Pulse2011

Performance Management of z/VM and Linux

Nigel Bland
Vic Cross
Roberto Calderon

AGENDA

- Introduction
- Monitoring requirements
  - Virtual Linux and z/VM performance considerations
  - Don't forget the hardware
  - Integration from hardware – systems – applications
    Persistent historical views
- Why IBM
- 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 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.

OMEGAMON XE on z/VM and Linux agents

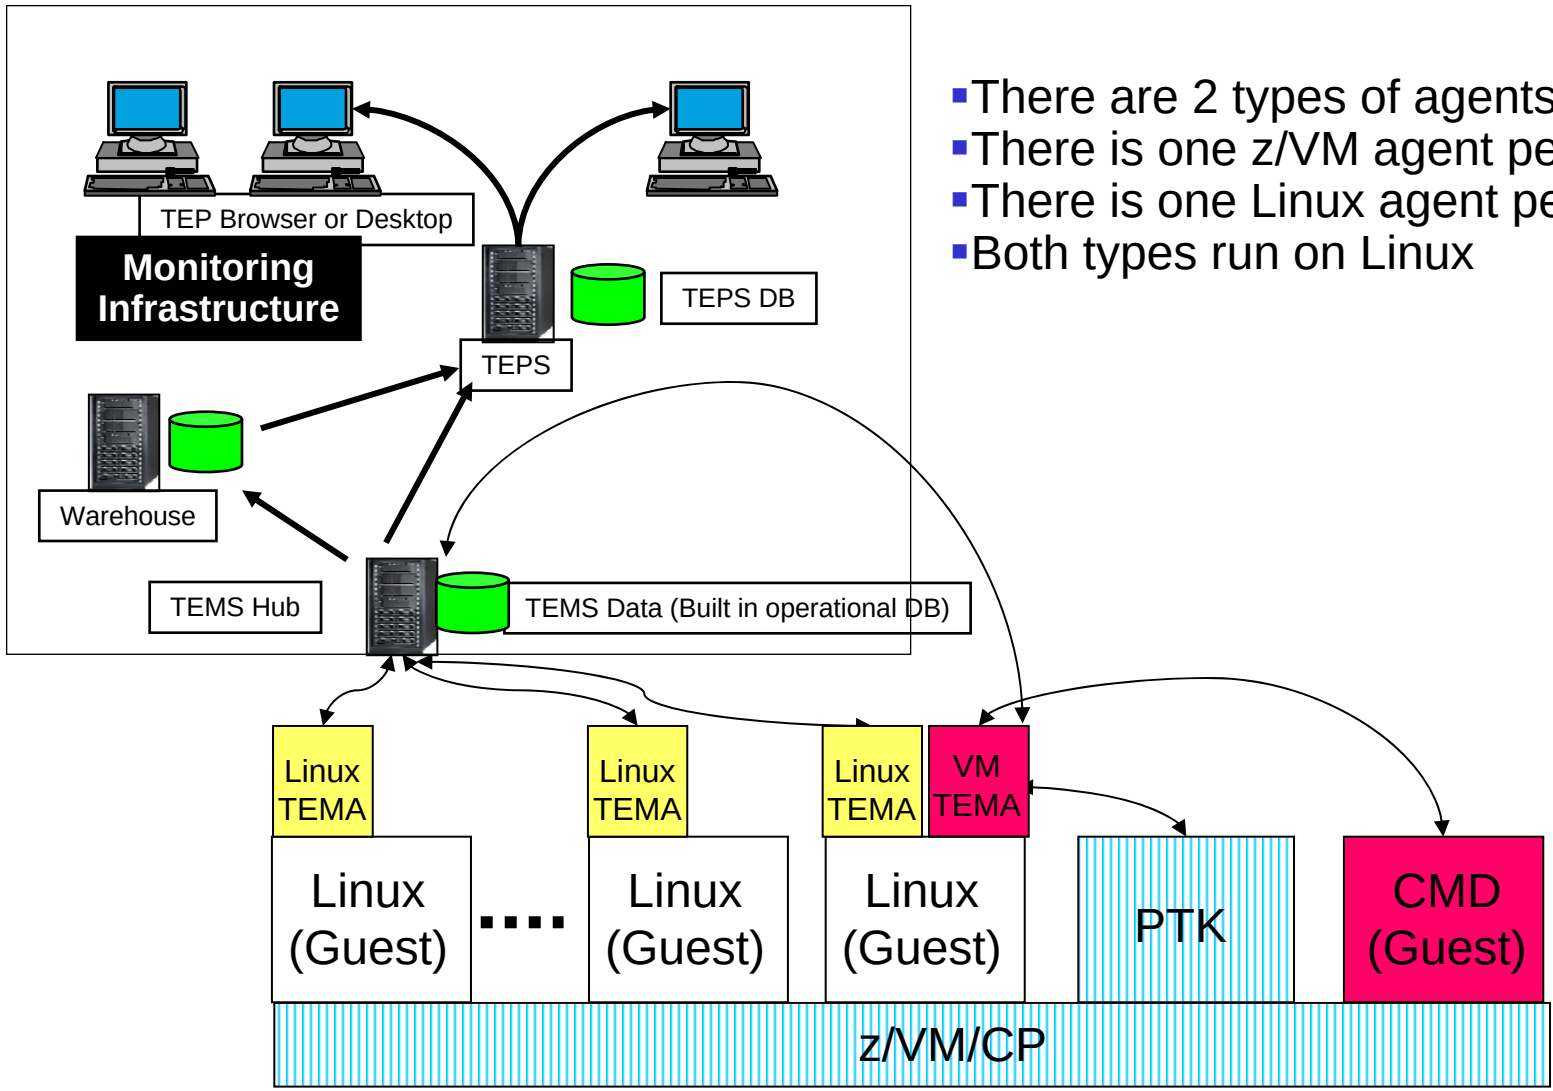

- There are 2 types of agents
- There is one z/VM agent per z/VM LPAR
- There is one Linux agent per Linux Guest
- Both types run on Linux

- Background Based Collectors Collect Based On Sampling Interval
  - Sampling Interval Can Be User Defined
    - XE MQ = 60 Seconds
    - XE MFN = 5 Minutes
    - XE z/VM Linux
      - CP Monitor Interval Can Be From 6 Secs To 1 HR – Default 5 Min
      - Linux TEMA Can Be Set On Interval
        - » Careful Not To Set Interval Less Than CP Monitor Interval
  - Situations Evaluate Data From The Last 'Take Sample'
  - Situation Intervals Need To Consider Collector Sampling Interval
  - There Is No Sync. Between Collector Start Time And Sit. Eval. Time
  - Same Data Could Be Evaluated More Than Once
  - Situation Impact On Performance Normally Is Less Than A Foreground Based Collector

