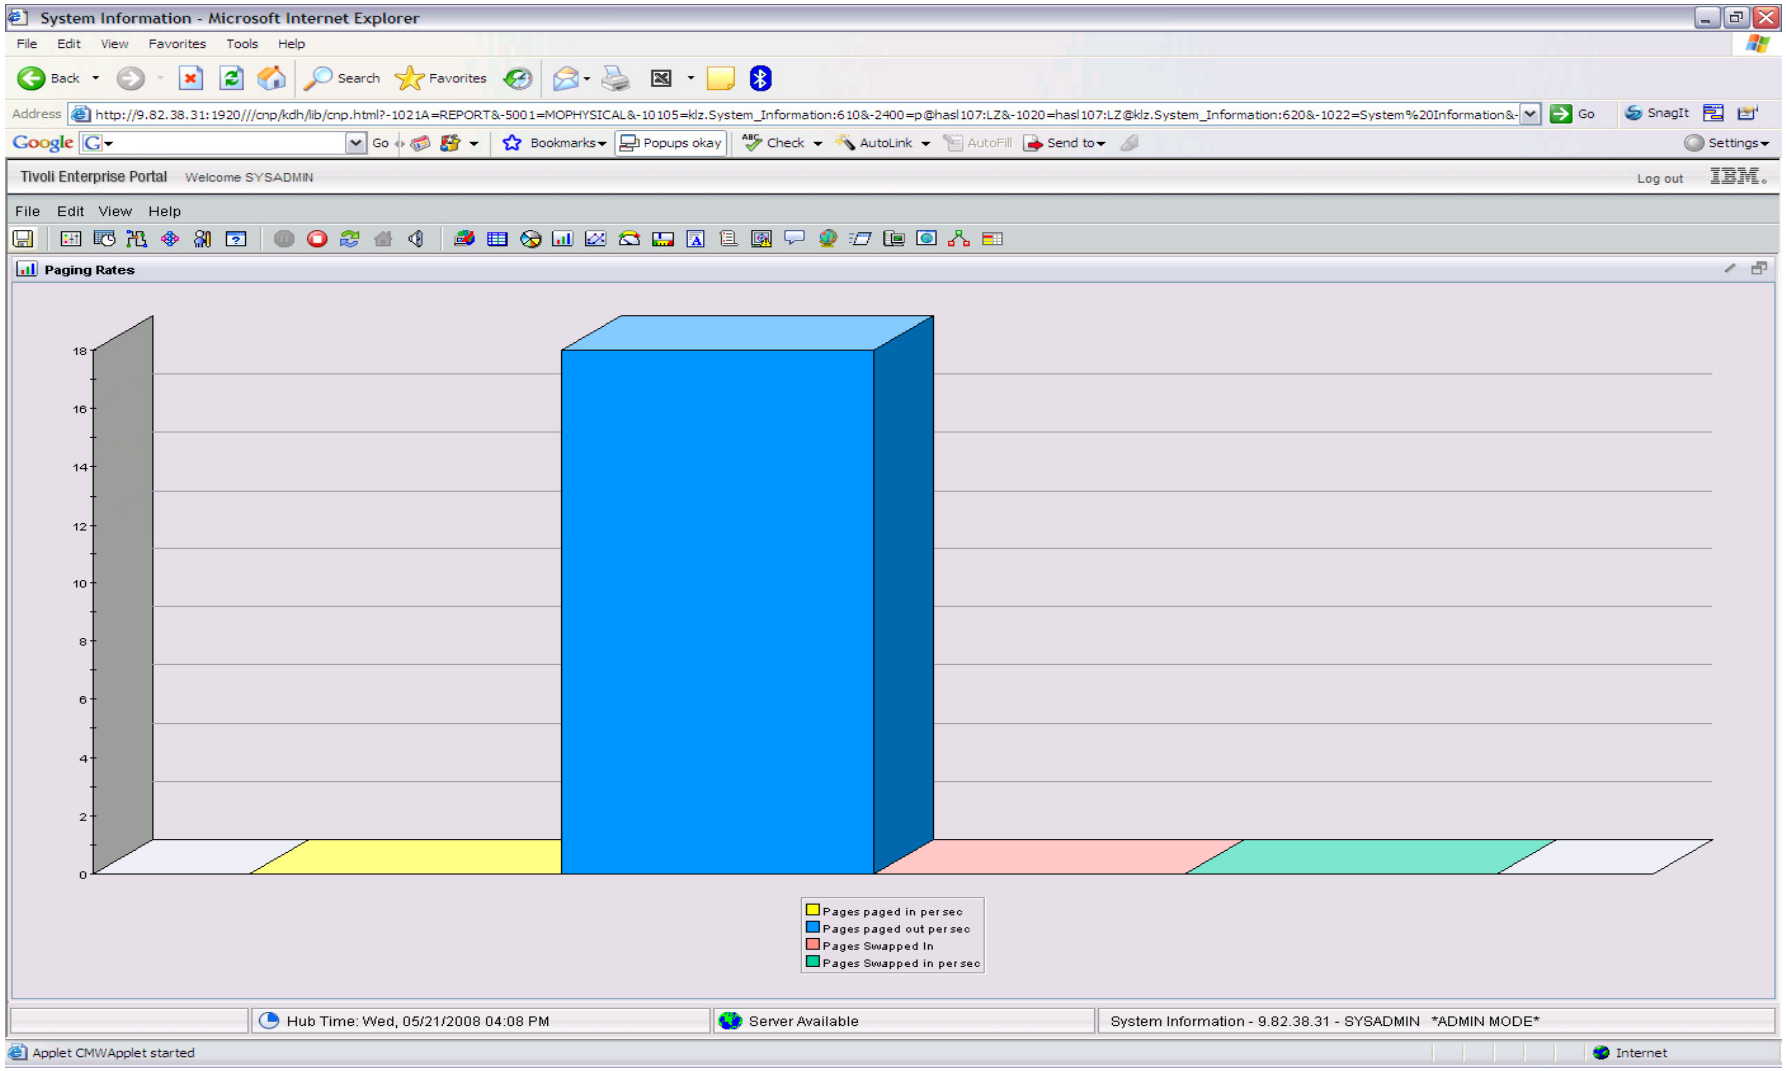

VDISK

- What is it?
  - FBA (Fixed Block Architecture disk) device emulated in-memory
    - Translation: Very fast “device”.
  - High performance paging device for Linux on z.
  - Memory is allocated by CP from the Dynamic Paging Area
  - Allocated only when referenced
    - Allocating a 10 MB device does NOT instantly consume 10 MB of pages.
    - Pages are allocated when needed.
  - Not recommended in a storage-constrained z/VM system.

VDISK Workspace

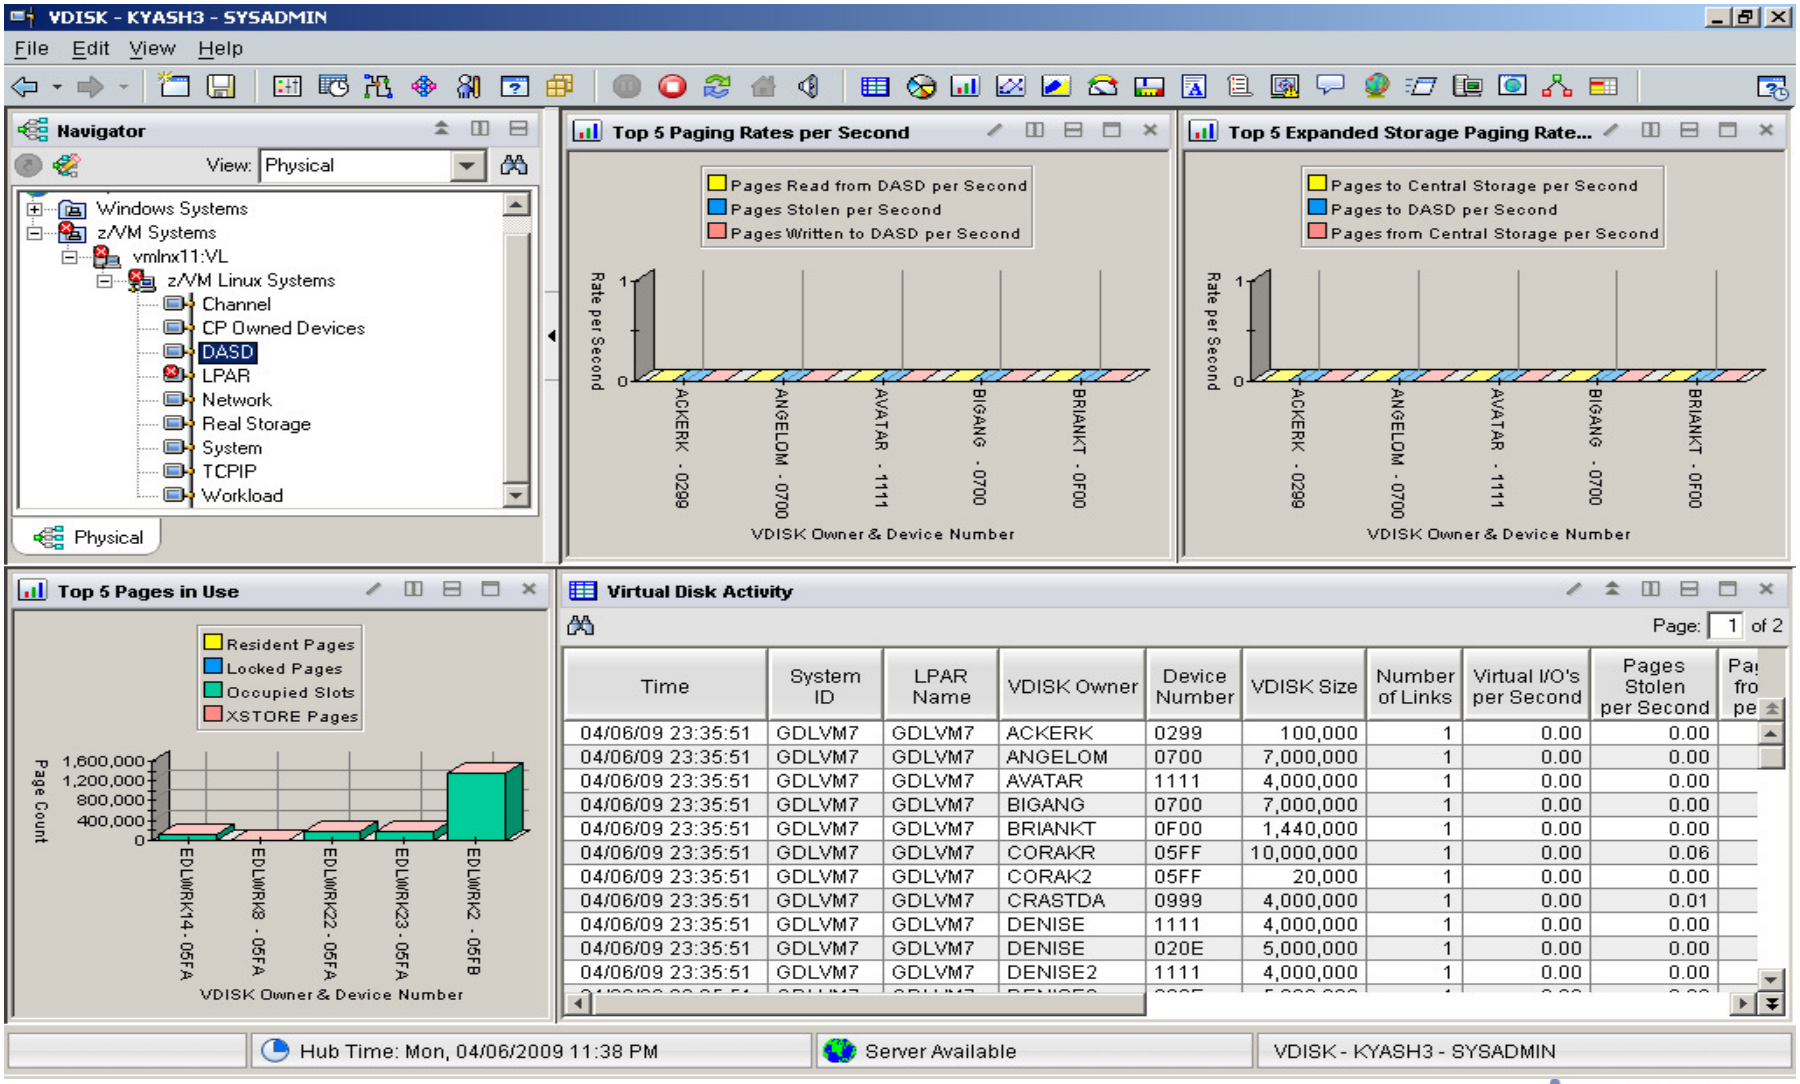

Memory Configuration

- Plan on a virtual to real (V:R) memory ratio in the range of 1.5:1 to 3:1.
- z/VM's architecture still benefits from expanded storage:
  - Serves as high speed cache.
  - Increases consistency of response time.
  - See **http://www.vm.ibm.com/perf/tips/storconf.html** for the gory details.
- Rule of Thumb - start with 20-25% of memory configured as expanded:
  - The lower the paging rate, the lower the amount of expanded storage required.
  - The greater the number of page frames available in central storage above 2GB, the higher the amount of expanded storage required.
  - Some workloads 2–4GB of expanded storage is sufficient, 1GB minimum. However, more and more Linux systems are running heavy workloads and the 20-25% rule still applies.
  - XSTORE obsolete for z/VM 6.3

OMEGAMON Memory Configuration

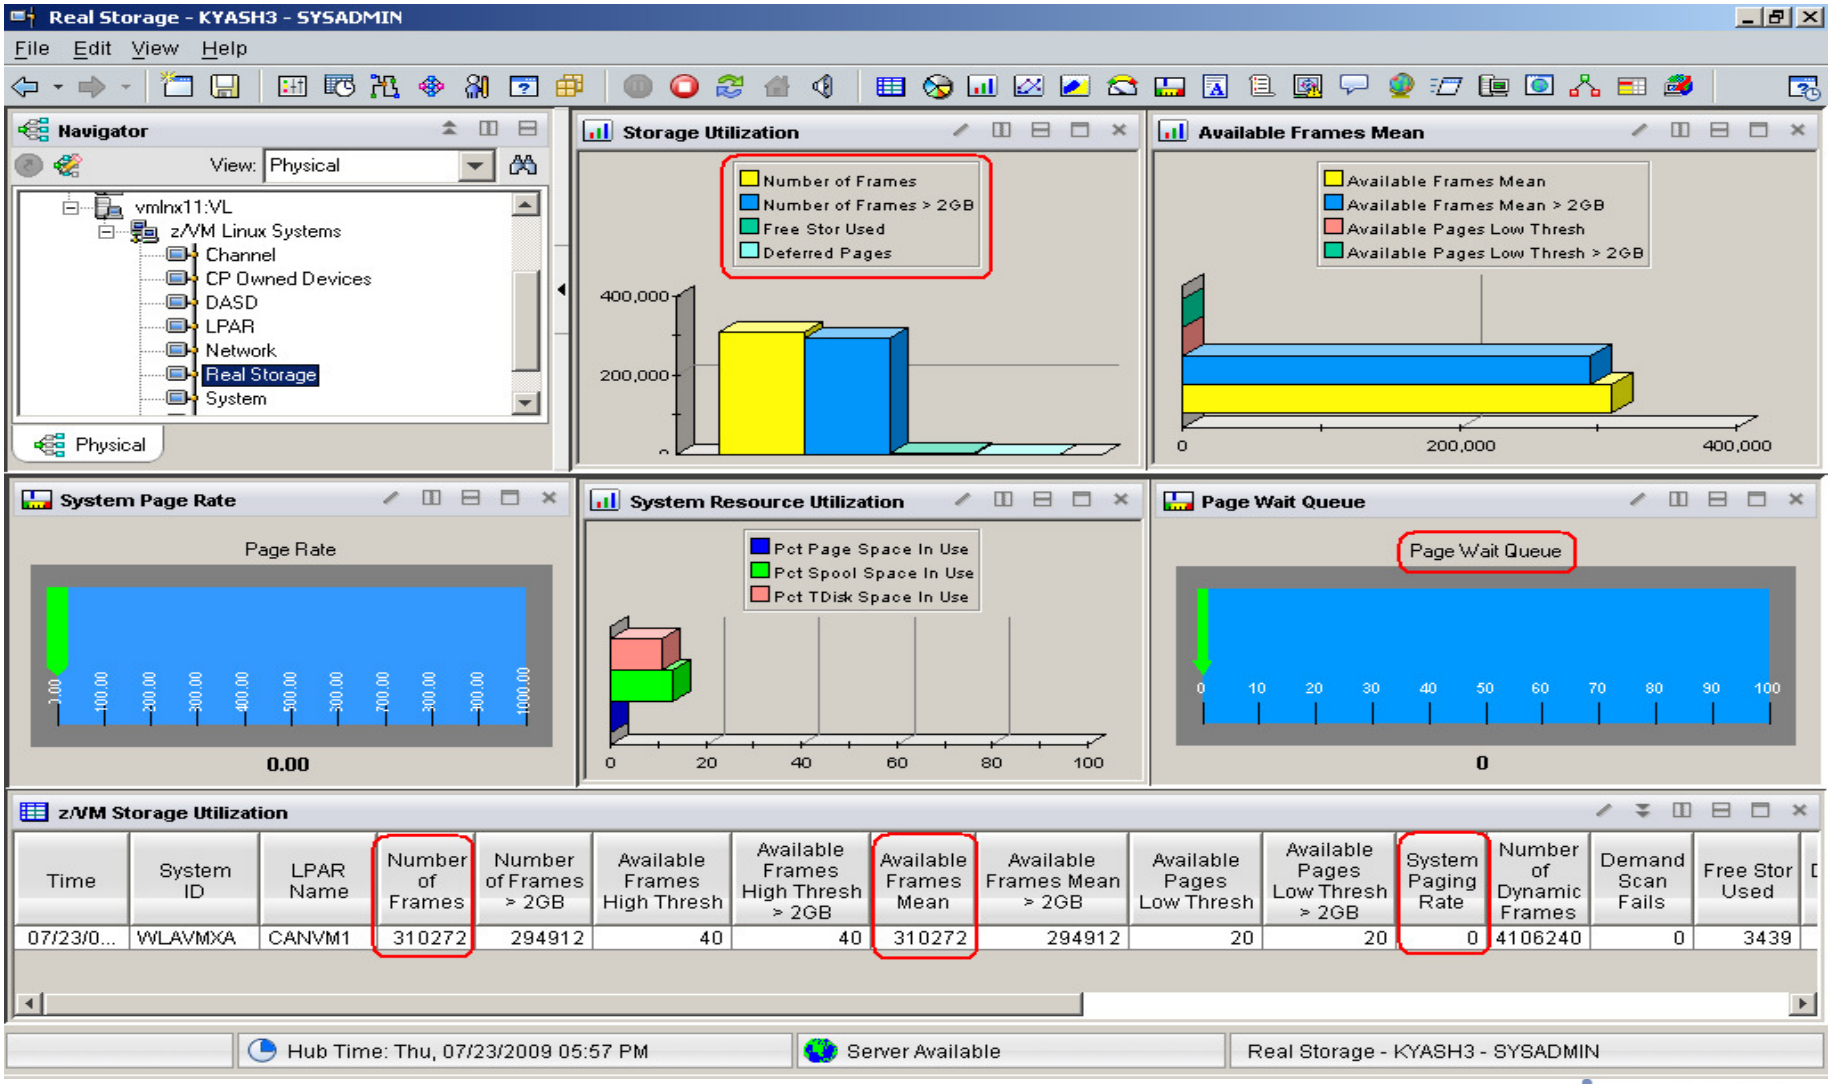

Paging Subsystem

- Plan for DASD page space utilization < 50%:
  - Page space tends to get fragmented over time.
  - Large contiguous free space allows for greater paging efficiency.
  - Monitor usage with OMEGAMON XE or Q ALLOC PAGE command.
- Do not mix page space with any other space on a volume.
- Recommend using devices of the same size/geometry.
- Calculation guidelines are located in the CP Planning and Administration Manual.
- z/VM 6.3 block paging changes eliminates the benefits of 50% page space, making monitoring even more important.

Planning for Large Memory

- **Normal best practices for migrating from an earlier release still apply.**
- **Change your paging XSTORE into central**
  - XSTORE provided an aging function. It helped catch reclaim selection "mistakes".
  - The new ***IBR*** concept and ***global aging list*** provide the same function but do so more efficiently in central.
- **Plan enough DASD paging space**
  - The system now pre-writes pages to DASD.
  - See z/VM manuals for detail recommendations.
- **Plan a robust paging DASD configuration**
  - Use plenty of paging volumes
  - Make the volumes all the same size
  - Put only paging space on the volumes you use for paging
  - Spread the paging volumes through your logical control units
  - Avoid logical control units that you know are hot on application I/O
  - Use plenty of chpids
  - Do not use ESCON chpids
  - Do not mix ECKD paging and SCSI paging
  - Leave reserved slots in the CP-owned list

OMEGAMON CP Owned Devices – Paging Subsystem

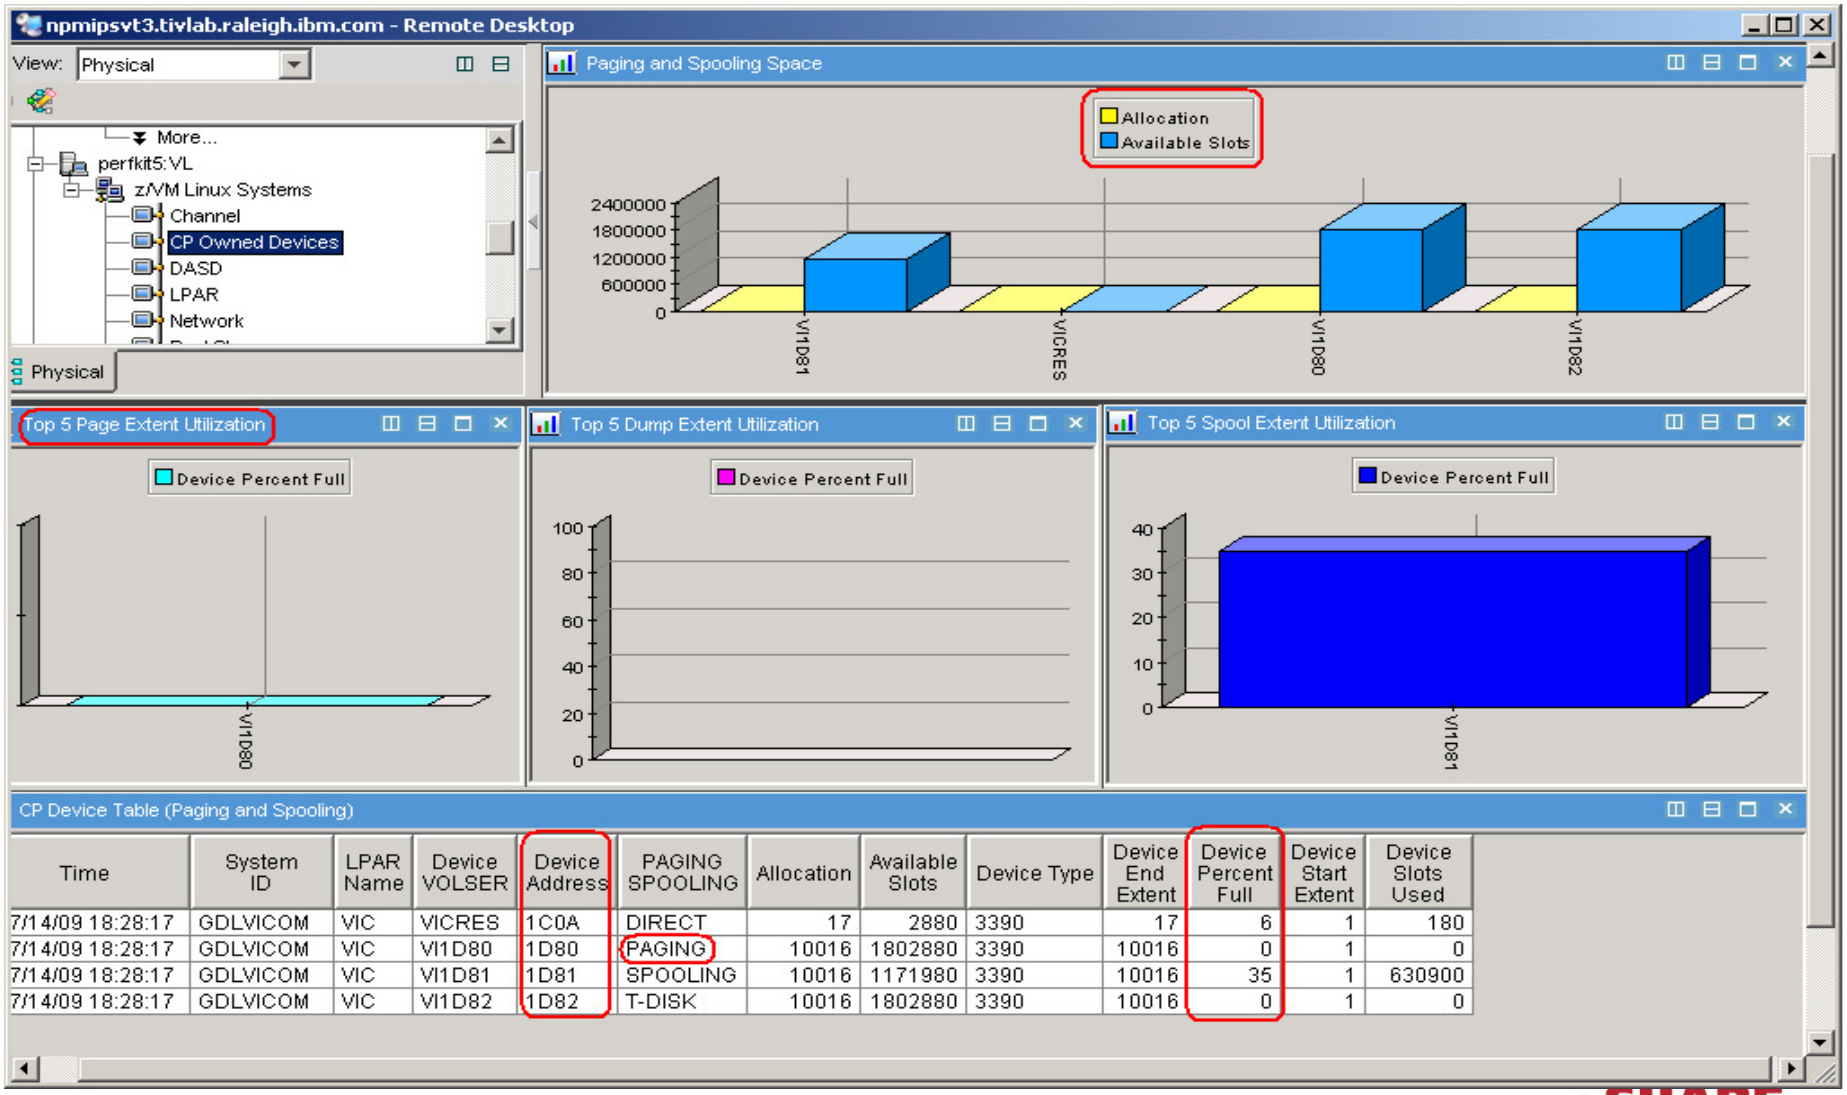

Time	System ID	LPAR Name	Device VOLSER	Device Address	PAGING SPOOLING	Allocation	Available Slots	Device Type	Device End Extent	Device Percent Full	Device Start Extent	Device Slots Used
7/14/09 18:28:17	GDLVICOM	VIC	VICRES	1C0A	DIRECT	17	2880	3390	17	6	1	180
7/14/09 18:28:17	GDLVICOM	VIC	VI1D80	1D80	PAGING	10016	1802880	3390	10016	0	1	0
7/14/09 18:28:17	GDLVICOM	VIC	VI1D81	1D81	SPOOLING	10016	1171980	3390	10016	35	1	630900
7/14/09 18:28:17	GDLVICOM	VIC	VI1D82	1D82	T-DISK	10016	1802880	3390	10016	0	1	0

z/VM Page Attributes

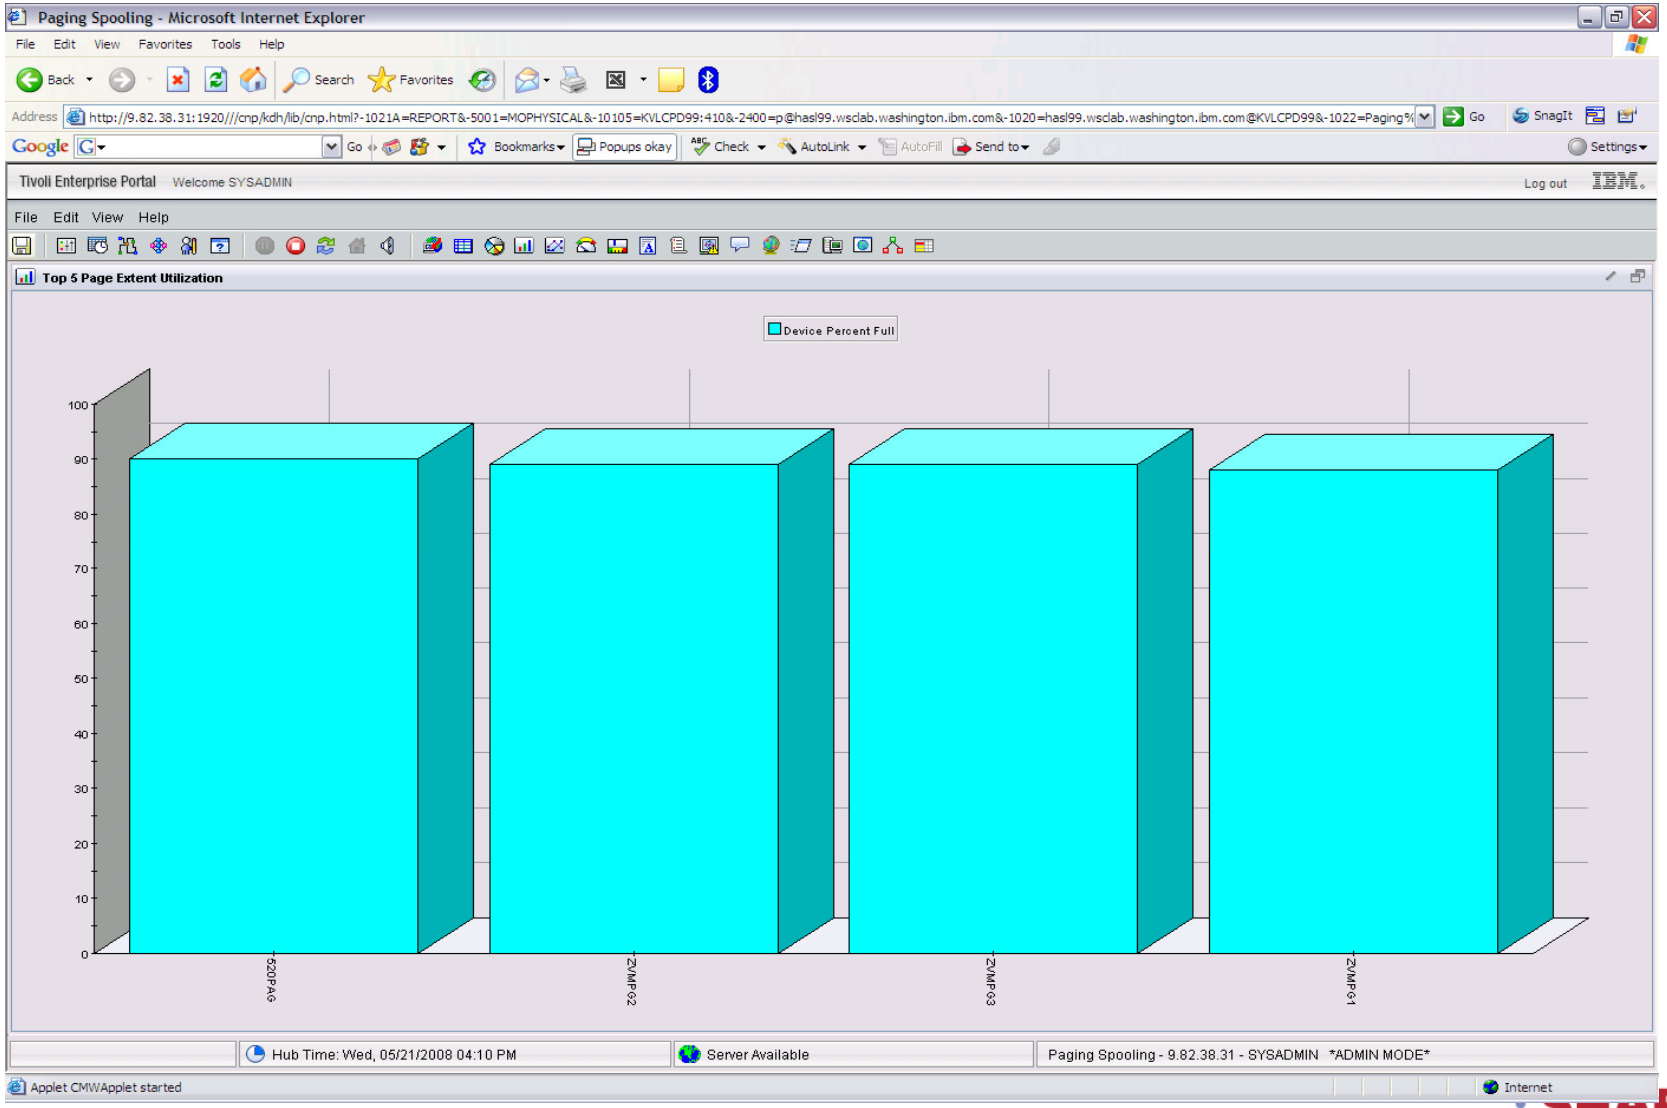

Minidisk Cache

- z/VM minidisk cache is a write-through cache:
  - Improves read I/O performance.
  - But it's not free.
- Not recommended for:
  - Memory constrained systems.
  - Linux swap file disks.
  - Flashcopy targets (see next chart)
- Default system settings are less than optimal.
- Recommended settings:
  - Eliminate MDC in expanded storage.
    - **SET MDC XSTORE 0M 0M**
  - Limit MDC in central storage – 10% is a good starting point.
    - **SET MDC STORE 0M 256M**
  - Monitor with product like OMEGAMON XE and/or the Q MDC command.

MDC and FlashCopy Interaction

- FlashCopy requests require z/VM to flush MDC for the entire minidisk.
- MDC Flush processing is very expensive even when there is no data in MDC to flush
  - System Time becomes very high.
- z/OS DFSMS and other utilities can make extensive use of FlashCopy for functions such as defragmentation
- Mitigations
  - Turn off MDC for minidisks that are FlashCopy targets

OMEGAMON MDISK Cache Allocations

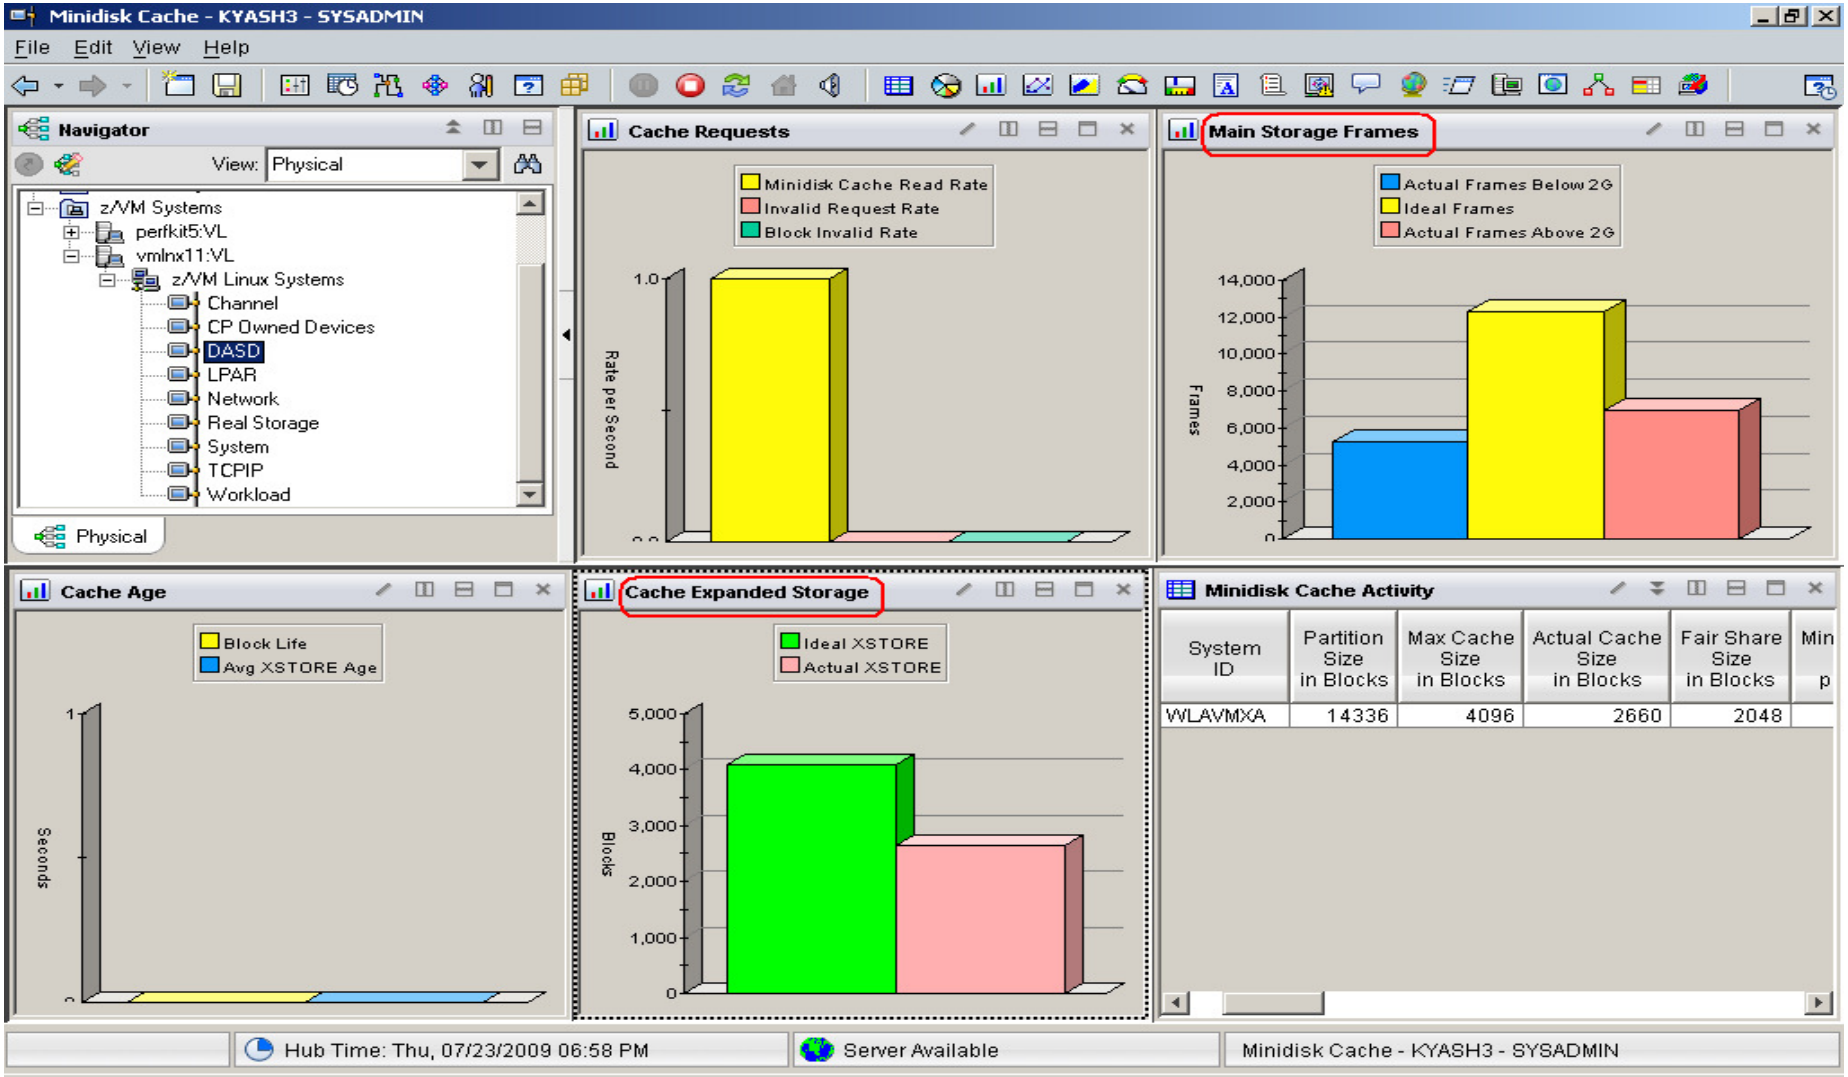

OMEGAMON MDISK Cache Allocations – p. 2

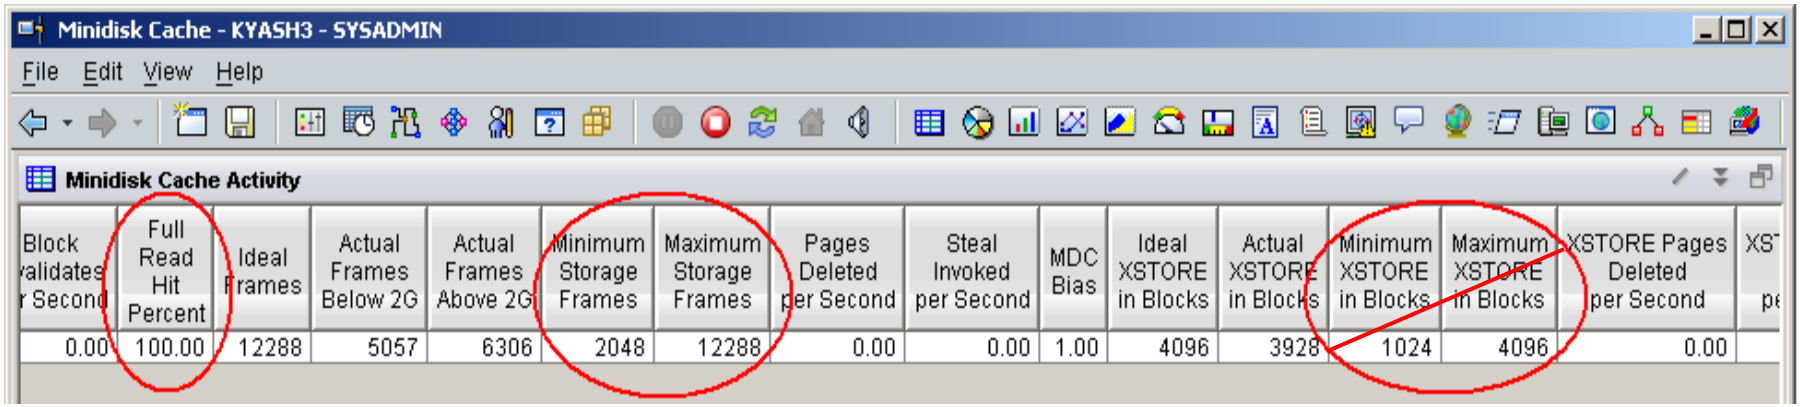

Block validates r Second	Full Read Hit Percent	Ideal Frames	Actual Frames Below 2G	Actual Frames Above 2G	Minimum Storage Frames	Maximum Storage Frames	Pages Deleted per Second	Steal Invoked per Second	MDC Bias	Ideal XSTORE in Blocks	Actual XSTORE in Blocks	Minimum XSTORE in Blocks	Maximum XSTORE in Blocks	XSTORE Pages Deleted per Second	XST pe
0.00	100.00	12288	5057	6306	2048	12288	0.00	0.00	1.00	4096	3928	1024	4096	0.00	

Direct Access Storage Devices (DASD)

- **Avg Pending Time for DASD**
  - Average pending time for real DASD I/Os. RoT: Should be less than 1 millisecond.
- Items worth keeping an eye on:
  - **Number of I/O's per Second, Percent Busy**
  - **Avg Service Time** Average service time for real DASD devices (sum of the pending, connect, and disconnect times).
  - **DASD I/O Rate** Rate of traditional real I/Os per second to real DASD devices. Worth monitoring.