- Assume A 5 Min. Collector Interval
- Situation Interval Of 1 Minute

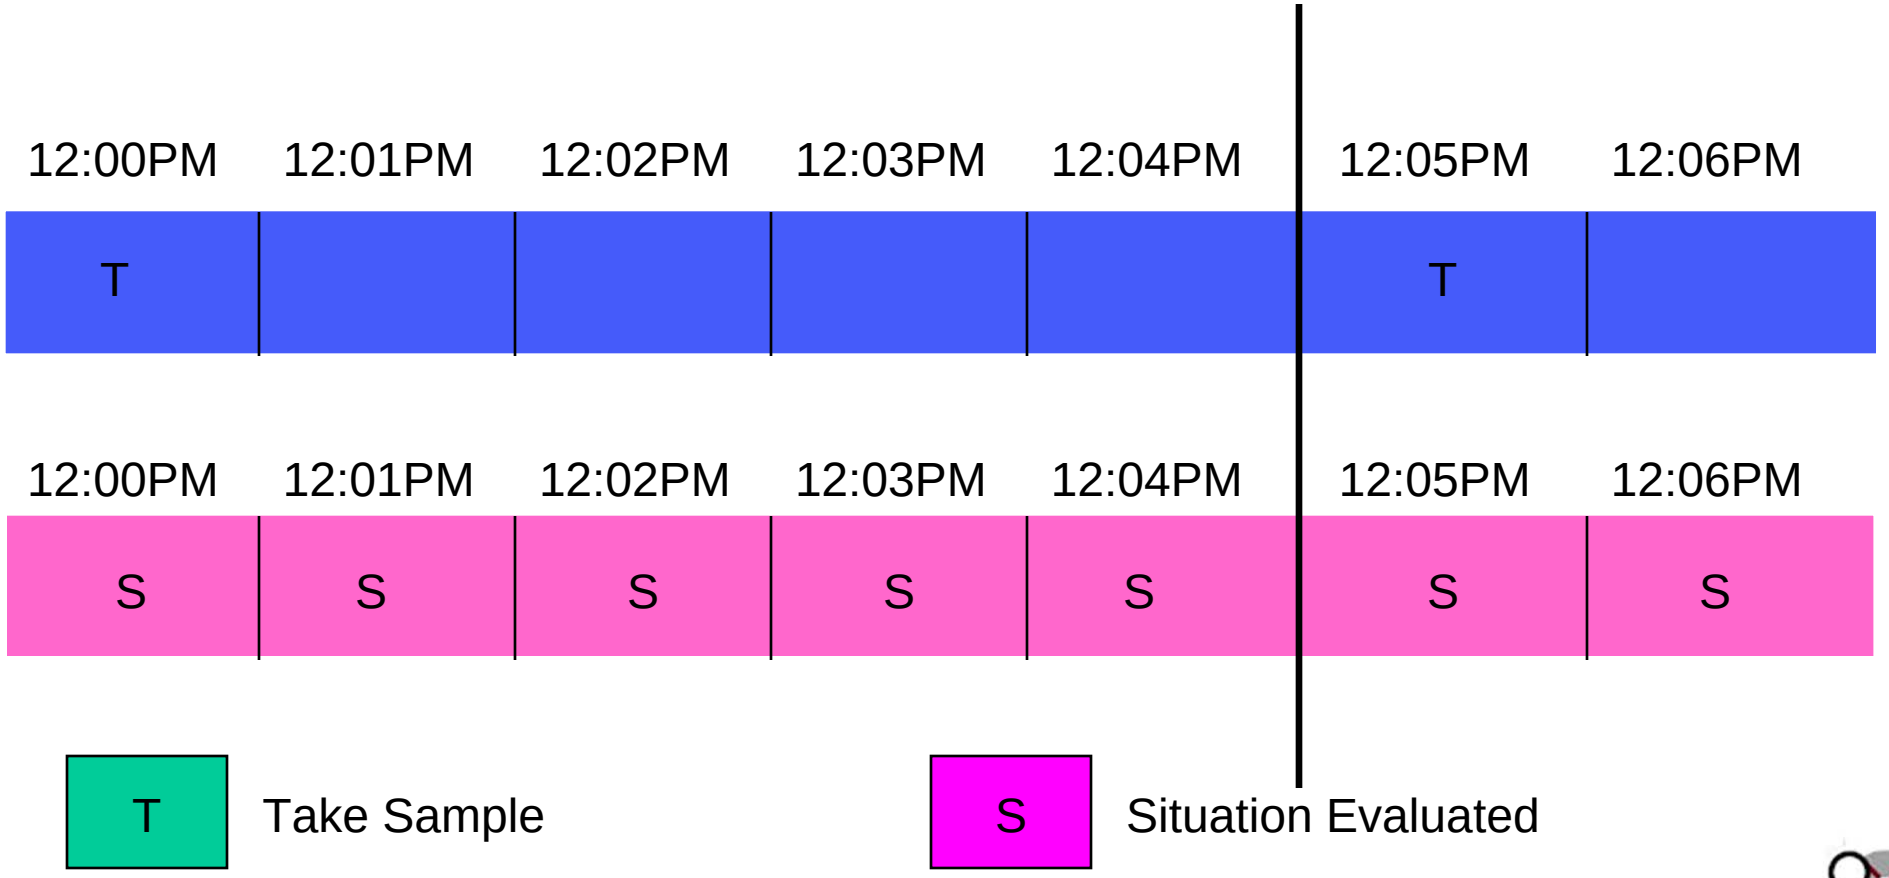

- Foreground Based Collectors Are Driven
  - Based On Situation Intervals
  - Based On Requests From TEP Users
    - It Can Be A Bad Idea To Have Frequent TEP Refresh Intervals
      - Or A User Pressing F5 Constantly
      - Or Reports Which Return Many Rows Refreshed Constantly
  - By UADVISOR (Historical) Data Collection
  - By CUA User Pressing F5 Or Auto Refresh
- 'Take Sample' Can Have A Bigger Impact On Performance
  - Each Take Sample Drives Data Collection From The Managed System

AGENDA

- Introduction
- Monitoring requirements
  - Virtual Linux and z/VM performance considerations
  - Don't forget the hardware
  - Integration from hardware – systems – applications
    Persistent historical views
- Why IBM
- Bringing it all together

OMEGAMON XE on z/VM and Linux An Integrated Monitoring Approach

- Provides performance monitoring for z/VM and Linux guests
- Linux agents gather performance data from Linux guests
- z/VM agent gathers performance data from z/VM
  - Including z/VM view of guests
  - Uses IBM Performance Toolkit for VM as its data source
- 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
  - Able to integrate z/VM and Linux into Enterprise Solution
  - Data warehousing for trend analysis