DASD I/O Workspace

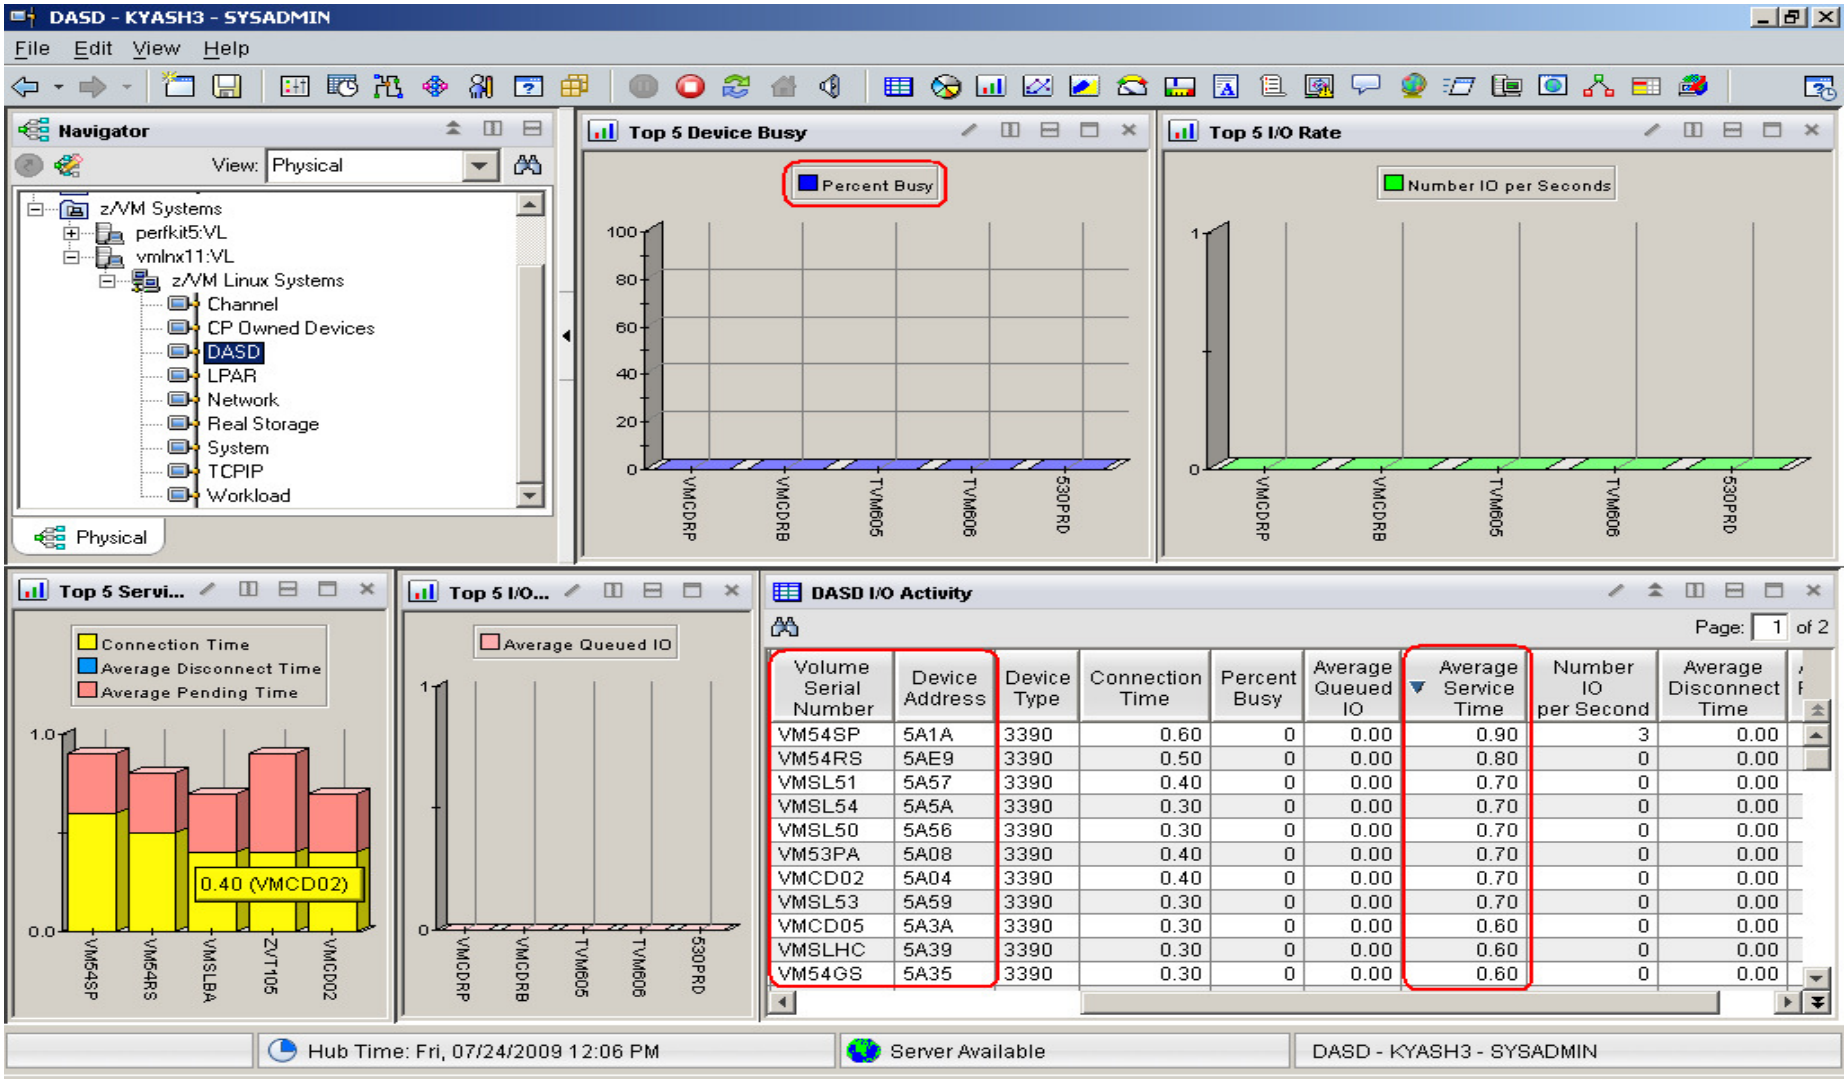

System Dump & Spool Space

- Dump Space
  - Ensure there is sufficient dump space defined to the system.
  - Dump space requirements vary according to memory usage.
    - Q DUMP – identifies allocated dump space.
    - Calculation guidelines are located in CP Planning and Administration Manual.
- Spool Space
  - Various uses:
    - User printer, punch, reader files (console logs)
    - DCSS, NSS
    - System files
    - Page space overflow
  - Spool Management:
    - Monitor with OMEGAMON, Operations Manager, Q ALLOC SPOOL cmd
    - SFPURGER utility:
      - *Rule based tool to clean up spool space.*
      - *Included in the no charge CMS Utilities Feature (CUF).*

VMDUMP Processing Concern

- VMDUMP is a very helpful command for problem determination.
- Some weaknesses:
  - Does not scale well, can take up to 40 minutes per GB.
  - It is not interruptible
    - APAR VM64548 is open to address this.
- Linux provides a disk dump utility which is much faster relative to VMDUMP.
  - It is disruptive
  - Does not include segments outside the normal virtual machine.
- See **http://www.vm.ibm.com/perf/tips/vmdump.html**
- If you increase central, make sure you also increase dump space
  - More guidance available on www.vm.ibm.com/techinfo/
  - Download updated *"Allocating Space for CP Hard ABEND Dumps"*

System Dump & Spool Space

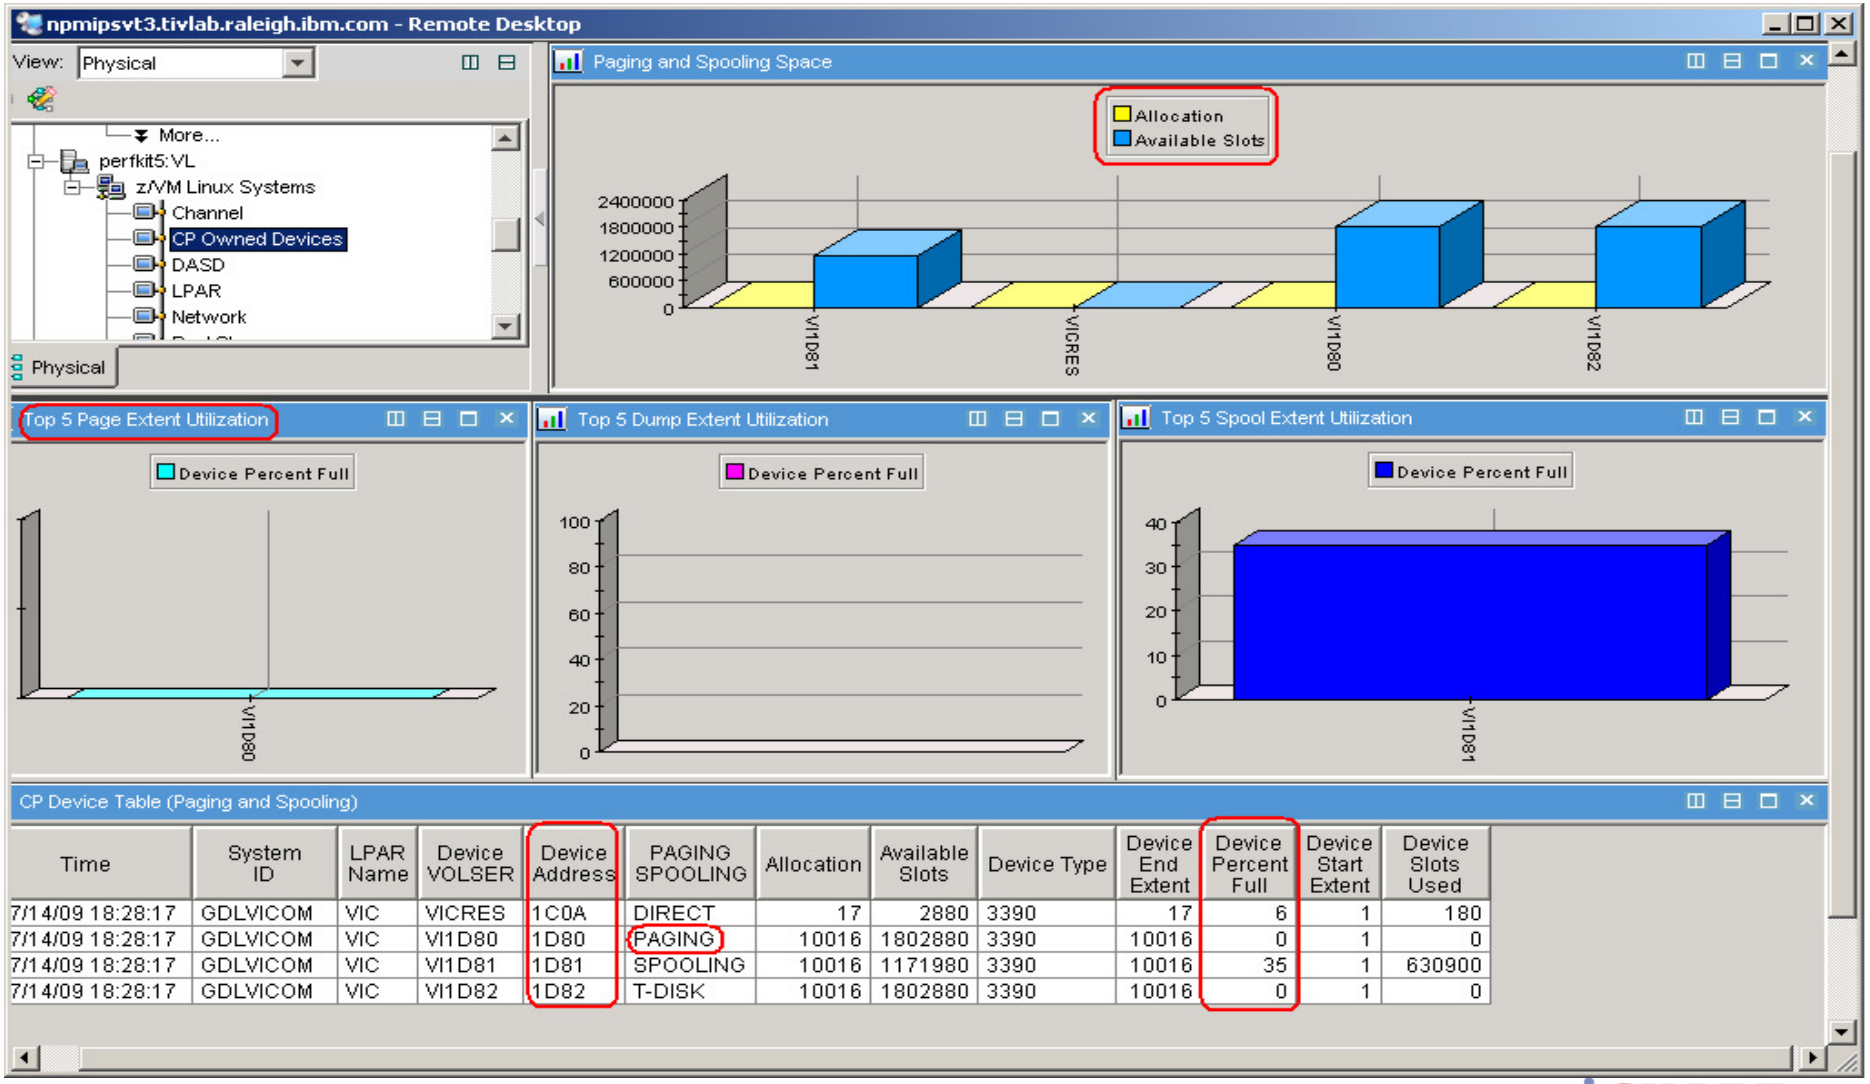

Time	System ID	LPAR Name	Device VOLSER	Device Address	PAGING SPOOLING	Allocation	Available Slots	Device Type	Device End Extent	Device Percent Full	Device Start Extent	Device Slots Used
7/14/09 18:28:17	GDLVICOM	VIC	VICRES	1C0A	DIRECT	17	2880	3390	17	6	1	180
7/14/09 18:28:17	GDLVICOM	VIC	VI1D80	1D80	PAGING	10016	1802880	3390	10016	0	1	0
7/14/09 18:28:17	GDLVICOM	VIC	VI1D81	1D81	SPOOLING	10016	1171980	3390	10016	35	1	630900
7/14/09 18:28:17	GDLVICOM	VIC	VI1D82	1D82	T-DISK	10016	1802880	3390	10016	0	1	0

Tips—Overall Health of Your System

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

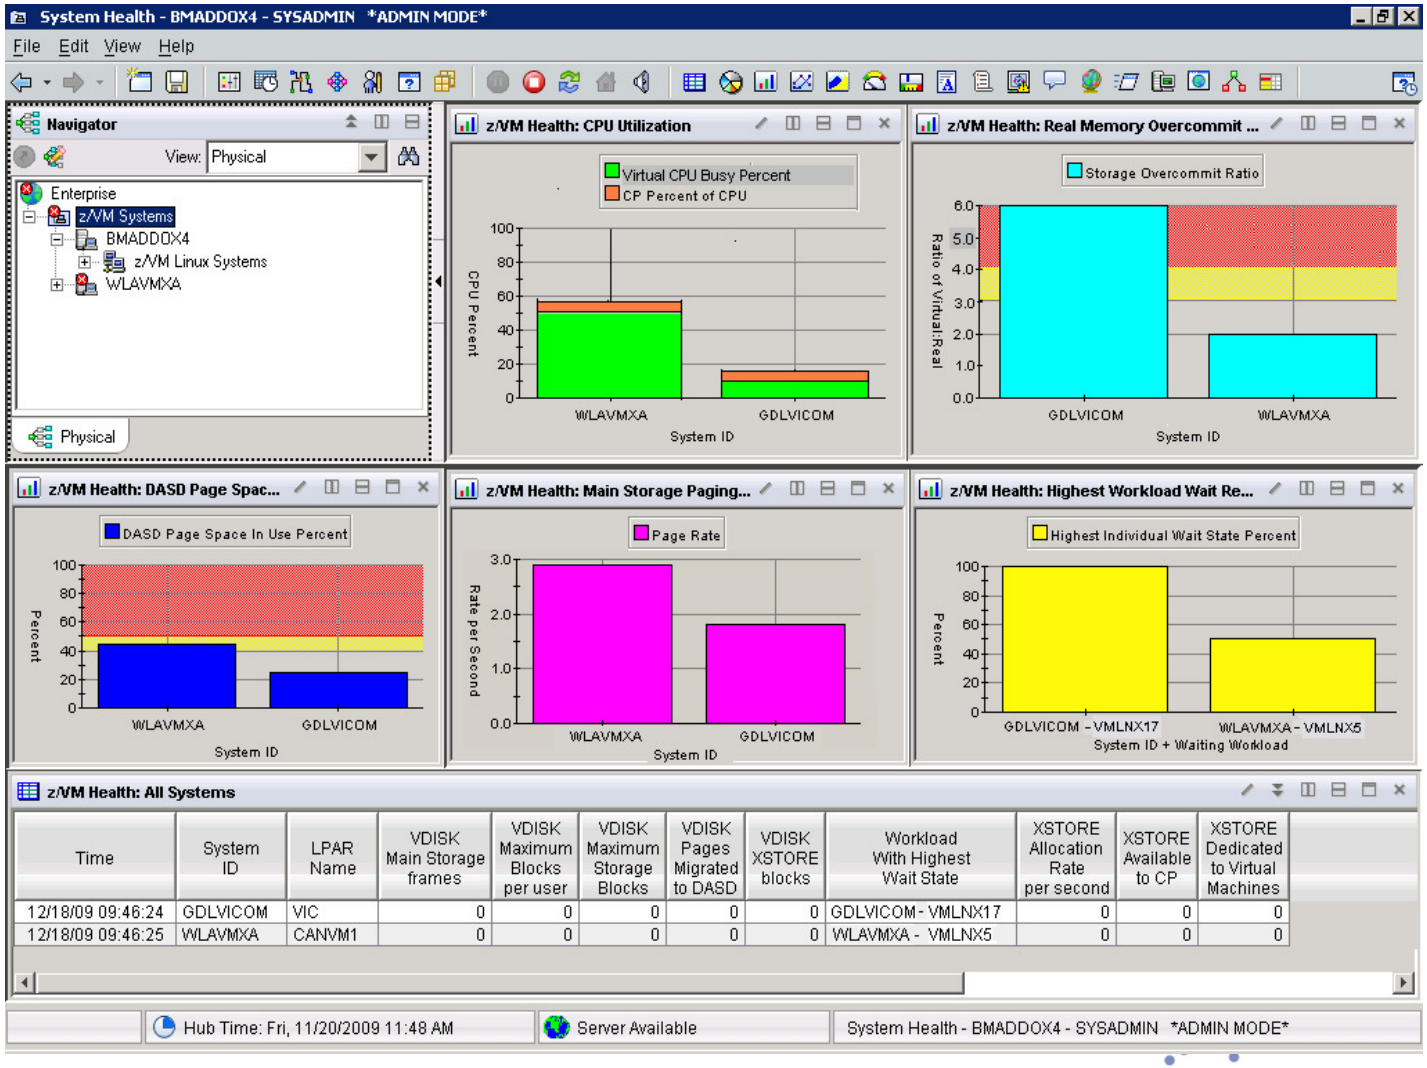

Resource Constraint Analysis (Waits)