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

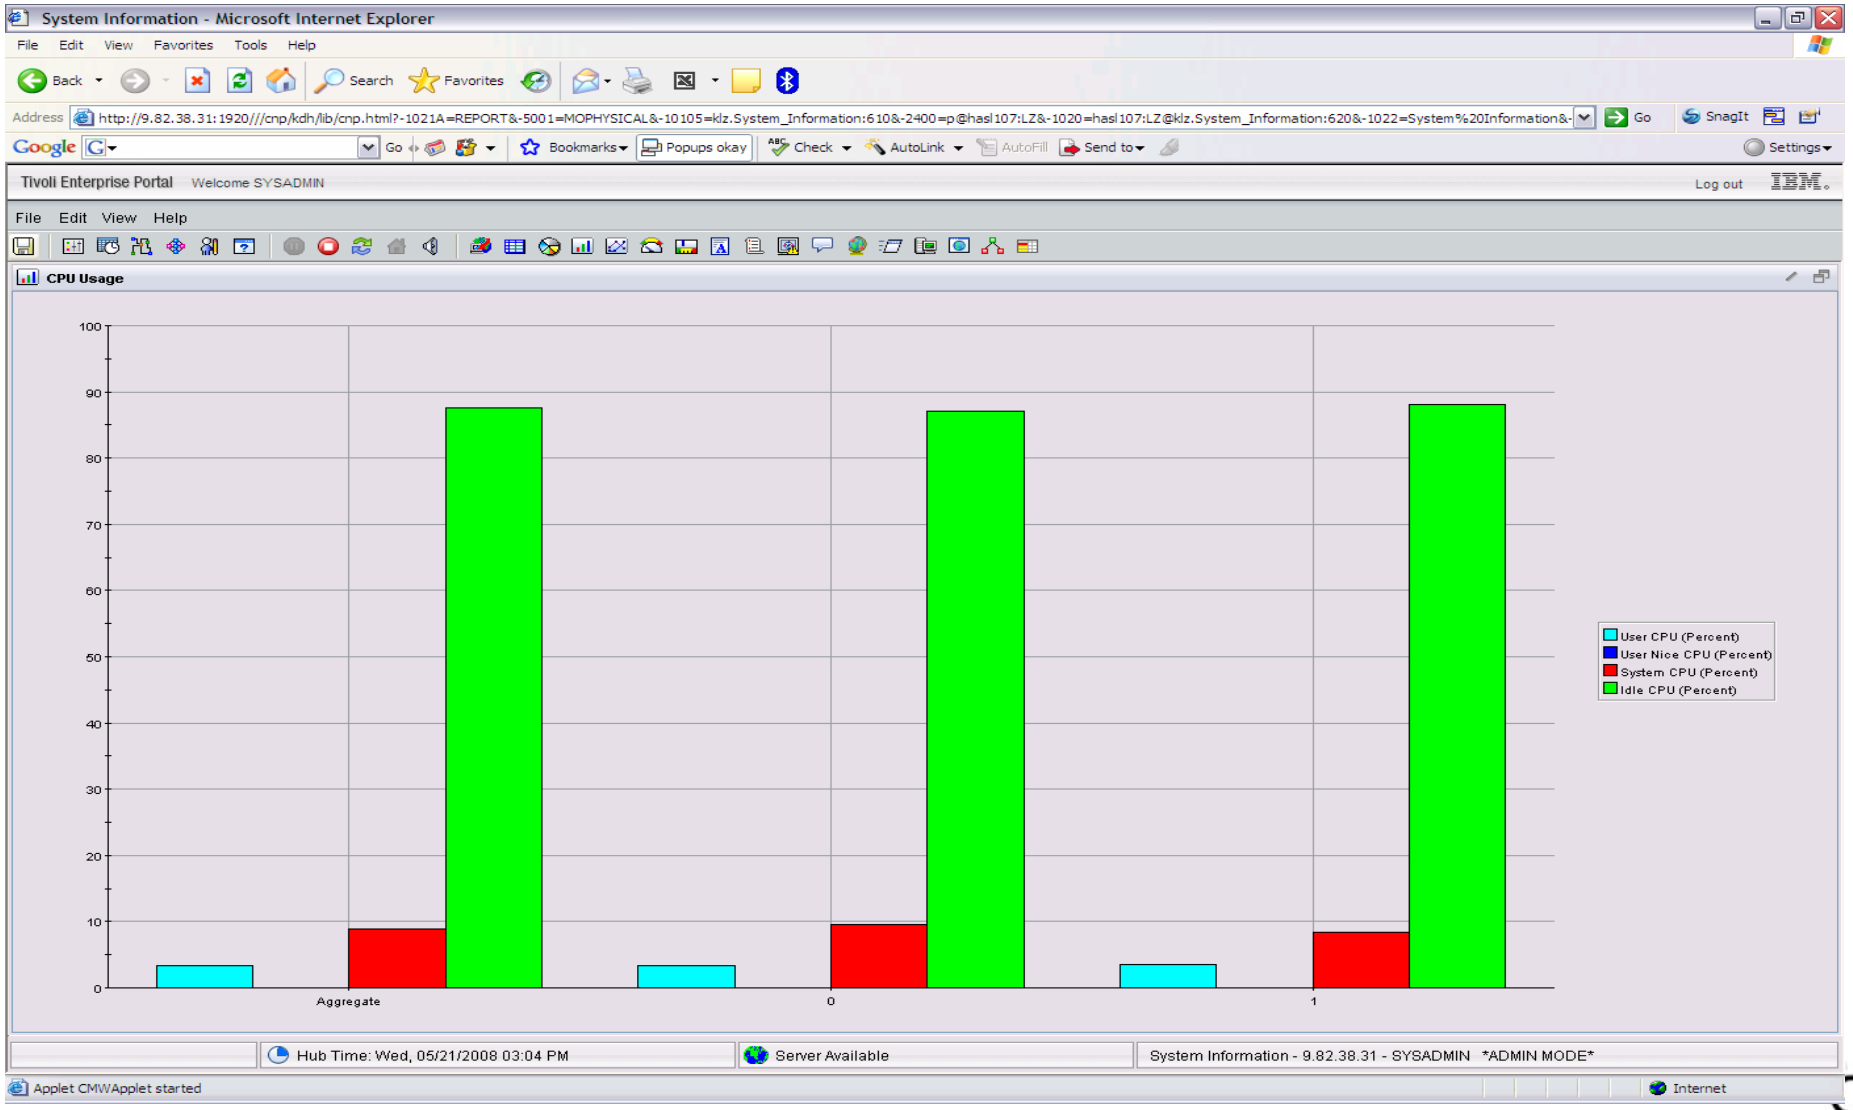

z/VM Processor Utilization

- 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