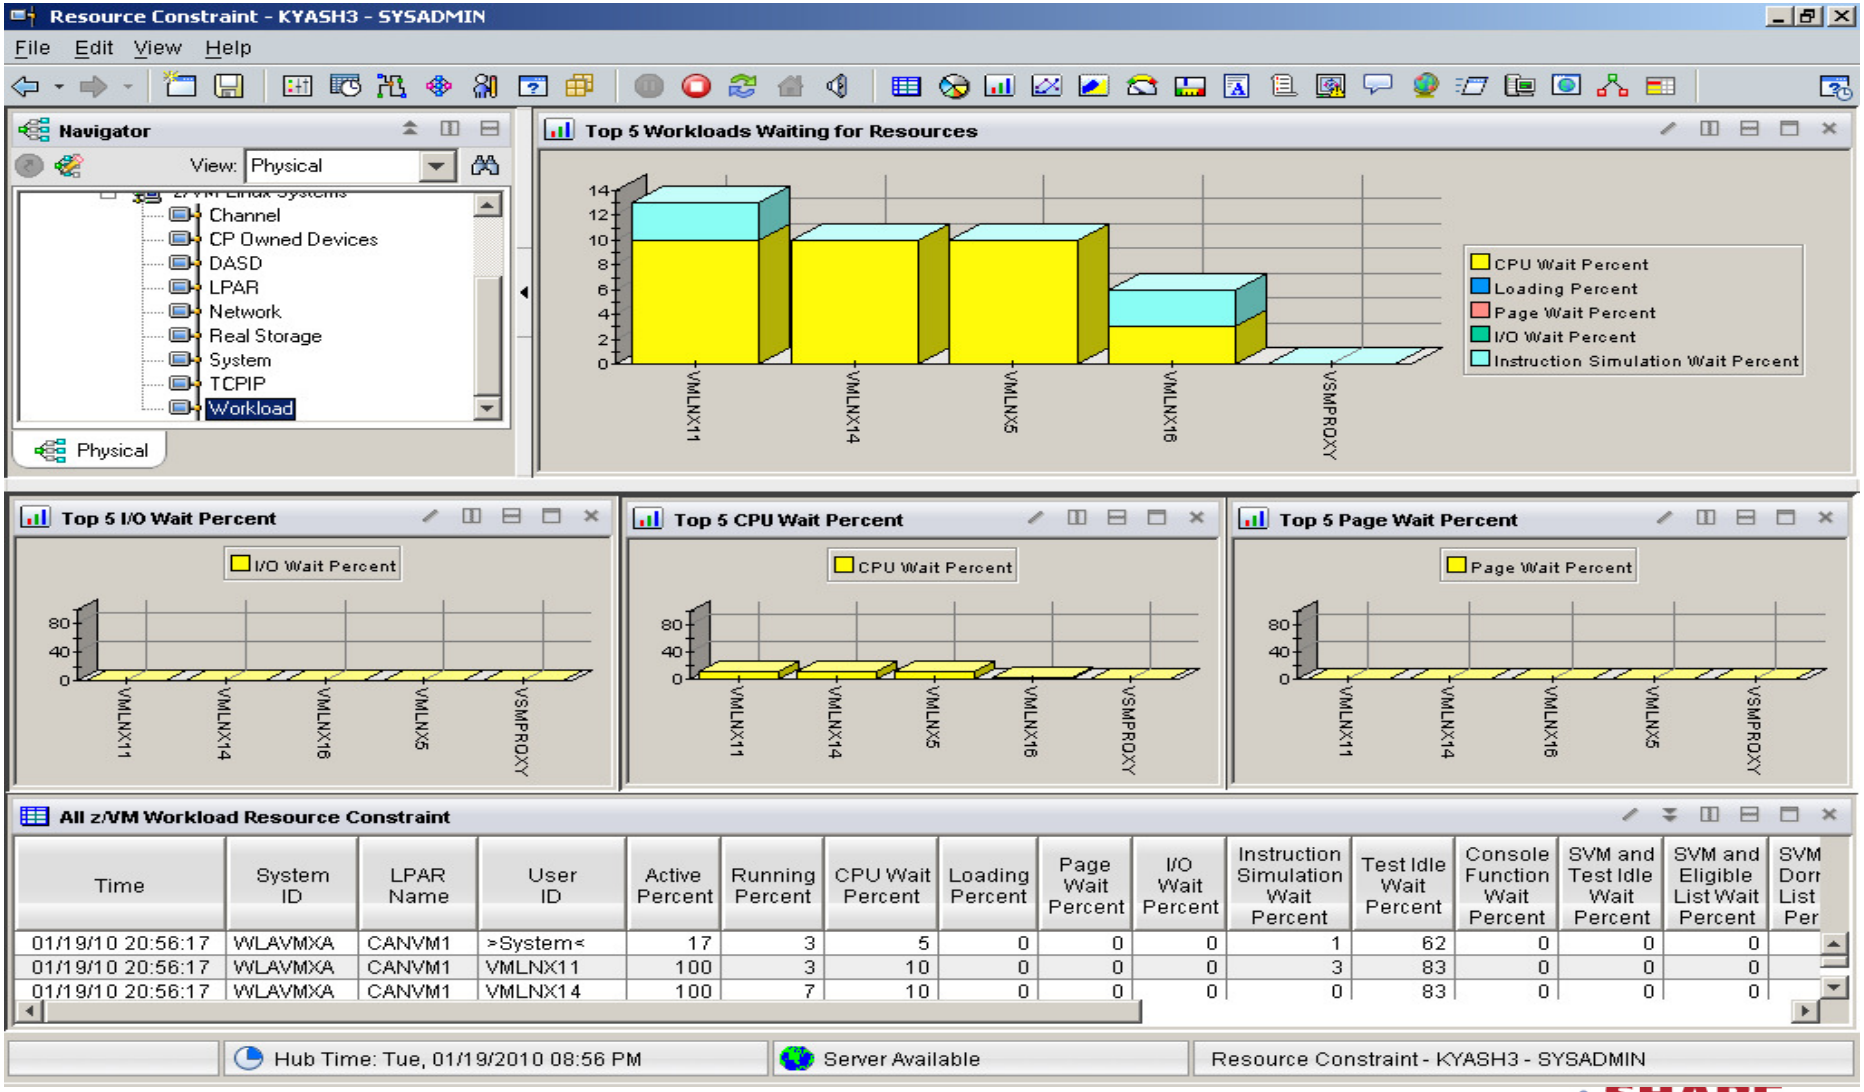

Time	System ID	LPAR Name	User ID	Active Percent	Running Percent	CPU Wait Percent	Loading Percent	Page Wait Percent	I/O Wait Percent	Instruction Simulation Wait Percent	Test Idle Wait Percent	Console Function Wait Percent	SVM and Test Idle Wait Percent	SVM and Eligible List Wait Percent	SVM Dorr List Per
01/19/10 20:56:17	WLAVMXA	CANVM1	>System<	17	3	5	0	0	0	1	62	0	0	0	
01/19/10 20:56:17	WLAVMXA	CANVM1	VMLNX11	100	3	10	0	0	0	3	83	0	0	0	
01/19/10 20:56:17	WLAVMXA	CANVM1	VMLNX14	100	7	10	0	0	0	0	83	0	0	0	

Do not ignore the hardware!

- Just because Linux resources are virtual, do not ignore the hardware!
  - Hardware is another potential layer of shared resources.
  - LPAR weight, CPU sharing, LPAR load, and other attributes need to be monitored for overall system performance.
  - The measurement should include the entire CEC and not just the LPAR hosting z/VM.

Processors

- Logical Processors
  - LPAR recommendation – no greater than a 4:1 logical to real ratio.
  - z/VM 5.1 - z/VM 5.2 support up to 24 processors.
  - z/VM 5.3 - z/VM 6.x support up to 32 processors.

LPAR Utilization Workspace

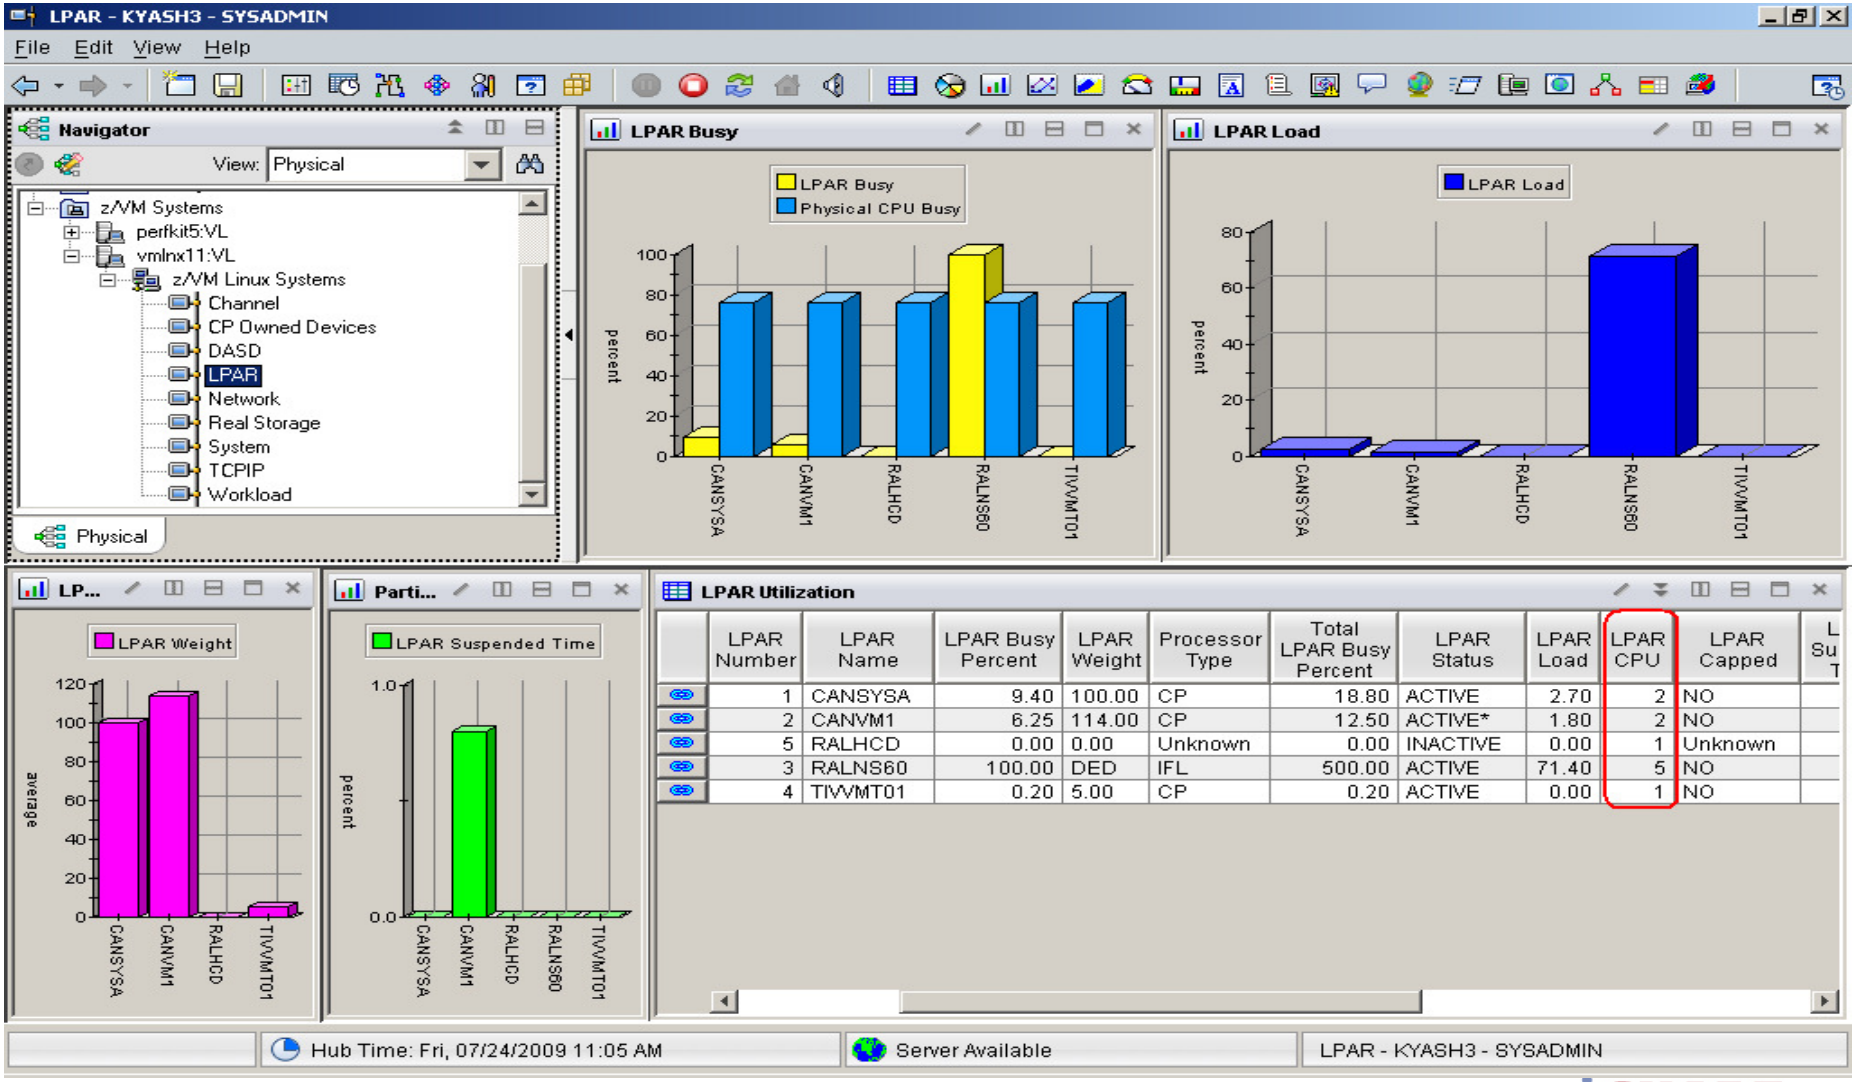

LPAR Number	LPAR Name	LPAR Busy Percent	LPAR Weight	Processor Type	Total LPAR Busy Percent	LPAR Status	LPAR Load	LPAR CPU	LPAR Capped
1	CANSYSA	9.40	100.00	CP	18.80	ACTIVE	2.70	2	NO
2	CANVM1	6.25	114.00	CP	12.50	ACTIVE*	1.80	2	NO
5	RALHCD	0.00	0.00	Unknown	0.00	INACTIVE	0.00	1	Unknown
3	RALNS60	100.00	DED	IFL	500.00	ACTIVE	71.40	5	NO
4	TIVVMT01	0.20	5.00	CP	0.20	ACTIVE	0.00	1	NO

Processor by LPAR name workspace

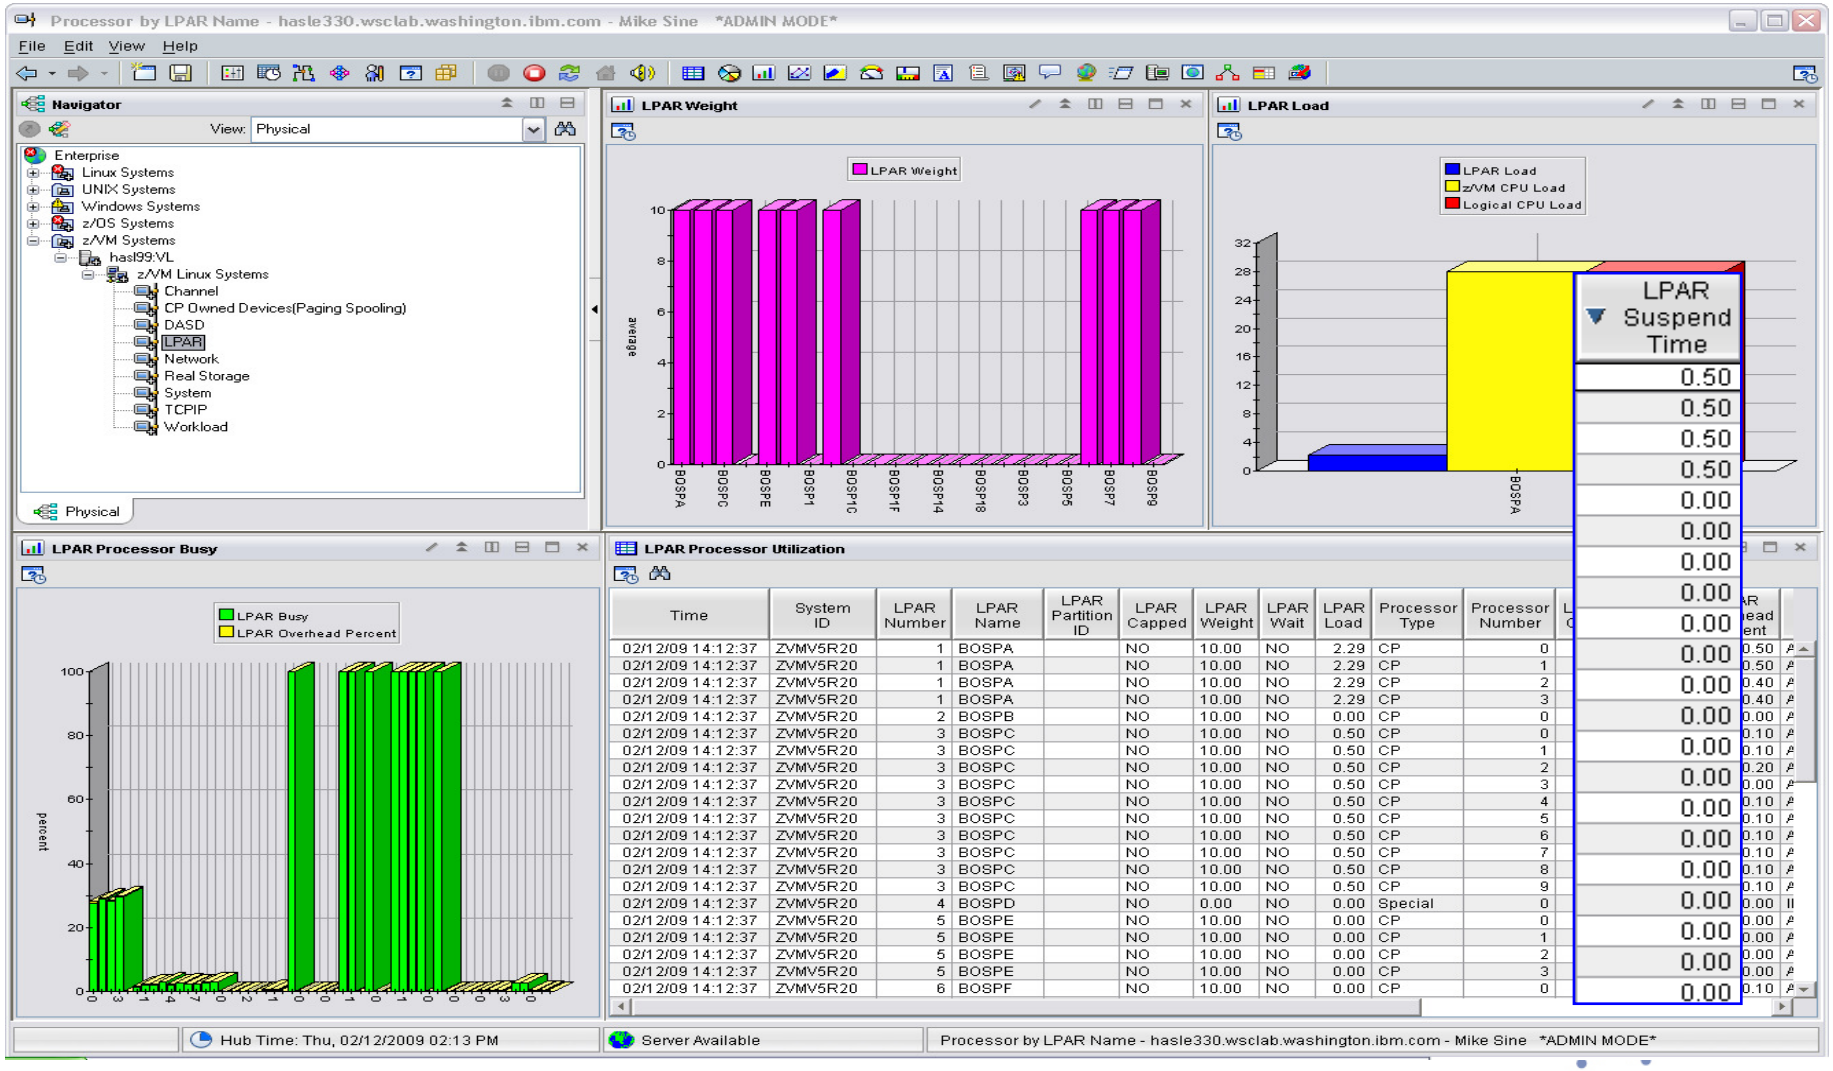

CPC workspace

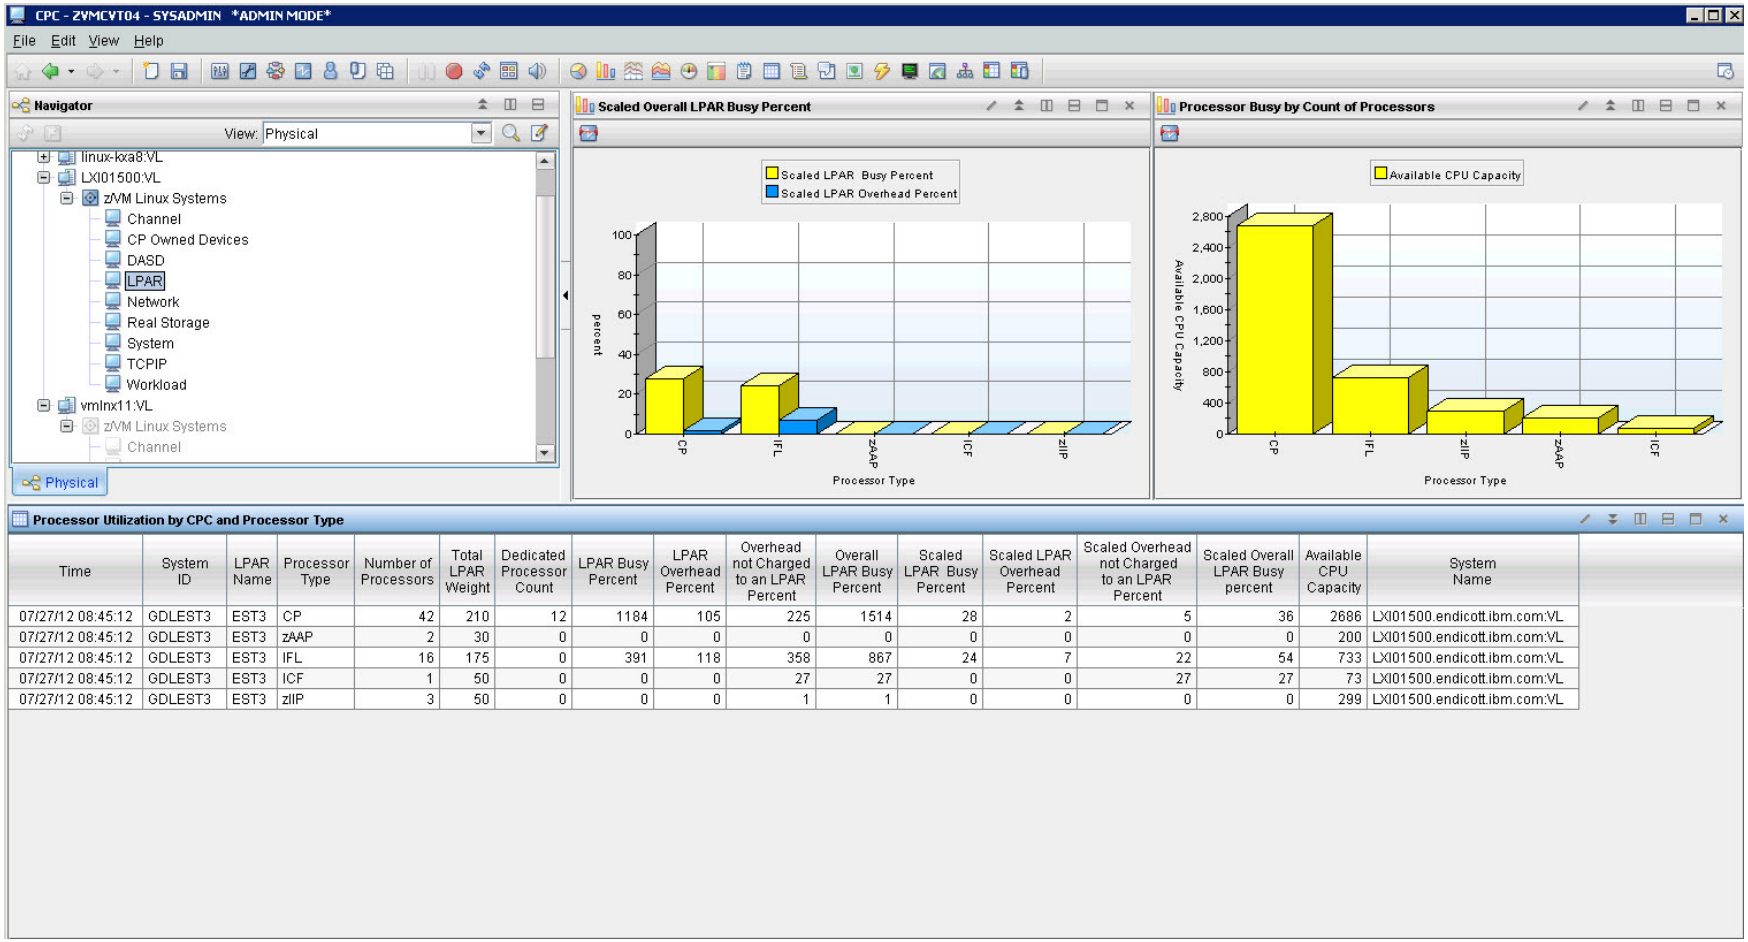

Time	System ID	LPAR Name	Processor Type	Number of Processors	Total LPAR Weight	Dedicated Processor Count	LPAR Busy Percent	LPAR Overhead Percent	Overhead not Charged to an LPAR Percent	Overall LPAR Busy Percent	Scaled LPAR Busy Percent	Scaled LPAR Overhead Percent	Scaled Overhead not Charged to an LPAR Percent	Scaled Overall LPAR Busy percent	Available CPU Capacity	System Name
07/27/12 08:45:12	GDLEST3	EST3	CP	42	210	12	1184	105	225	1514	28	2	5	36	2686	LXI01500.endicott.ibm.com:VL
07/27/12 08:45:12	GDLEST3	EST3	zAAP	2	30	0	0	0	0	0	0	0	0	0	200	LXI01500.endicott.ibm.com:VL
07/27/12 08:45:12	GDLEST3	EST3	IFL	16	175	0	391	118	358	867	24	7	22	54	733	LXI01500.endicott.ibm.com:VL
07/27/12 08:45:12	GDLEST3	EST3	ICF	1	50	0	0	0	27	27	0	0	27	27	73	LXI01500.endicott.ibm.com:VL
07/27/12 08:45:12	GDLEST3	EST3	zIIP	3	50	0	0	0	1	1	0	0	0	0	299	LXI01500.endicott.ibm.com:VL

History On-Demand with Persistent Historical Views

This makes it easier to see anomalies, or match spikes. Capturing performance data as a base line is a must:

- General history data – business as usual.
- Detailed raw monitor data prior to and following any major changes.
- Ability to review attributes of a past incident through the enterprise view!
- On-Demand through the Portal or Batch

On-Demand: Persistent Historical Views

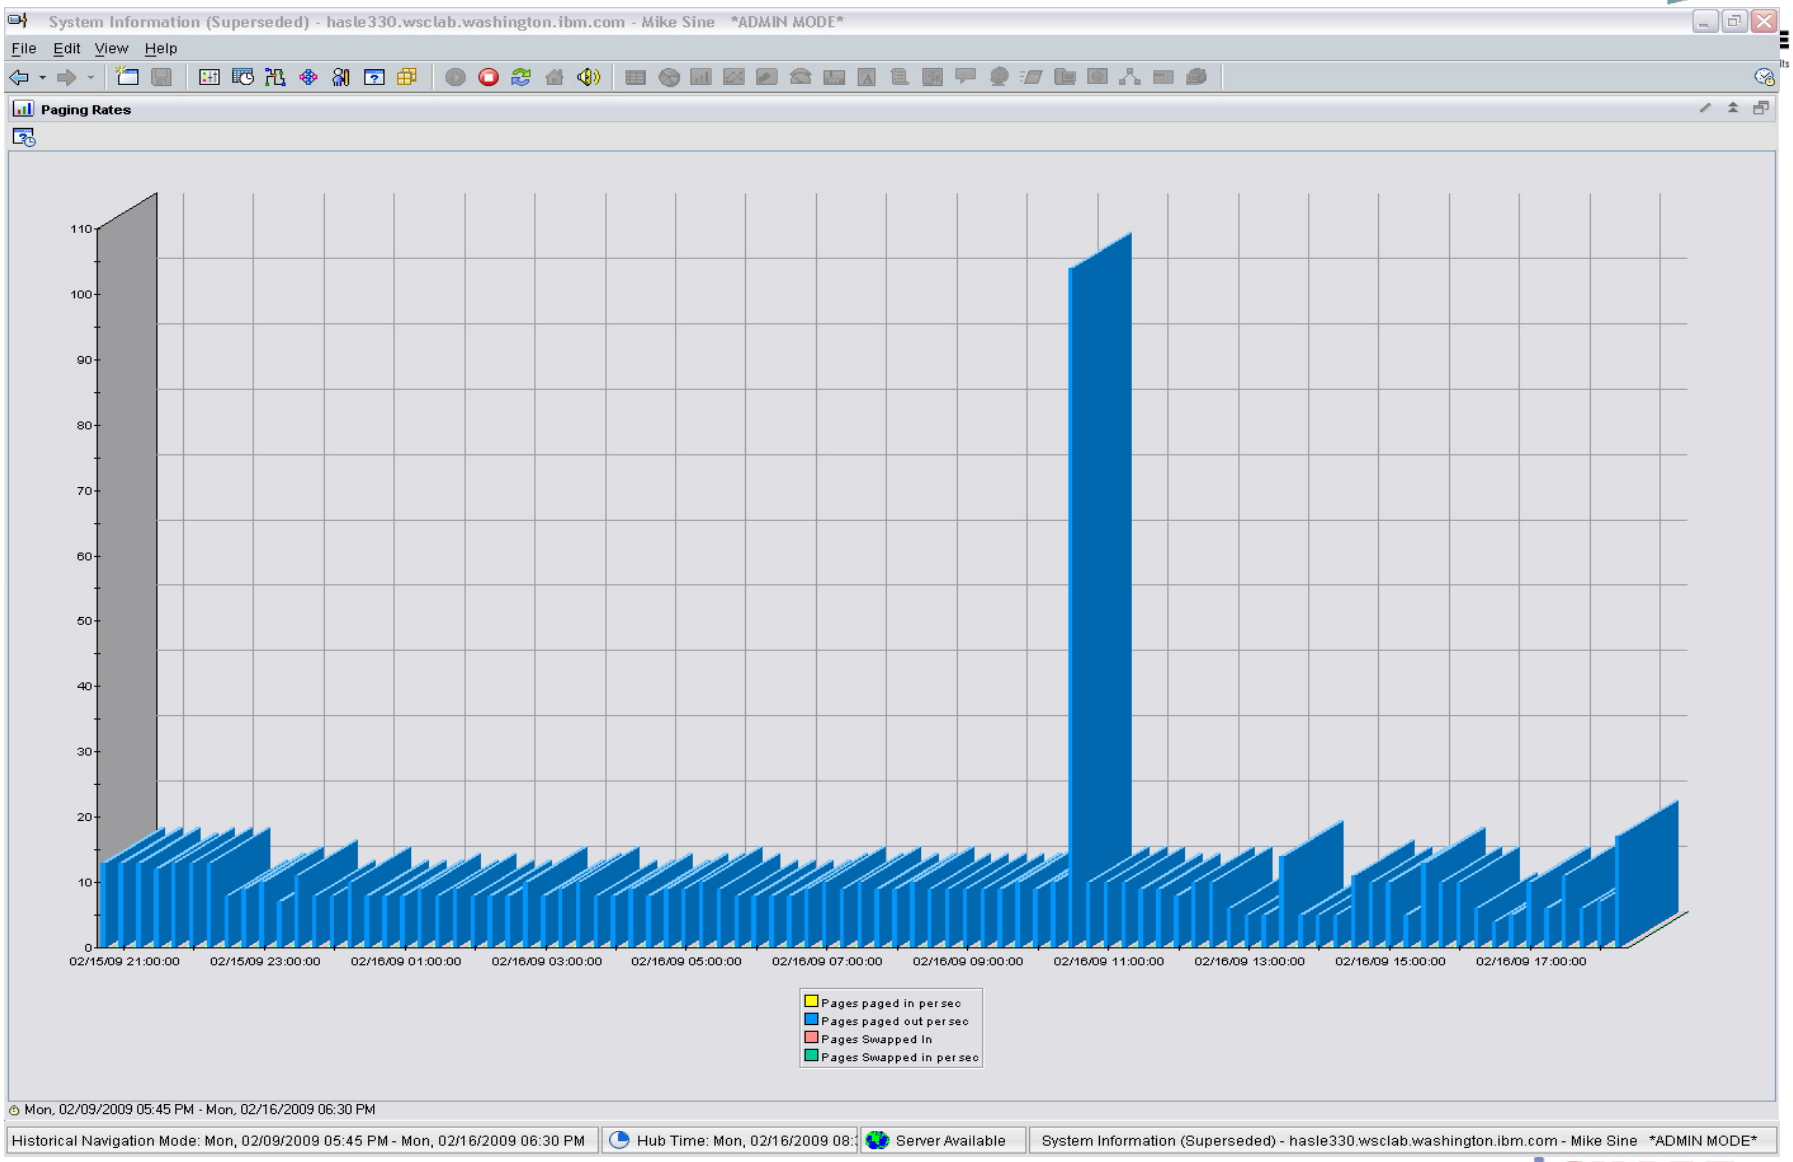

On-Demand: Persistent Historical Views

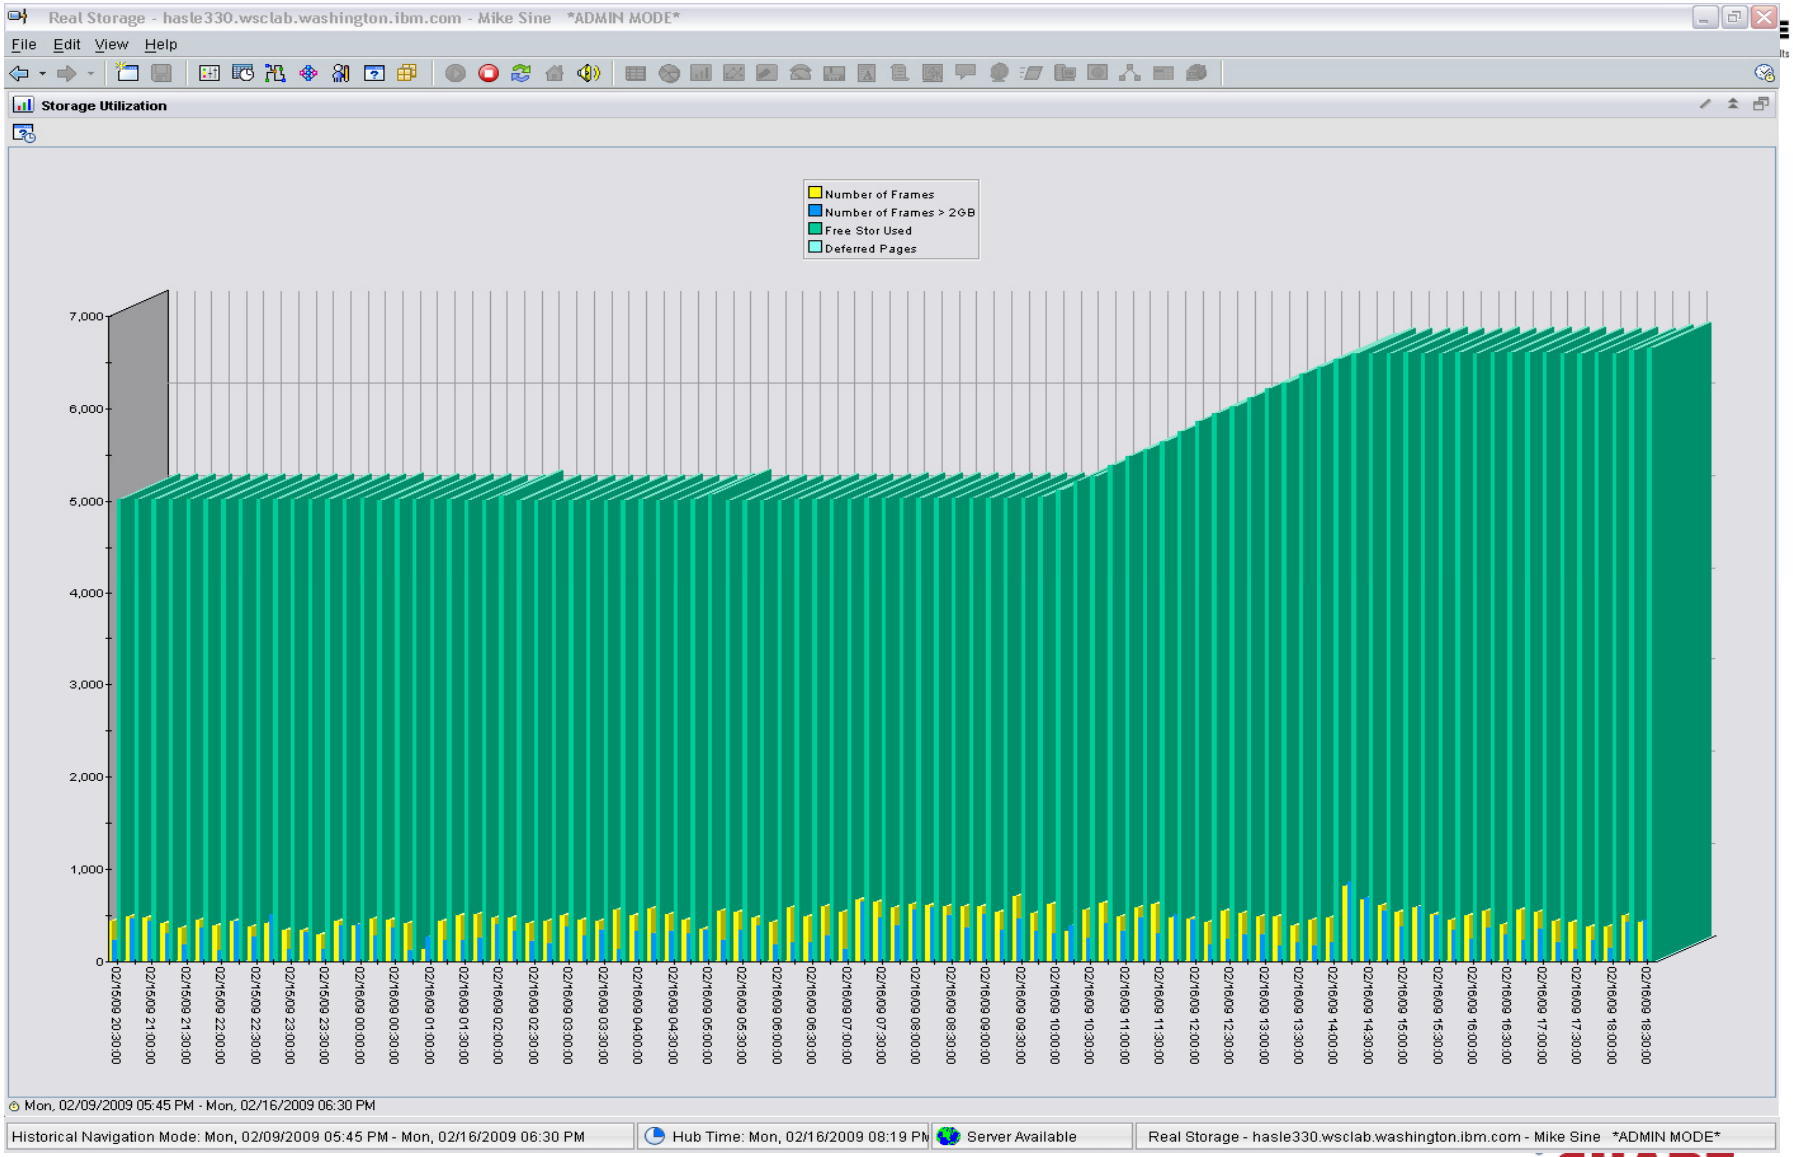

Max and Avg CPU example:

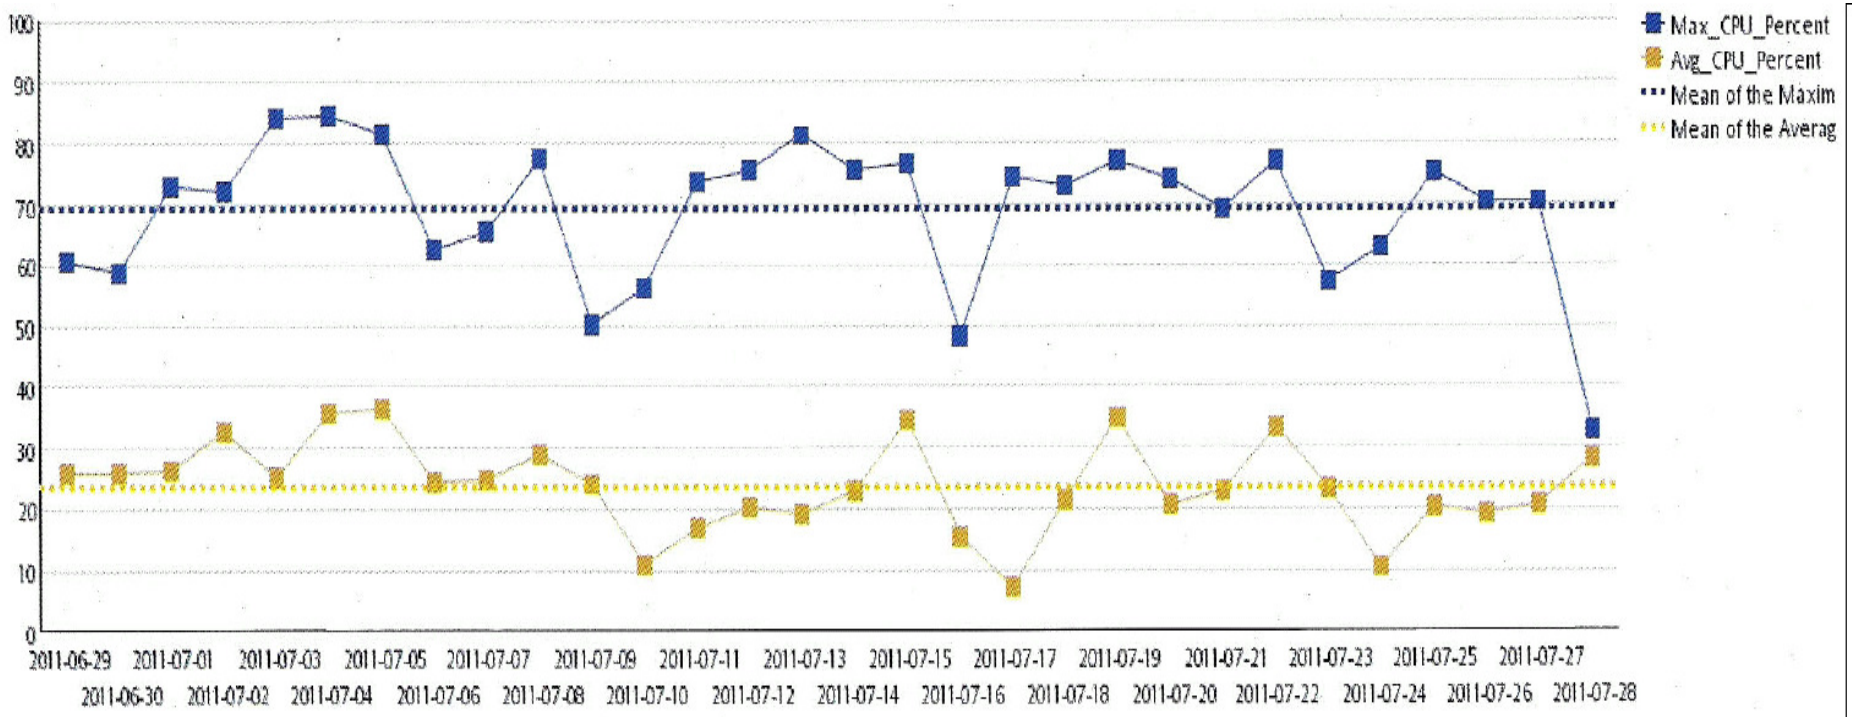

Legend:

Max_CPU_Percent:	Maximum CPU for the day as a percent of the number of virtual CPUs
Avg_CPU_Percent:	Average CPU for the day as a percent of virtual CPUs
Mean of the Maximum:	30 day average for Maximum CPU percentages
Mean of the Averages:	30 day average for the average CPU percentages
AVG_Main_Memory_Util:	Average main memory utilization for the day as a percent
AVG_Cache_Used:	Average size of memory used to cache buffers in megabytes
AVG_Page_Alloc_Rate:	Average number of pages obtained from available list in 4 kilobyte pages per second
AVG_Swap_Used:	The percent of swap space used.