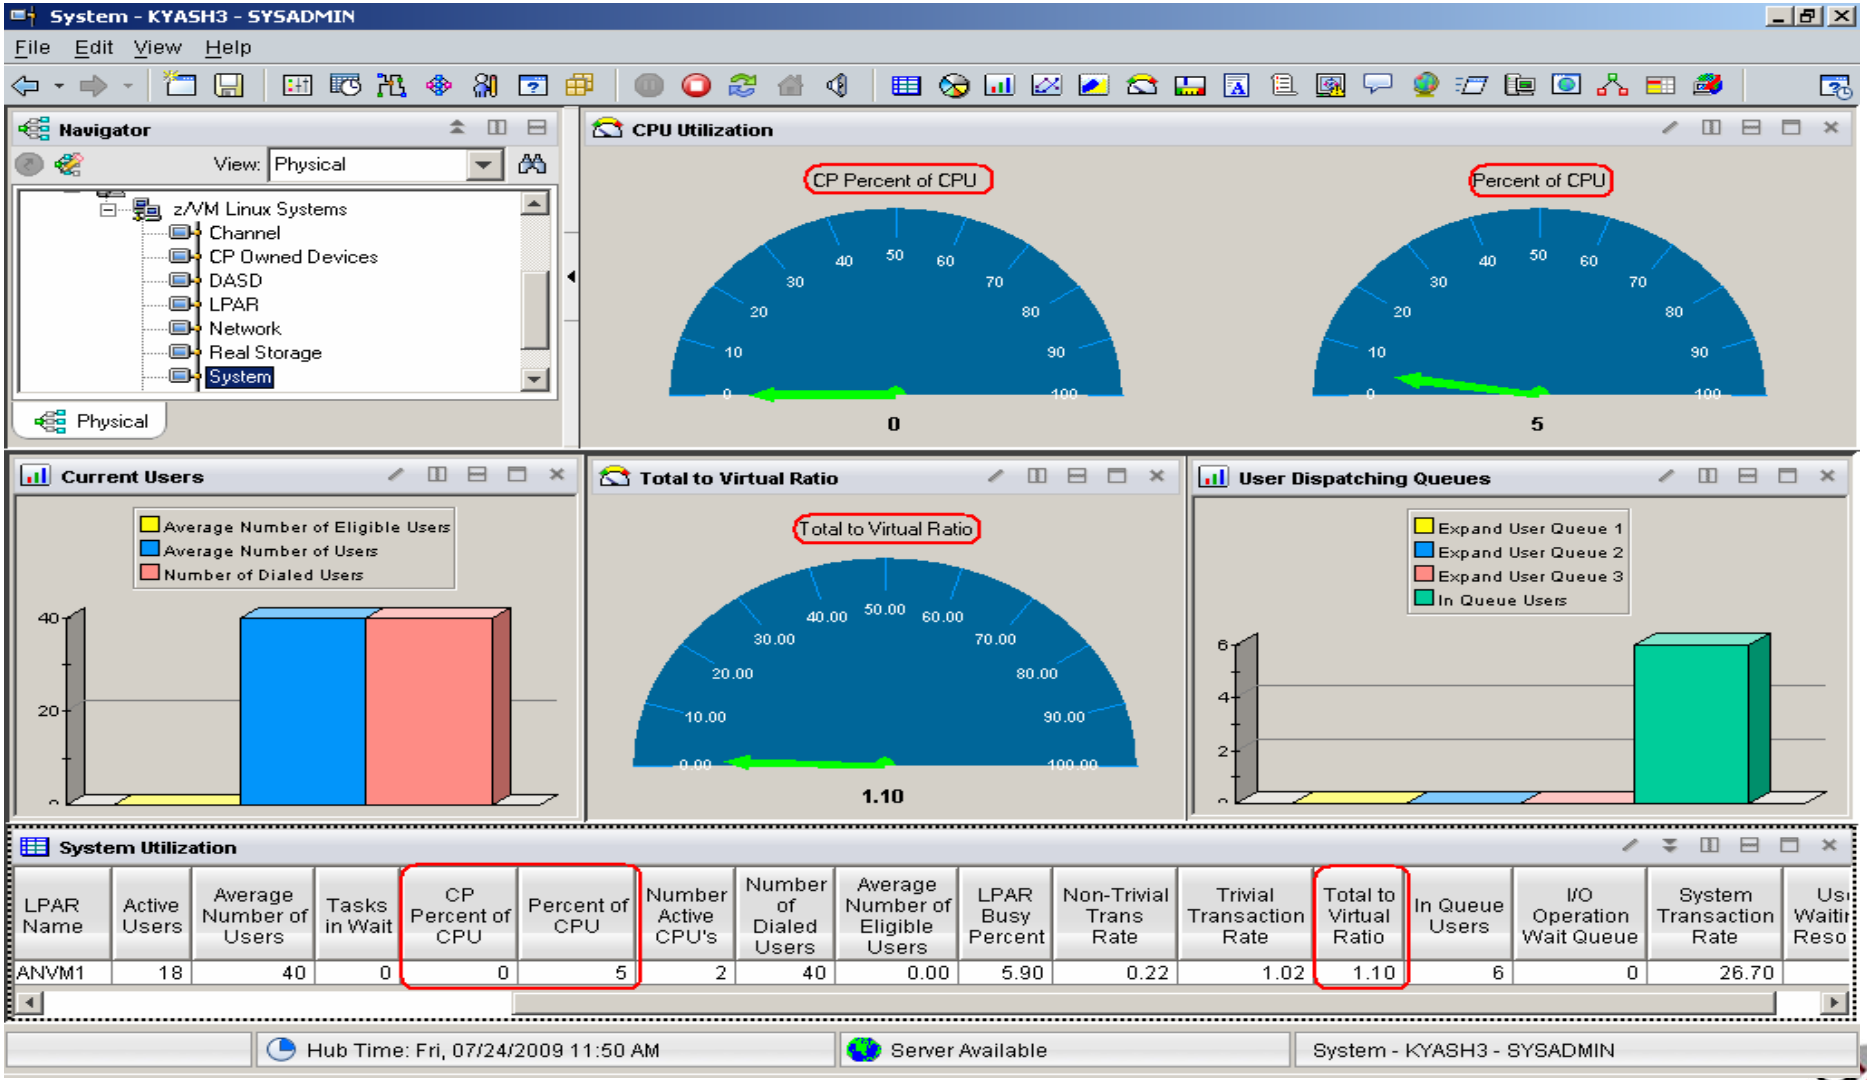

z/VM Workload Workspace

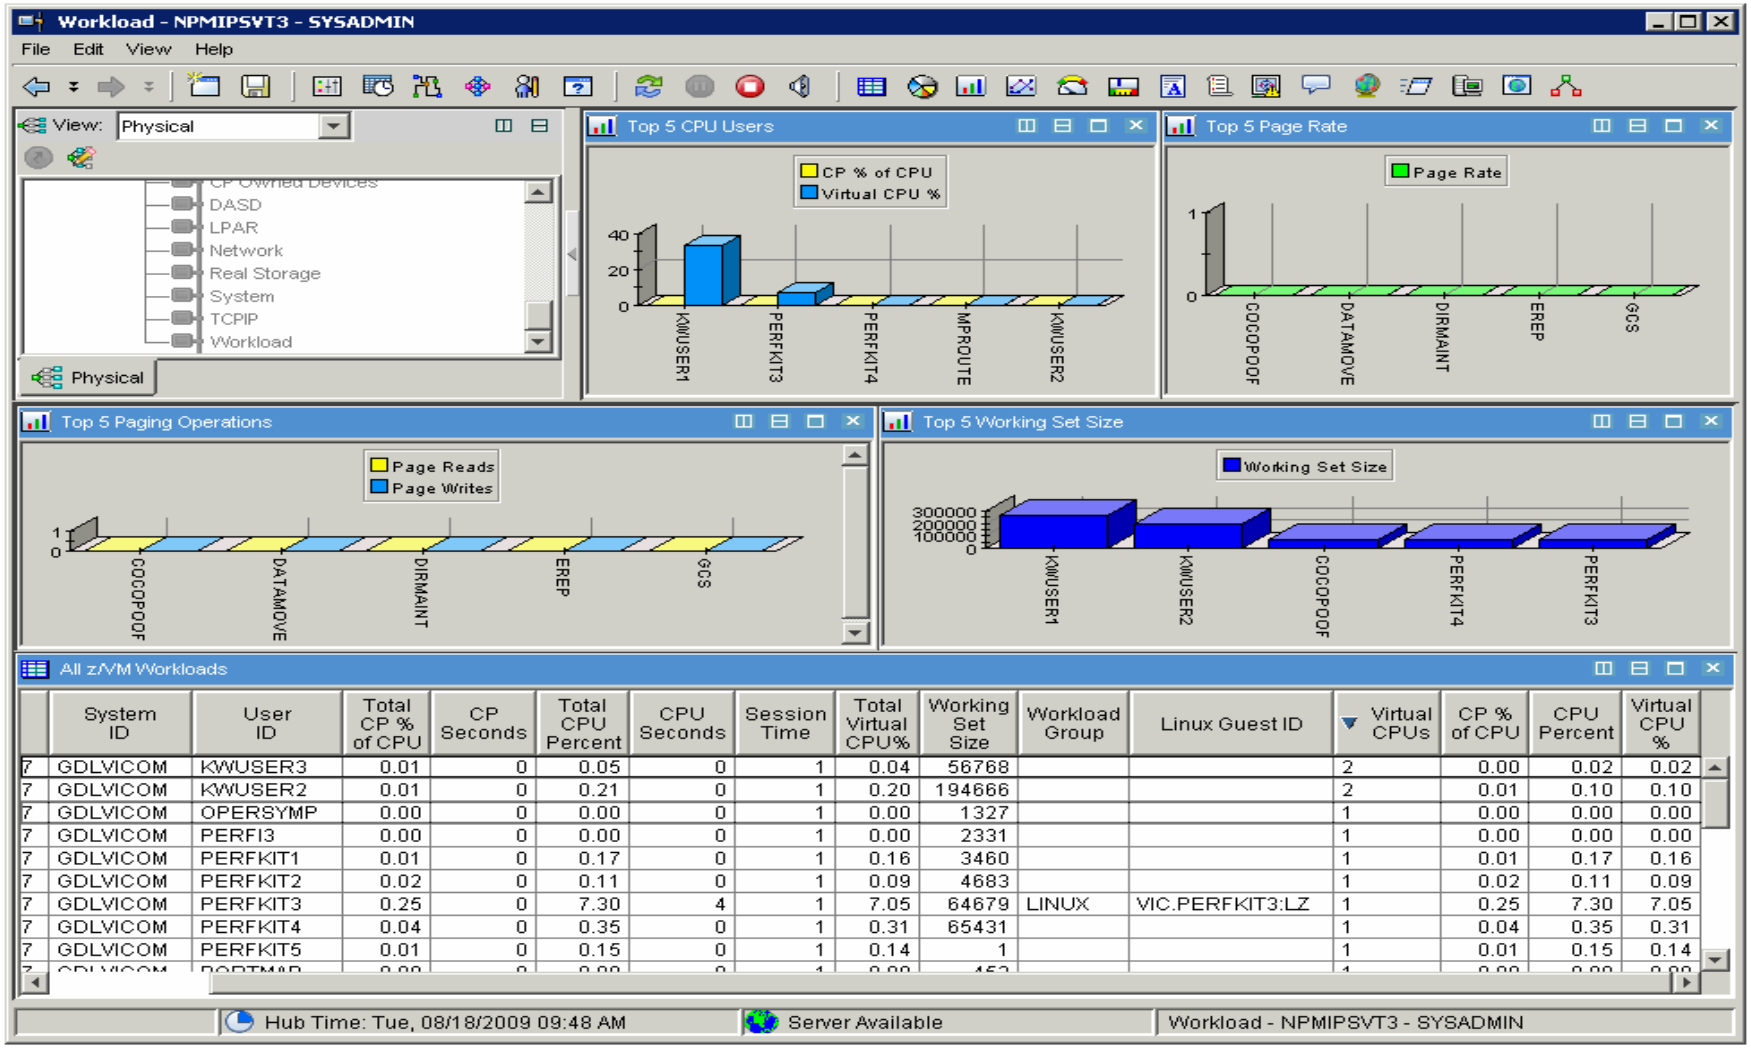

	System ID	User ID	Total CP % of CPU	CP Seconds	Total CPU Percent	CPU Seconds	Session Time	Total Virtual CPU%	Working Set Size	Workload Group	Linux Guest ID	Virtual CPUs	CP % of CPU	CPU Percent	Virtual CPU %
7	GDLVICOM	KWUSER3	0.01	0	0.05	0	1	0.04	56768			2	0.00	0.02	0.02
7	GDLVICOM	KWUSER2	0.01	0	0.21	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