Avg Linux Memory breakdown example:

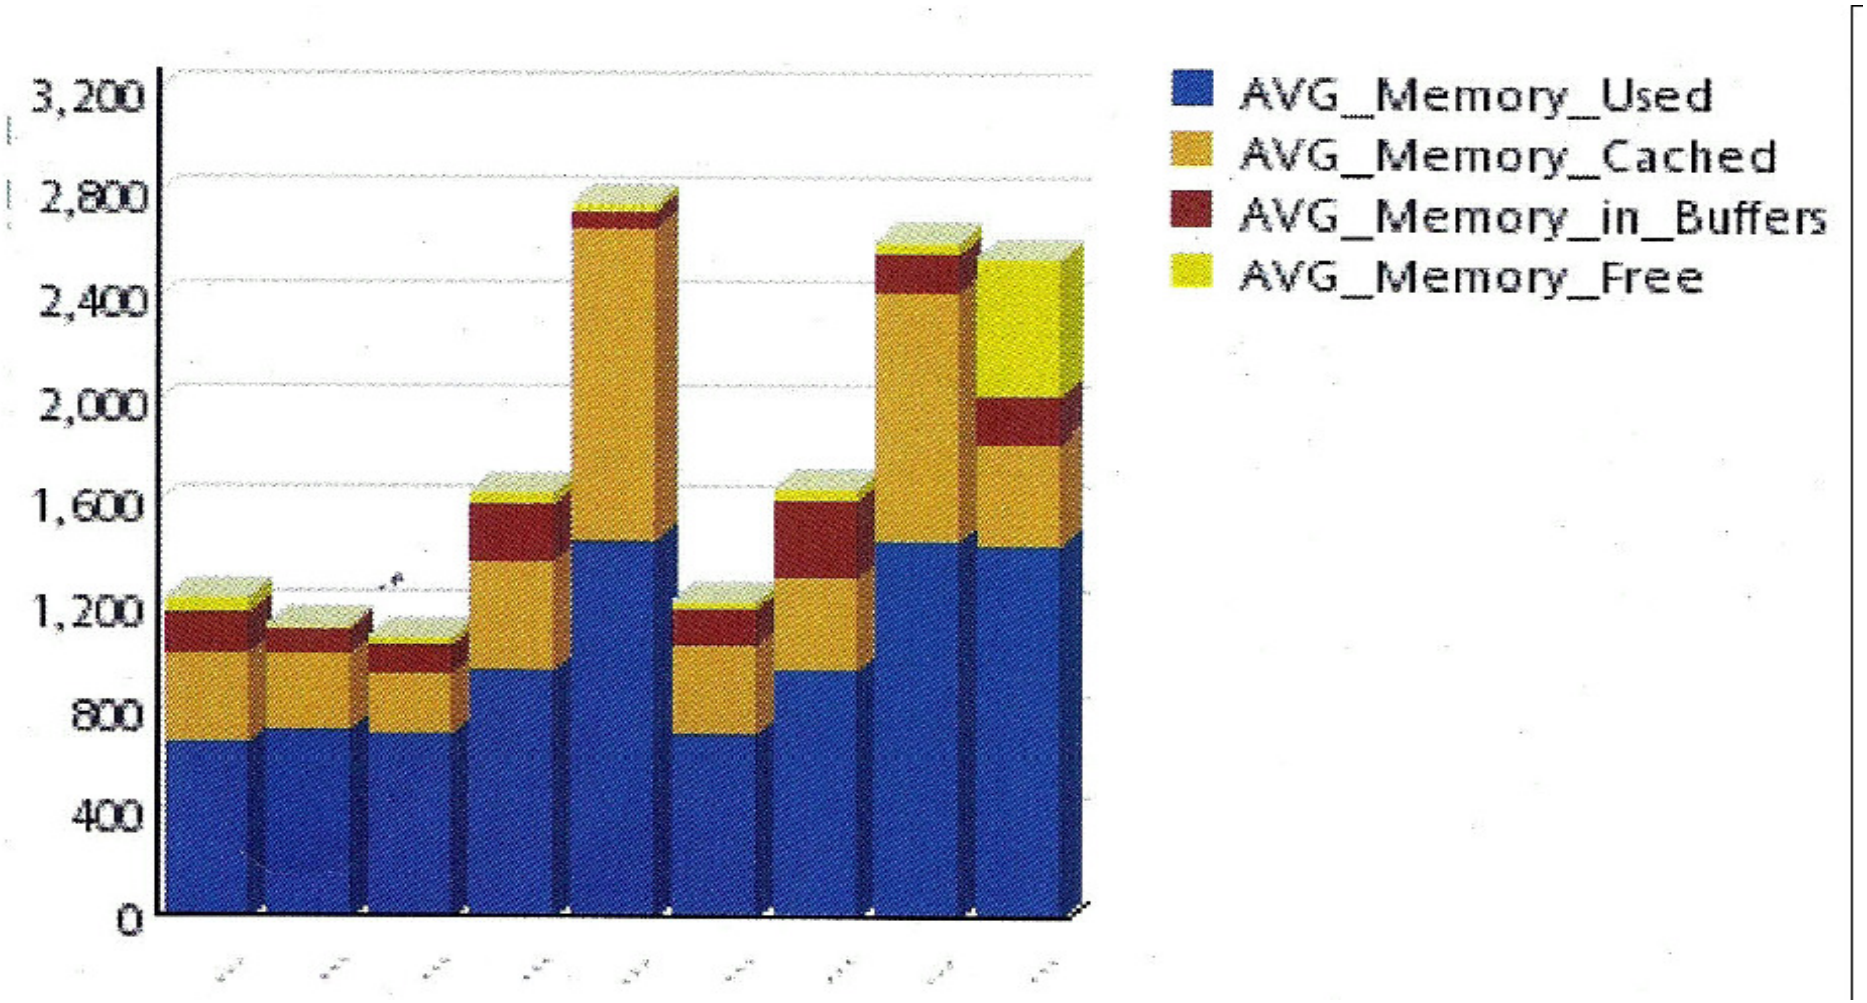

Tivoli Common Reporting (TCR)

- TCR reports available on the OPAL website
  - http://www-18.lotus.com/wps/portal/topal
- What is TCR?
  - Tivoli Common Reporting.
  - Consistent approach to viewing and administering reports.
  - Cognos based.
  - Flexible development environment (Eclipse based) for creating report definitions.
  - Five templates provided for download.
  - Taking suggestions for more

Sample Reports Available

- z/VM VM System CPU Utilization
- z/VM VM System Paging Utilization
- z/VM Linux System CPU Utilization
- z/VM VM System CP-Owned Device Utilization
- z/VM VM System TCP Server Statistics

Tivoli

IBM

SHARE
Technology · Connections · Results

z/VM System CPU Utilization

Report Period	All	**Significant Resources Selected**	5
Start Date	Dec 31, 1969 12:00 AM	**End Date**	Nov 30, 2007 11:59 PM
System ID	All	**LPAR Name**	All

LPAR Busy

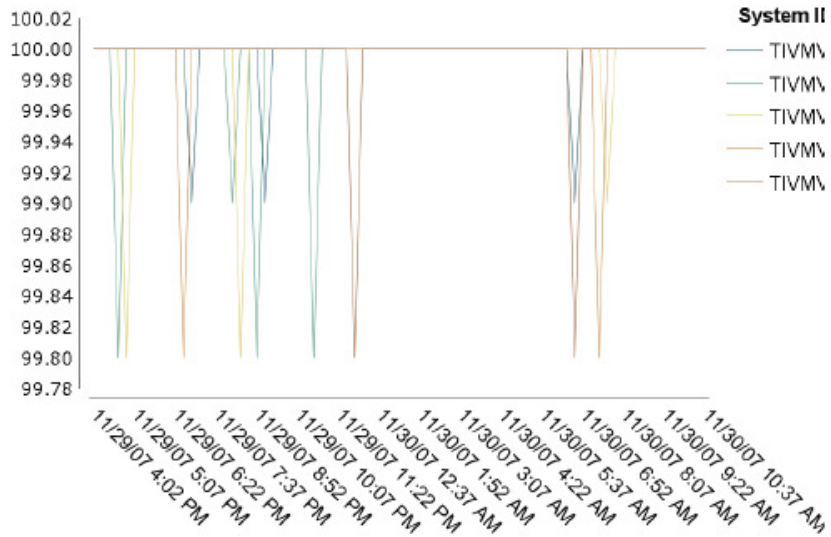

Available Summarization Time Periods:
Hourly
Daily
Weekly
Monthly
Not Summarized Data

System = TIVMVS6					
LPAR Name	**LPAR Busy**	**LPAR Load**	**LPAR Suspend Time**	**LPAR Overhead Time**	**Date/Time**

System = TIVMVS6

LPAR Name	LPAR Busy	LPAR Load	LPAR Suspend Time	LPAR Overhead Time	Date/Time
RALNS31	100	4.2	0	.6	Nov 29, 2007 4:02 PM
RALNS32	100	4.2	0	.6	Nov 29, 2007 4:02 PM
RALNS61	100	4.2	0	.6	Nov 29, 2007 4:02 PM
TIVMVS1	100	2.09	0	.6	Nov 29, 2007 4:02 PM
TIVMVS10	100	2.09	0	.6	Nov 29, 2007 4:02 PM
RALNS31	100	4.2	0	.6	Nov 29, 2007 4:08 PM
RALNS32	100	4.2	0	.6	Nov 29, 2007 4:08 PM
RALNS61	100	4.2	0	.6	Nov 29, 2007 4:08 PM
TIVMVS1	100	2.09	0	.6	Nov 29, 2007 4:08 PM
TIVMVS10	100	2.09	0	.6	Nov 29, 2007 4:08 PM
RALNS31	100	4.2	0	.6	Nov 29, 2007 4:22 PM
RALNS32	100	4.2	0	.6	Nov 29, 2007 4:22 PM
RALNS61	100	4.2	0	.6	Nov 29, 2007 4:22 PM
TIVMVS1	100	2.09	0	.6	Nov 29, 2007 4:22 PM
TIVMVS10	100	2.09	0	.6	Nov 29, 2007 4:22 PM
RALNS31	100	4.2	0	.6	Nov 29, 2007 4:37 PM
RALNS61	100	4.2	0	.6	Nov 29, 2007 4:37 PM
TIVMVS1	100	2.09	0	.6	Nov 29, 2007 4:37 PM
TIVMVS10	100	2.09	0	.6	Nov 29, 2007 4:37 PM
RALNS32	99.8	4.2	0	.6	Nov 29, 2007 4:37 PM
RALNS31	100	4.2	0	.6	Nov 29, 2007 4:52 PM
RALNS32	100	4.2	0	.6	Nov 29, 2007 4:52 PM
TIVMVS1	100	2.09	0	.6	Nov 29, 2007 4:52 PM
TIVMVS10	100	2.09	0	.6	Nov 29, 2007 4:52

AGENDA

- Introduction
- Monitoring Requirements
  - Virtual Linux and z/VM performance considerations
  - Don’t forget the hardware
  - Integration from hardware – systems – applications Persistent historical views
- Enterprise Management
- Operational Requirements
  - Centralized Control
  - Including all Enterprise Virtual Machines
- Integrating Monitoring and Operations
- Bringing it all together

Addressing the trend – zEnterprise will enable management of diverse resources across diverse platforms as a single Workload

- **A Platform Workload is a grouping mechanism and "management view" of virtual servers supporting a business application**
- **Provides the context within which associated platform resources are presented, monitored, reported, and managed**
- **Management policies are associated to Platform Workload**
  - Currently supports Performance Policy

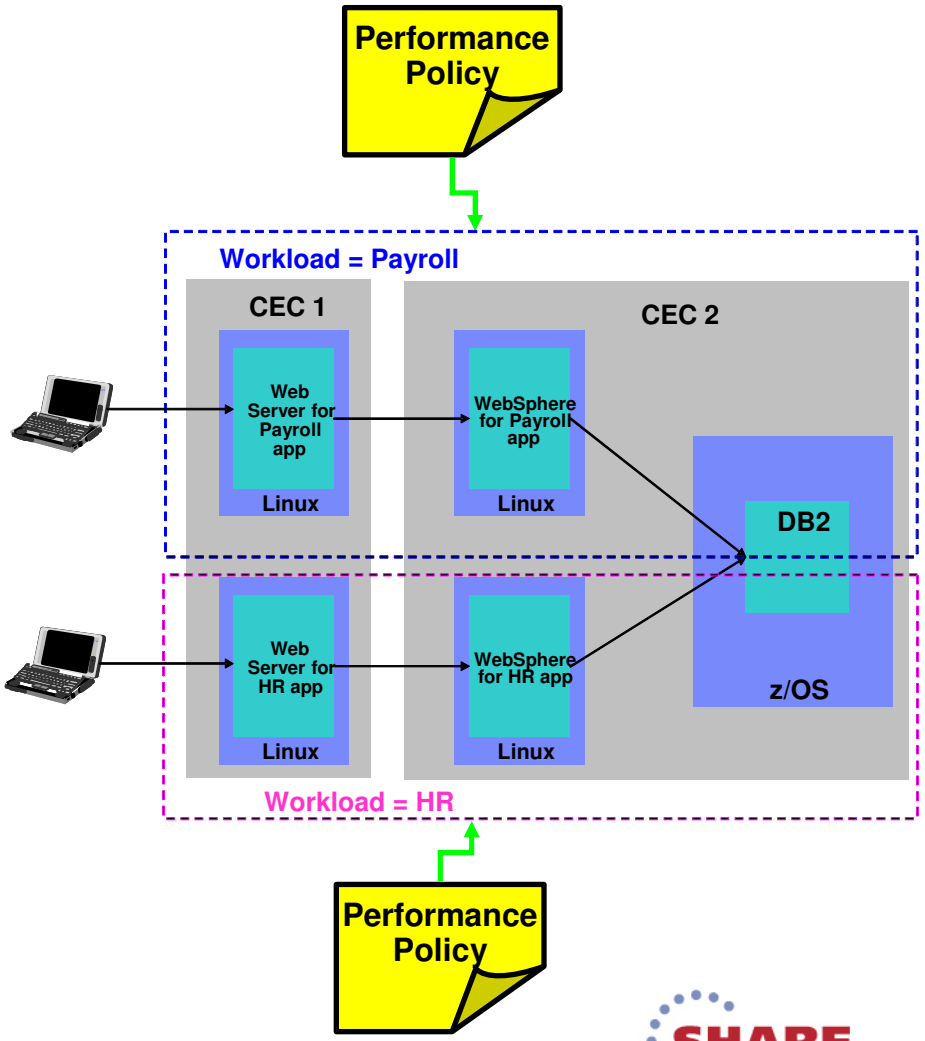

Looking at managing the zEnterprise aka “systems of systems“ getting yourself organized..

Visibility

See your Business

Control

Manage service risk and compliance

Automation

Optimize business service delivery

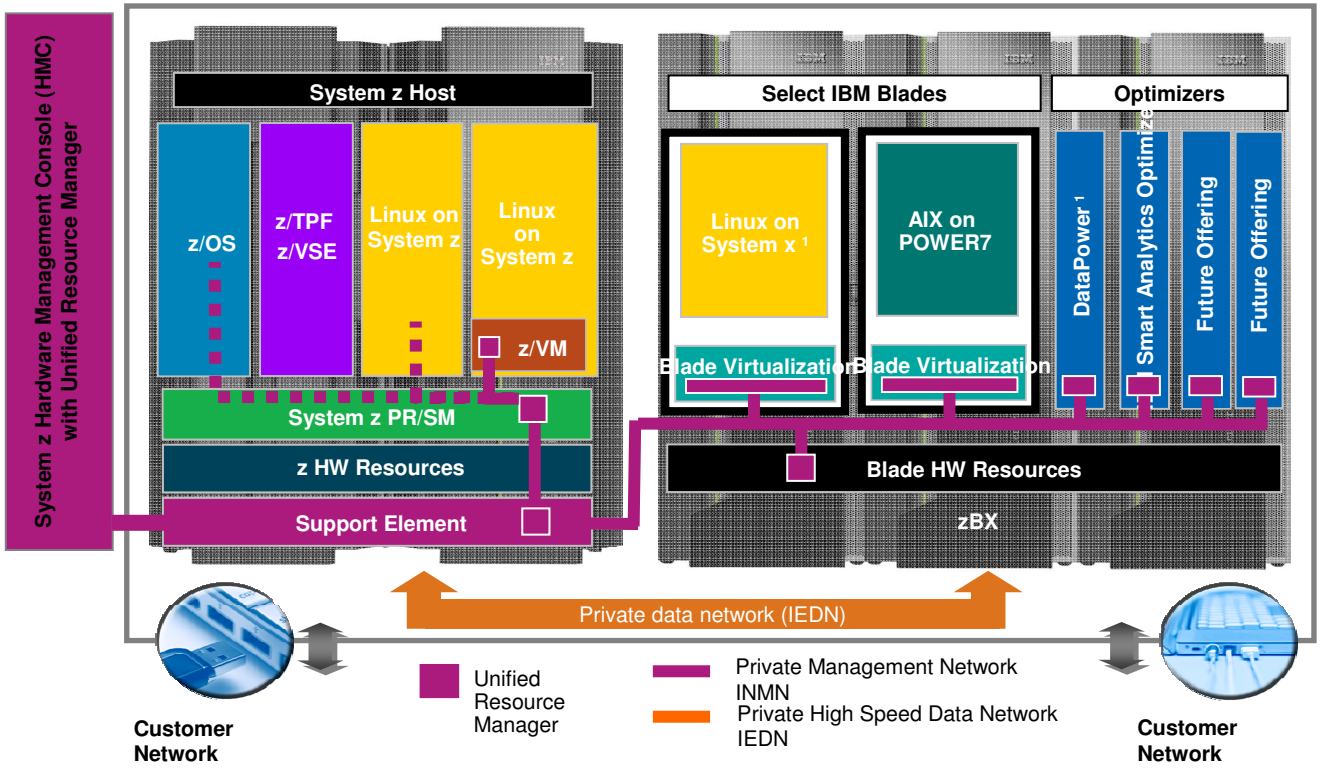

Manage different Hypervisors as Centralized resource.

Monitoring and Managing the Enterprise – zEnterprise will enable the management of Resources across Virtual Servers

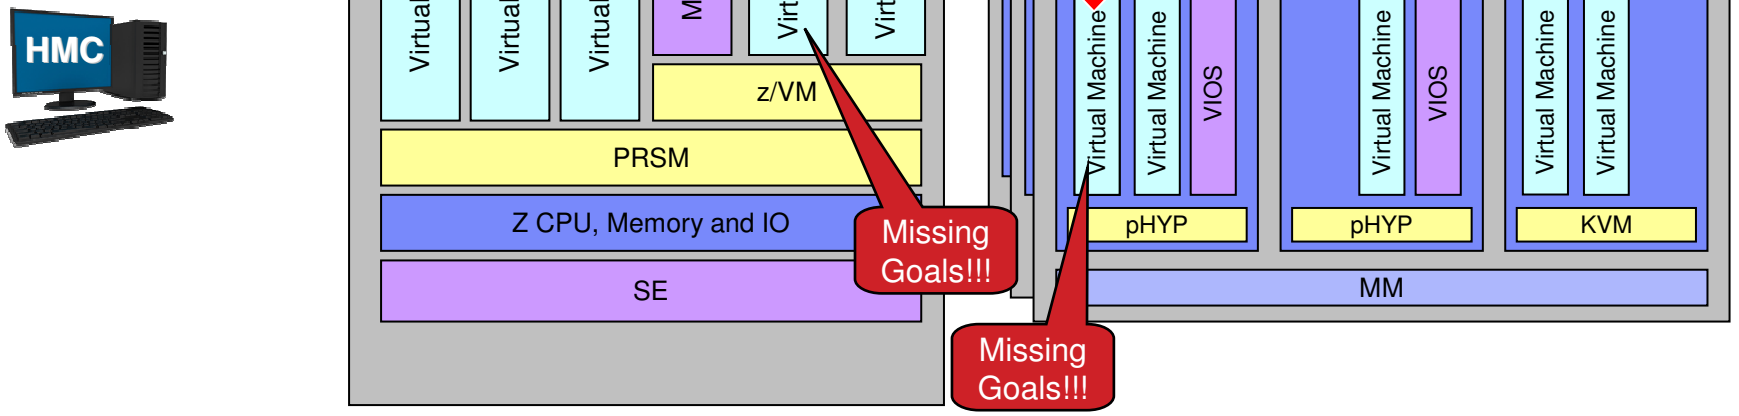

- Manage resources across virtual servers to achieve workload goals
  - Detect that a virtual server is part of Workload not achieving goals
  - Determine that the virtual server performance can be improved with additional resources
  - Project impact on all effected Workloads of moving resources to virtual server
  - If good trade-off based on policy, redistribute resources
  - Initially support CPU management

Business views across the zEnterprise

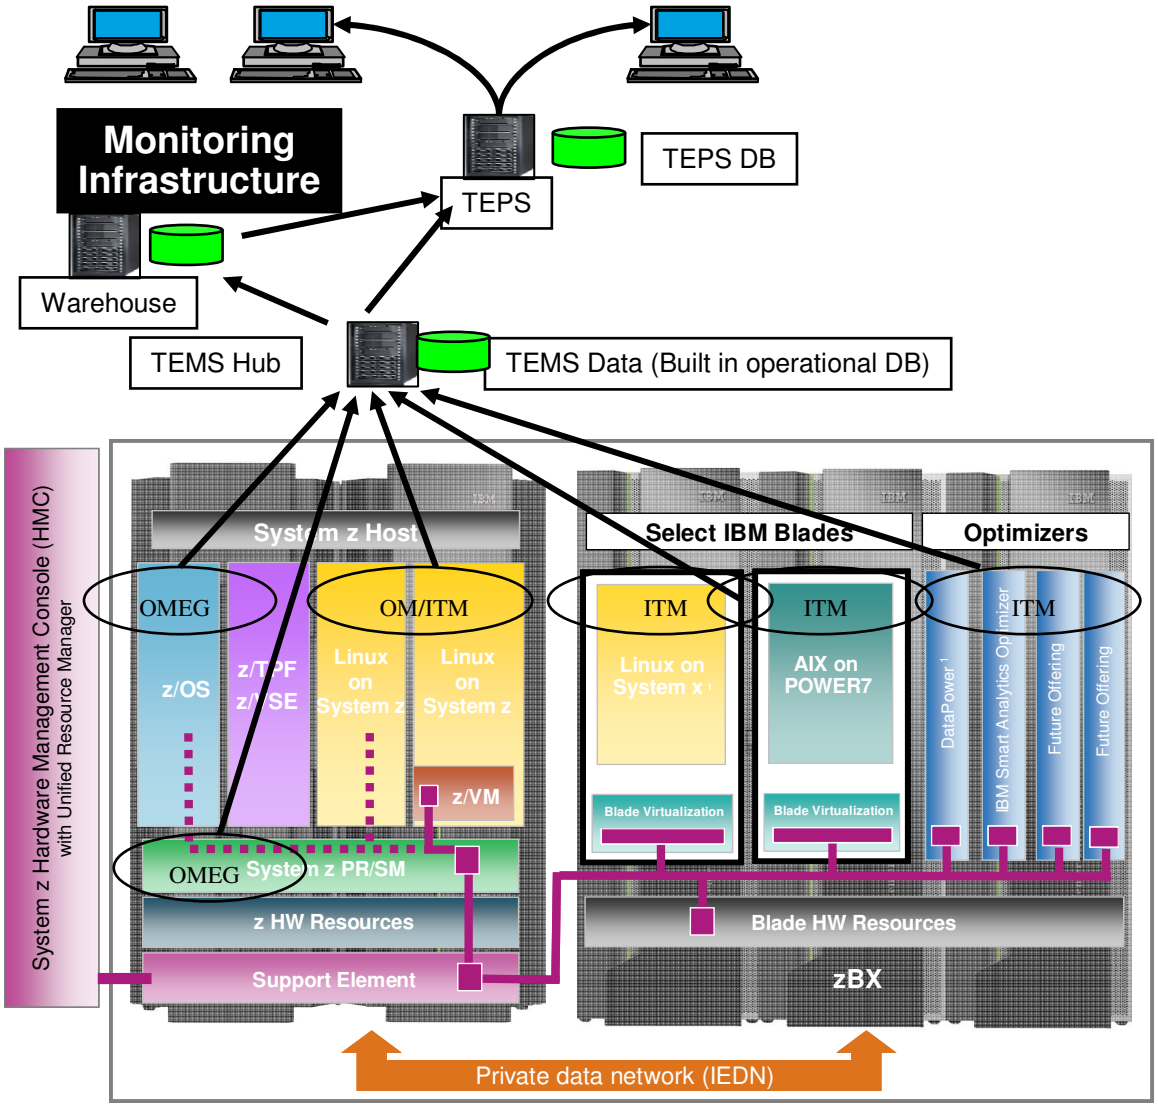

- ITM Infrastructure is shown separate to highlight components, however, each of these ITM components can reside on the zEnterprise.
- OMEGAMON agents can monitor z/OS system and subsystems, z/VM system and LPAR components, and Linux on z.
- ITM agents can monitor Linux on System z, Linux on System x, and AIX on Power7, and supported applications and databases.

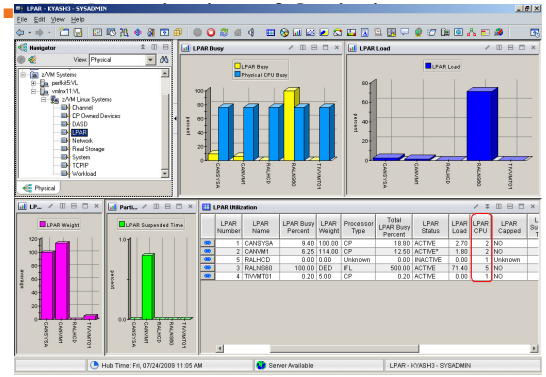

Common Interface across the zEnterprise

The future is ensembles and multiple hypervisors

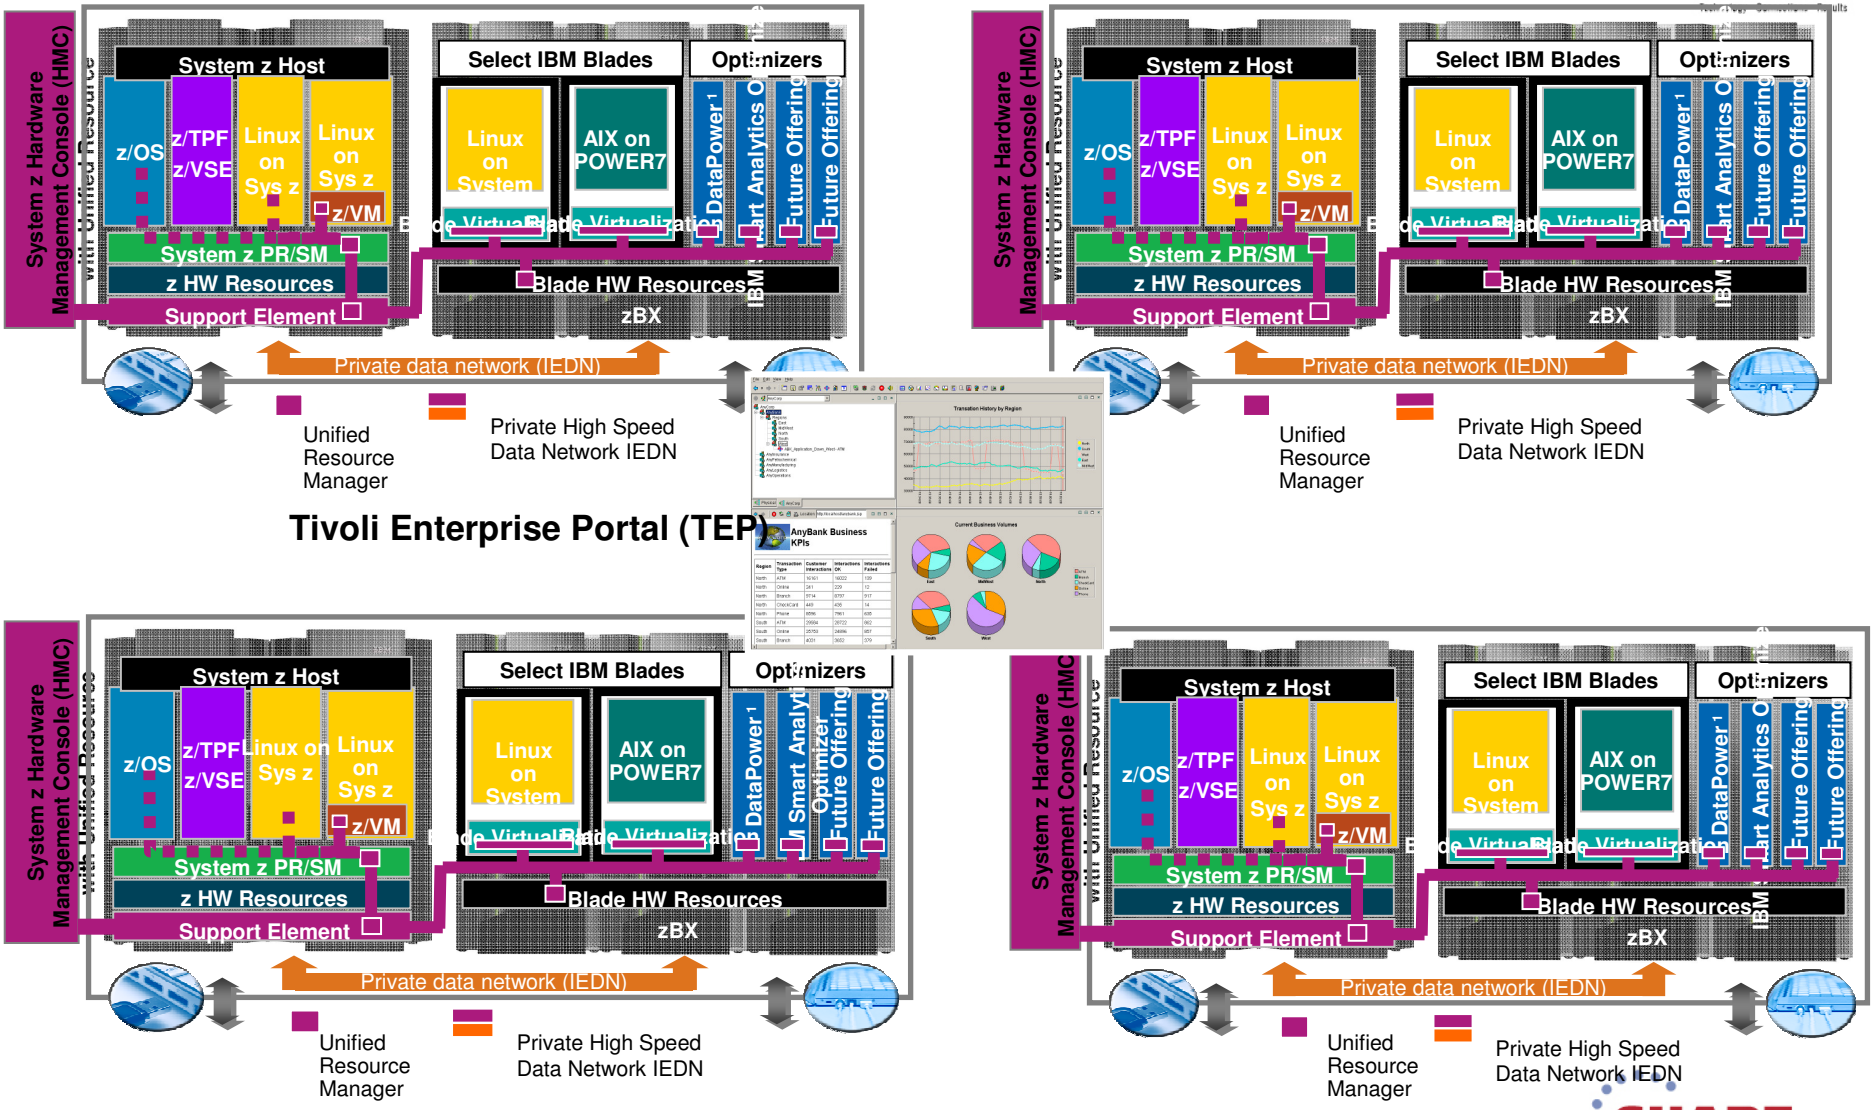

Manage different Hypervisors as Centralized resource.

AGENDA

- Introduction
- Monitoring Requirements
  - Virtual Linux and z/VM performance considerations
  - Don’t forget the hardware
  - Integration from hardware – systems – applications Persistent historical views
- Enterprise Management
- Operational Requirements
  - Centralized Control
  - Including all Enterprise Virtual Machines
- Integrating Monitoring and Operations
- Bringing it all together

Operations Manager for z/VM

Increase productivity
- Authorized users view and interact with monitored virtual machines without logging onto them
- Multiple users view/interact with a virtual machine simultaneously

Improve system availability
- Monitor virtual machines and processes
- Take automated actions based on console messages
- Reduce problems due to operator error

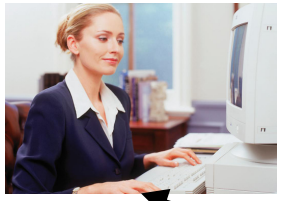

- View & interact with consoles
- View spool files

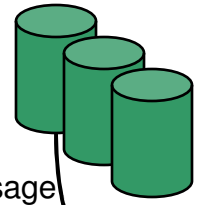

Monitor spool usage

Console monitoring

Service Virtual Machine being monitored

Operations Manager for z/VM

Take action

Service Virtual Machine being monitored

Schedule tasks

Console monitoring

Automation
- Routine activities done more effectively with minimal operations staff
- Schedule tasks to occur on a regular basis

Integration
Fulfill take action requests from OMEGAMON XE on z/VM and Linux

Features and Functions

- Monitor service machine consoles
- Monitor spool usage
- Monitor system events
- View and interact with monitored consoles from authorized user IDs
- Find and view spool files
- Schedule events/actions
- Dynamic configuration
- Separation of access control

Monitor Service Machine Consoles

Operations Manager

Test Data

```
TEST Message 1
TEST Message 2
TEST Message 3
…
```

OPERATOR

```
OPER Message 1
OPER Message 2
OPER Message 3
…
```

LINUX

```
LNX Message 1
LNX Message 2
LNX Message 3
…
```

TCP/IP

```
TCP Message 1
TCP Message 2
TCP Message 3
…
```

syslog data

```
slog Message 1
slog Message 2
slog Message 3
…
```

Data space 1

```
TEST Message 1
TEST Message 2
…
```

Data space 2

```
OPER Message 1
OPER Message 2
…
```

Data space 3

```
LNX Message 1
LNX Message 3
…
```

Data space 4

```
TCP Message 1
TCP Message 2
…
```

Data space 5

```
slog Message 1
slog Message 2
slog Message 3
…
```

Data space 6

```
OPER Message 1
LNX Message 1
LNX Message 2
TCP Message 1
slog Message 1
slog Message 2
TEST Message 1
OPER Message 2
…
```

Daily log

```
DIRM Message 1
LNX Message 1
LNX Message 2
TCP Message 1
DIRM Message 2
TCP Message 2
```

Enterprise level console/syslog management:

- View & interact with consoles
- View spool files

Monitor spool usage

Console monitoring

Service Virtual Machine being monitored

Operations Manager for z/VM

Take action

Service Virtual Machine being monitored

Schedule tasks

Console monitoring

- Centralized console/syslog management.
- Message log console for operations and automation.
- Similar to z/OS console management for host based operations.

System z Hardware Management Console (HMC) with Unified Resource Manager

System z Host

z/OS

z/TPF z/VSE

Linux on System z

Linux on System z

Syslog

Console

z/VM

System z PR/SM

z HW Resources

Support Element

Select IBM Blades

Optimizers

Linux on System x [1]

AIX on POWER7

Syslog

Syslog

Blade Virtualization

Blade Virtualization

DataPower [1]

IBM Smart Analytics Optimizer

Future Offering

Future Offering

Blade HW Resources

zBX

Private data network (IEDN)

Data space 6

```
OPER Message 1
LNX Message 1
LNX Message 2
TCP Message 1
slog Message 1
slog Message 2
TEST Message 1
OPER Message 2
…
```

Daily log

```
DIRM Message 1
LNX Message 1
LNX Message 2
TCP Message 1
DIRM Message 2
TCP Message 2
```

Monitor Service Machines

- Define rules to
  - Scan console messages for text matching
    - Includes column, wildcard, and exclusion support
    - Optionally restrict to specific user ID(s)
  - Take actions based on matches
- Multiple rules can apply to one message
  - Rules processed in order of definition in the configuration file
  - FINAL option available to indicate no additional rules should be evaluated

View and Interact with Consoles

- Authorized users can view live consoles of monitored service machines and guests
  - Multiple users can view the same console simultaneously
  - No need to logon to the service machine to see its console
  - Test data and Linux syslog data treated as a “console”
  - Views can be defined to look at a group of consoles in one view
- Full screen mode
  - Scroll up and down to view and search historical data
  - Auto scroll (on or off) as new output is displayed on the console
  - From command line, issue commands back to the monitored console
- Amount of data that is visible depends on specified or default data space size
- Rules/actions may modify the view
  - Suppress messages from the console
  - Hold or highlight messages with color, blinking, etc.
- Authorized users can view the log file
  - Can also request a copy of the log file from today or a previous day

Monitor and View Spool Files

- Create spool monitors to trigger actions when
  - Percent of spool usage falls within a specified range
  - Percent of spool usage increases at a specified rate
- Actions triggered can be the same actions used by console monitoring
- Authorized users can
  - Display a list of spool files based on one or more attributes
    - Owner
    - Size
    - Date created
  - From the list the user can
    - View the contents of an individual spool file
    - Transfer, change, or purge a spool file

Schedule Events and Actions

- Define schedules
  - Hourly, daily, weekly, monthly, or yearly, nth weekday of the month
  - Once on specified month, day, year, and time
  - At regular intervals
    - Every x hours and y minutes
  - Within a specified window of time
    - Specify start time
    - Specify conflicting schedules
    - Specify maximum time to defer this schedule
  - Within limits
    - Restrict to specific days of the week: Monday through Sunday plus holidays
    - Restrict to certain hours of the day
- Specify the action associated with the schedule
  - Actions specified are the same as those for console and spool monitoring

Respond to System Events

- Create monitors for z/VM system events (*VMEVENT) related to user IDs
  - Logon
  - Logoff
  - Failure condition (typically CP READ)
  - Logoff timeout started
  - Forced sleep started
  - Runnable state entered (VM READ)
  - Free storage limit exceeded
- Optionally restrict to specific user ID(s)
- Specify the action associated with the event
  - Actions specified are the same as those for schedules and console and spool monitors

Dynamic Configuration

- Initial configuration file loaded at startup
  - May imbed other configuration files
- Most configuration options can be updated while Operations Manager is running
  - Add, delete, or change:
    - Rules, actions, monitors, schedules, holidays, groups, user authorization
  - Suspend or resume rules, monitors, schedules
- Multiple methods
  - GOMCMD command interface
  - Load a new or updated configuration file
  - Commands in DEFACTN statements

Operations Manager

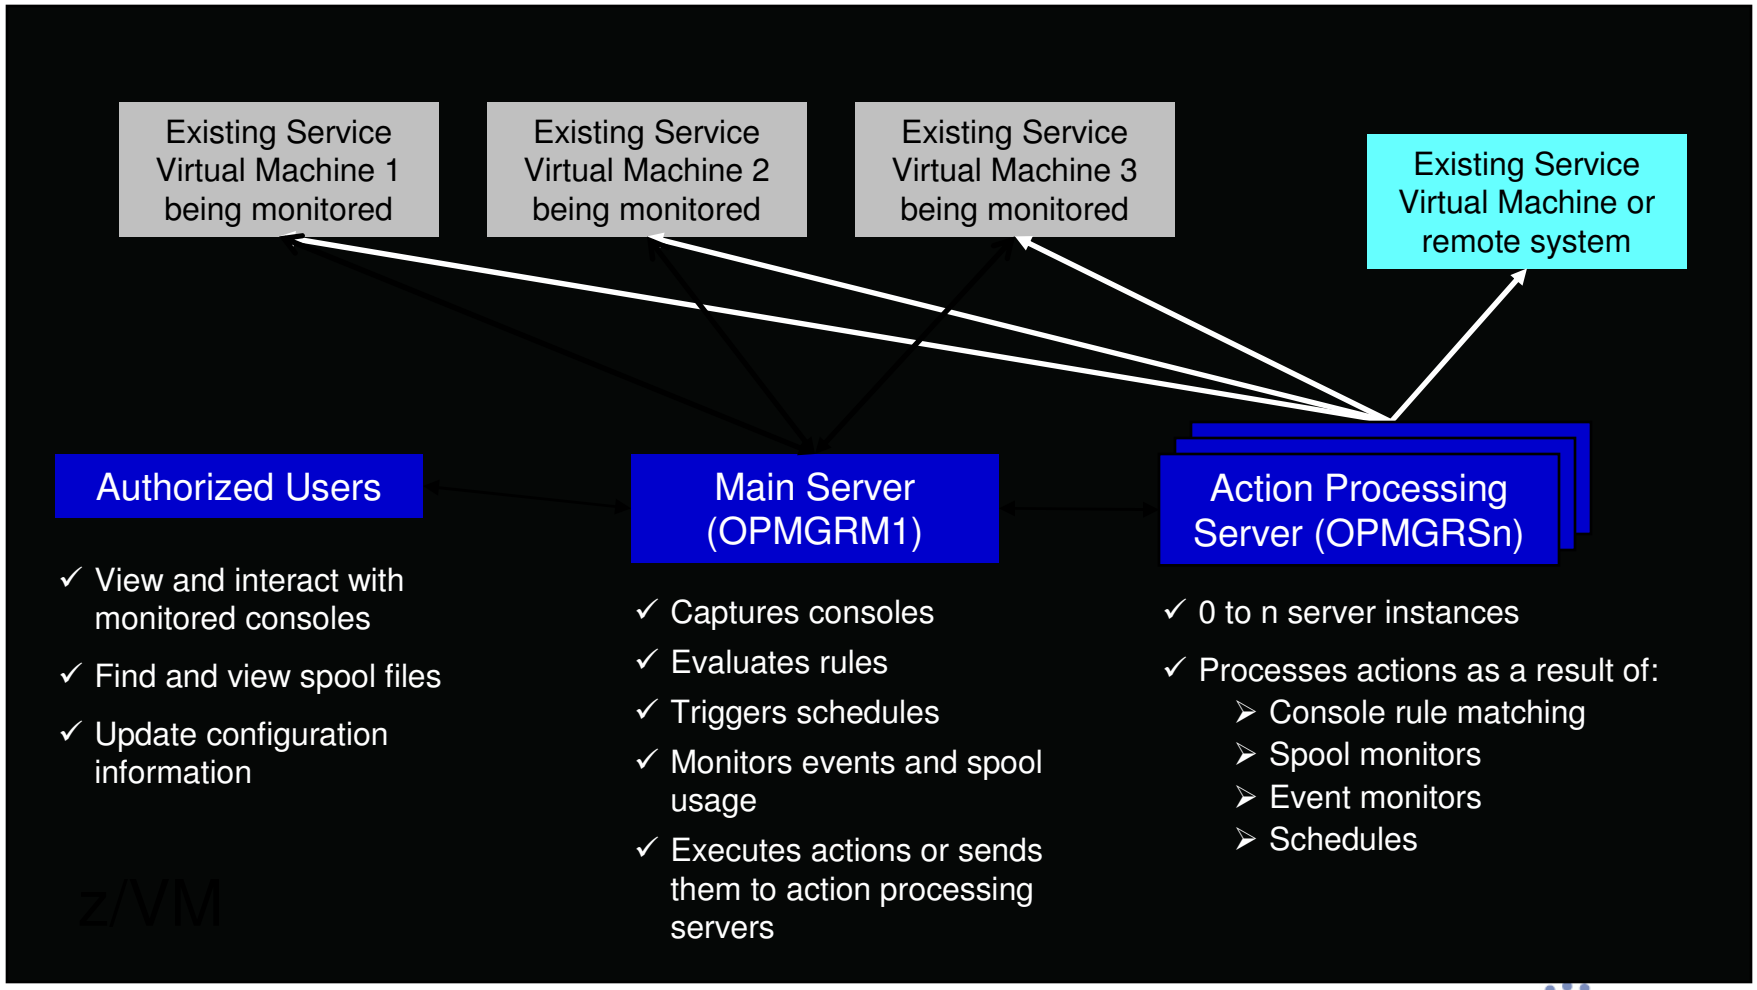

Summary

- Use Operations Manager to
  - Automate daily operations
  - Prevent problems rather than react to them
  - Automate reactions to problems when they can't be prevented
  - Improve problem determination procedures
  - Increase programmer and operator productivity

AGENDA

- Introduction
- Monitoring Requirements
  - Virtual Linux and z/VM performance considerations
  - Don’t forget the hardware
  - Integration from hardware – systems – applications Persistent historical views
- Enterprise Management
- Operational Requirements
  - Centralized Control
  - Including all Enterprise Virtual Machines
- Integrating Monitoring and Operations
- Bringing it all together

Bring it all together

It is often that a unit of work is serviced by multiple applications and databases across multiple operating systems, including z/VM and Linux. Integrated views allow:

- Unit of work, or application tracking
- Business views
- Single skill sets to monitor dissimilar hardware, operating system, and application environments.

Application View: Scaling Scenario

- WebSphere MQ on Linux for System z receives “order requests” in the form of Queue messages, and places them on a queue.
- A WebSphere Application Server is invoked to periodically check the queue for messages and process them to a DB2 on z/OS database.
- The orders are coming too fast for the Websphere application to process.
- A second Linux server is started with another copy of Websphere application server to aid in the processing of requests.

Application View: Scaling Scenario

- Trigger: Queue Depth
- Options for triggering actions can be based on things such as:
  - The number of orders received but not yet processed (the number of messages on the queue)
  - The amount of time it is taking to process the orders
  - The response time of the web application
  - The CPU usage of the z/VM Guest
  - Other things I haven't given much thought to yet.

MQ Series Queue growth started

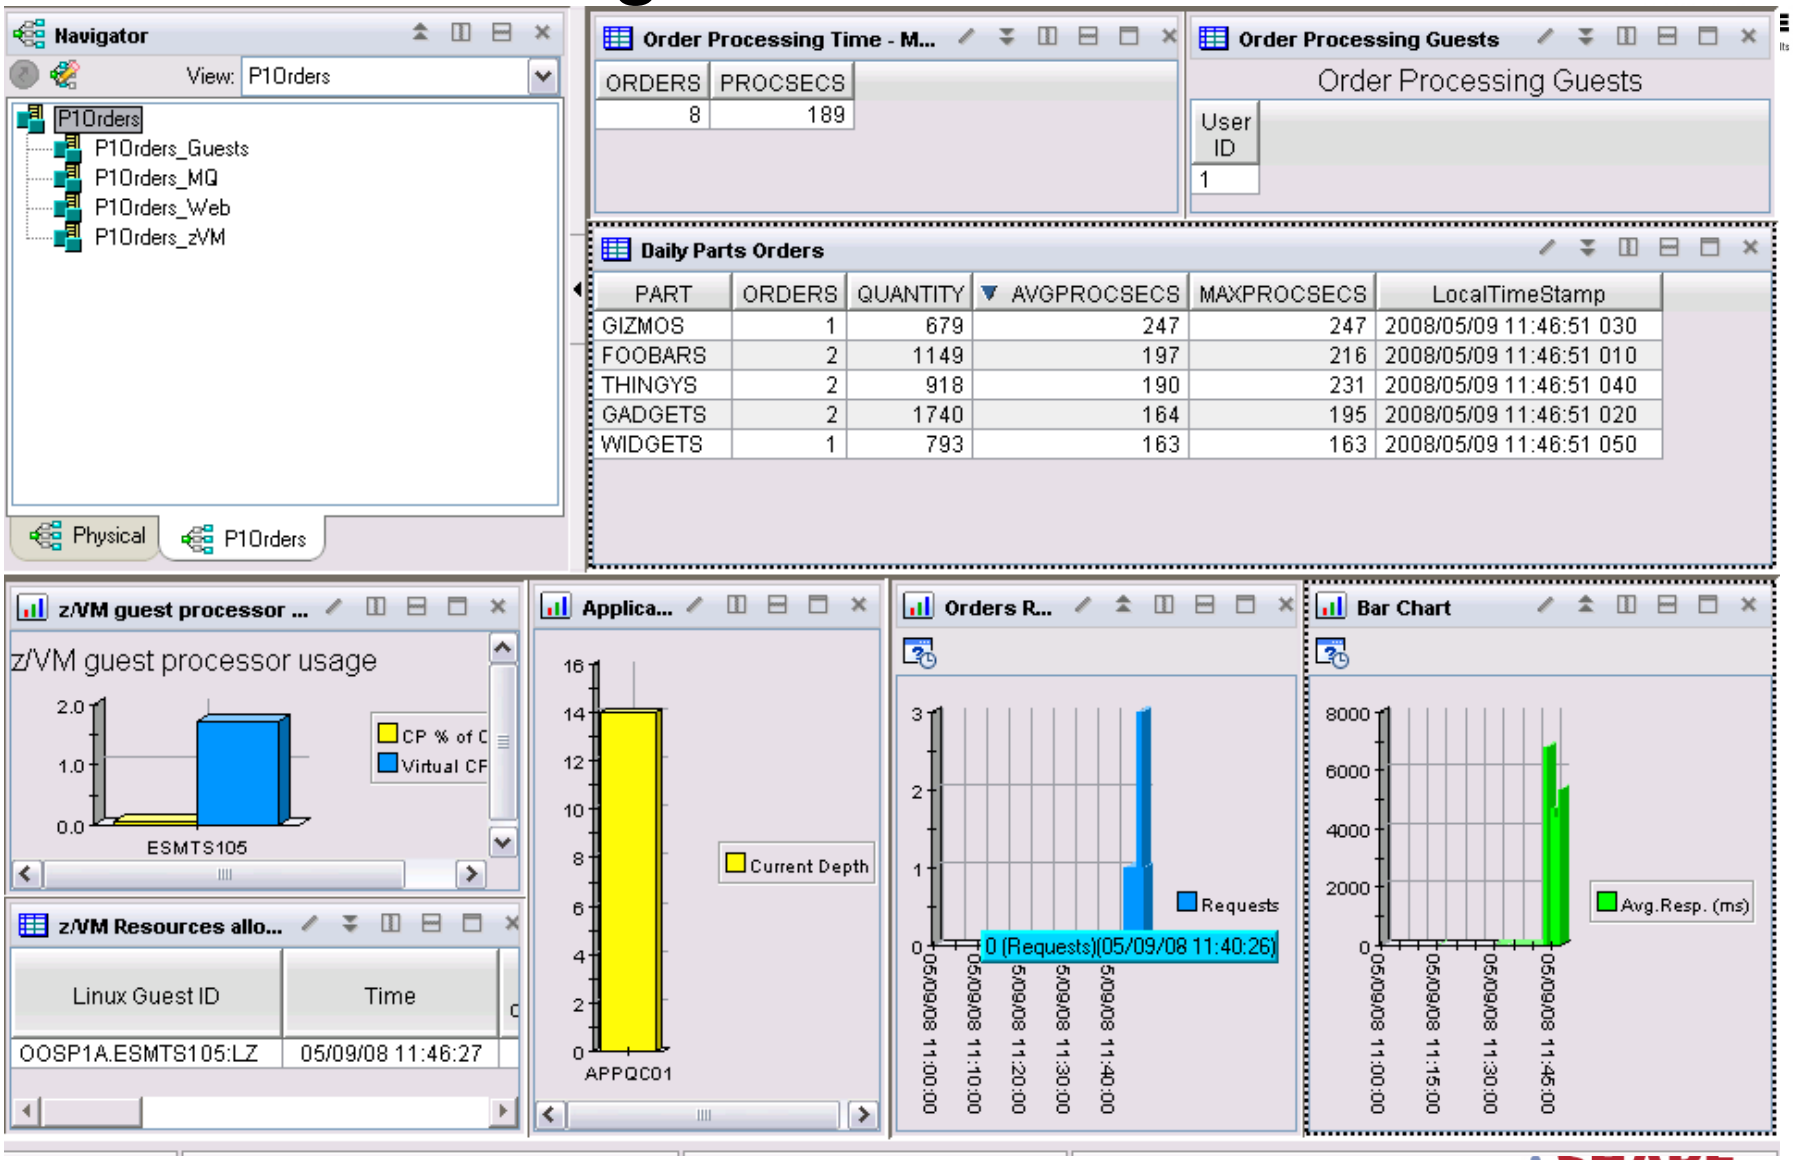

Scaling Scenario

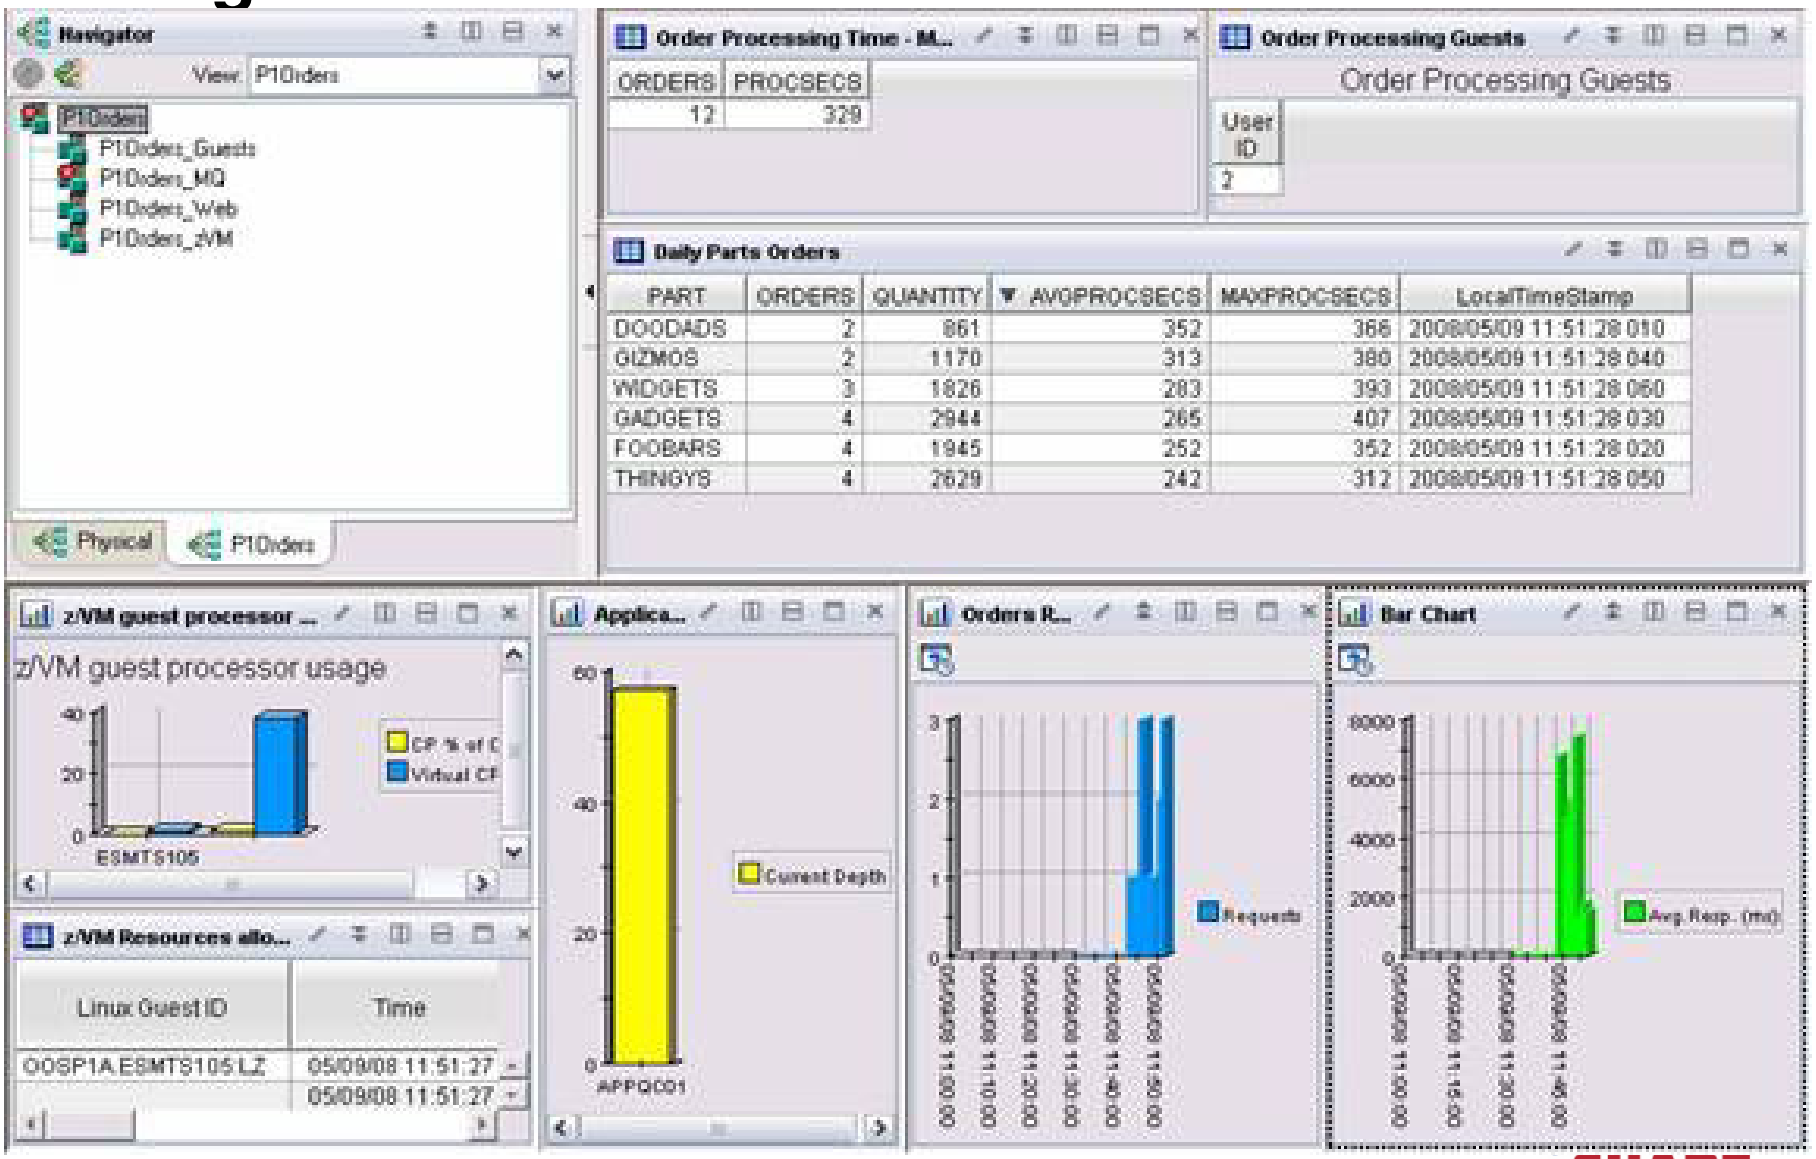

Adjusting Resources for a Linux Guest

- Virtual CPU consumption is high for a Linux guest
- Detect the alert
  - Automation receives the message
- Action is triggered by a rule in Operations Manager
- Operations Manager issues CP commands to tune the guest
  - SET QUICKDSP
  - SET SHARE
- Ability to monitor the output is key

Adjusting resources for a Linux guest

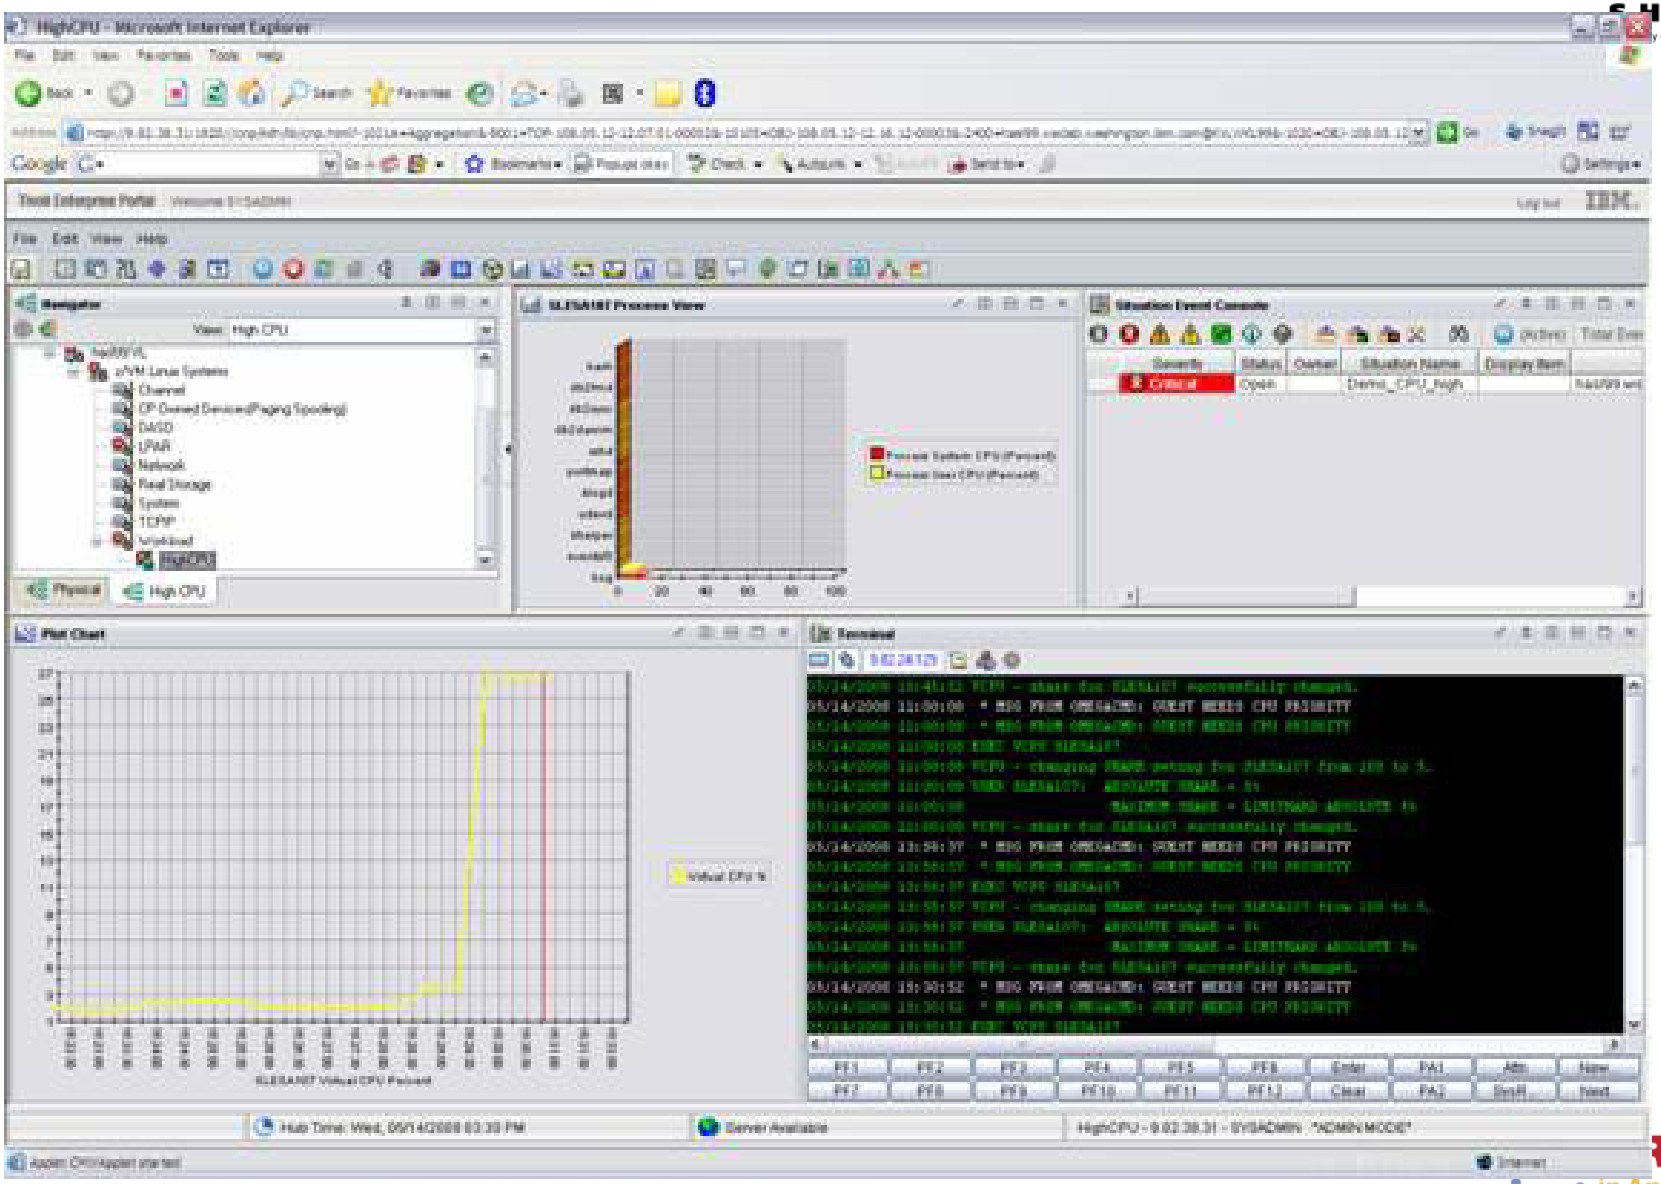

OMEGAMON Configuration

- Define a situation (alert) to detect high CPU consumption for Linux virtual machines.
- Define the automated “Take Action” to:
  - Direct a message to console monitored by Operations Manager.
  - Include in the message keywords to trigger Operations Manager rule.
    - Guest Name
    - Guest need CPU priority text
    - Any unique data desired for specific customer environment.

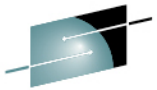

धन्यवाद

Hindi

多謝

Traditional Chinese

감사합니다

Korean

Спасибо

Russian

Gracias

Spanish

Thank You

English

شكراً

Arabic

Obrigado

Brazilian Portuguese

Danke

German

Grazie

Italian

多谢

Simplified Chinese

Merci

French

நன்றி

Tamil

ありがとうございました

Japanese

ขอบคุณ

Thai

Managing z/VM and Linux Performance Best Practices

Mike Sine
IBM

March 13, 2014
Session Number 14637