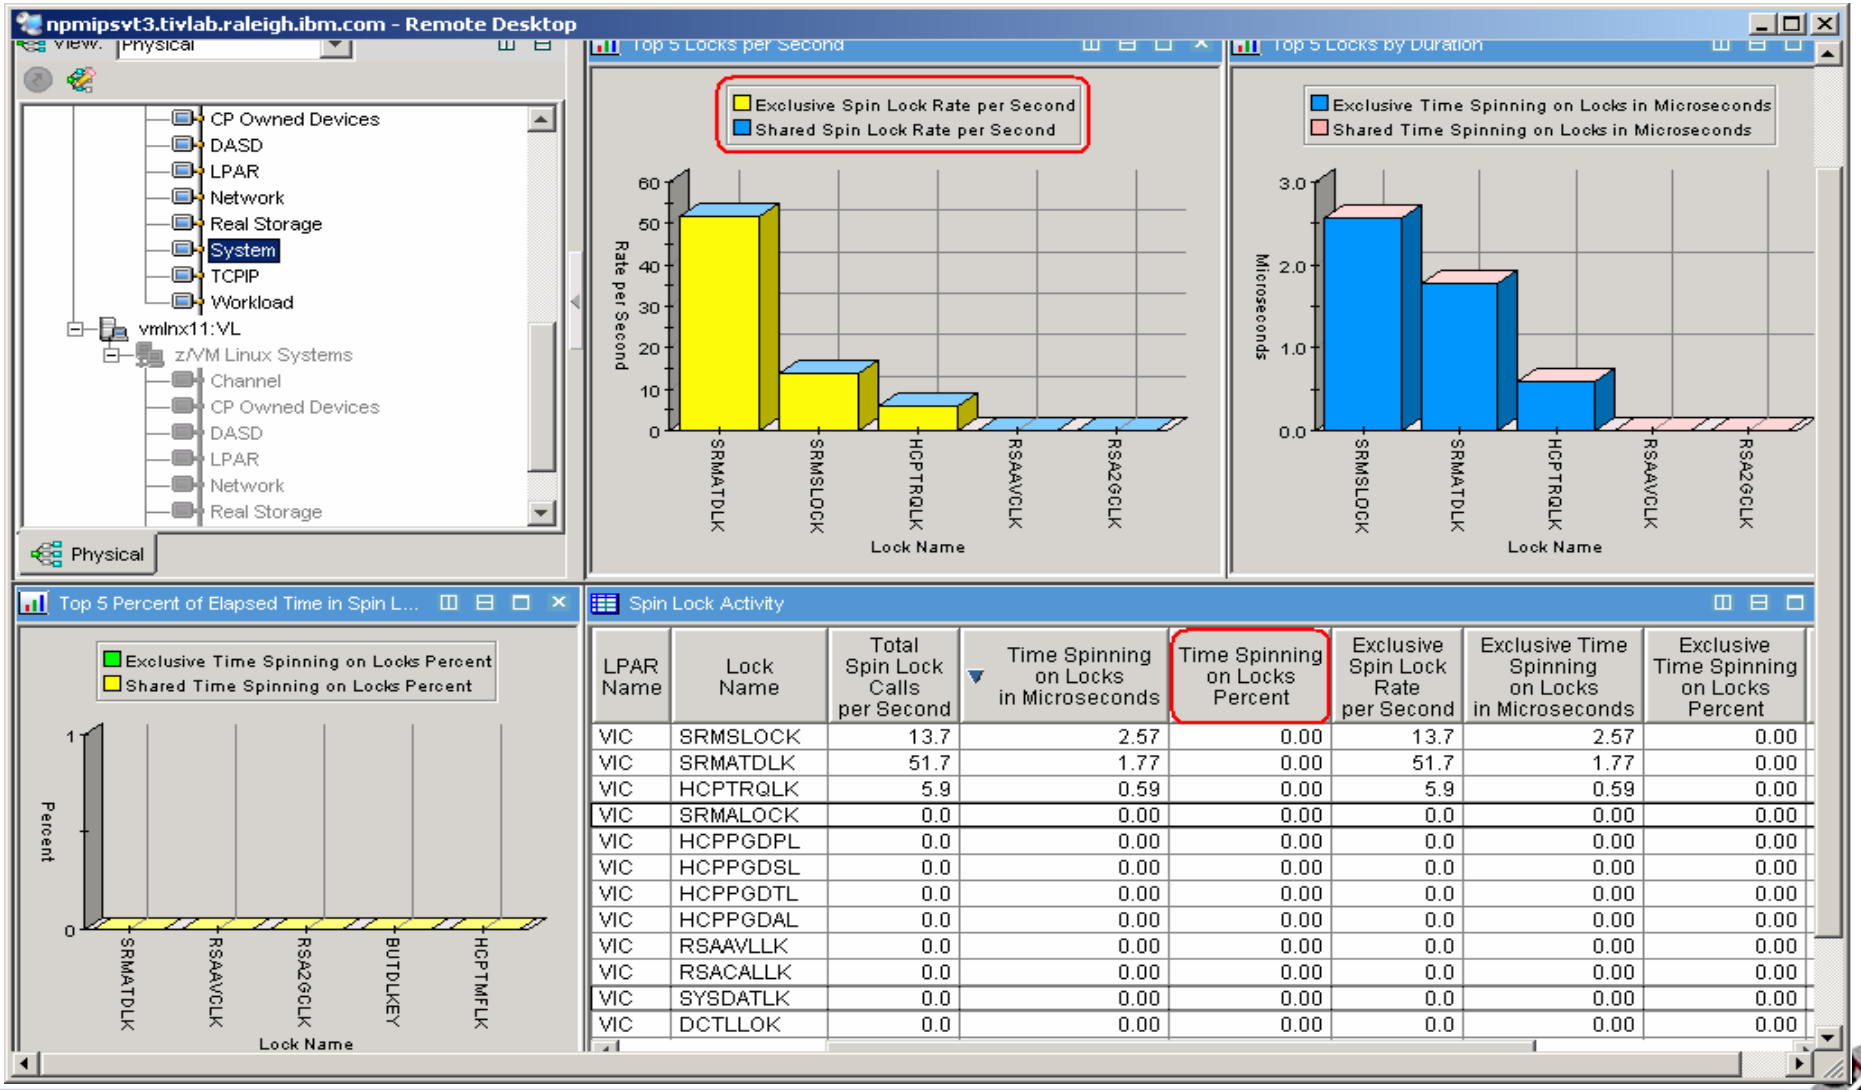

Is my Linux guest sized correctly?

- 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).
  - A good exercise is to compare Linux memory usage to z/VM working set size for the guest.

Need breakdown of memory use

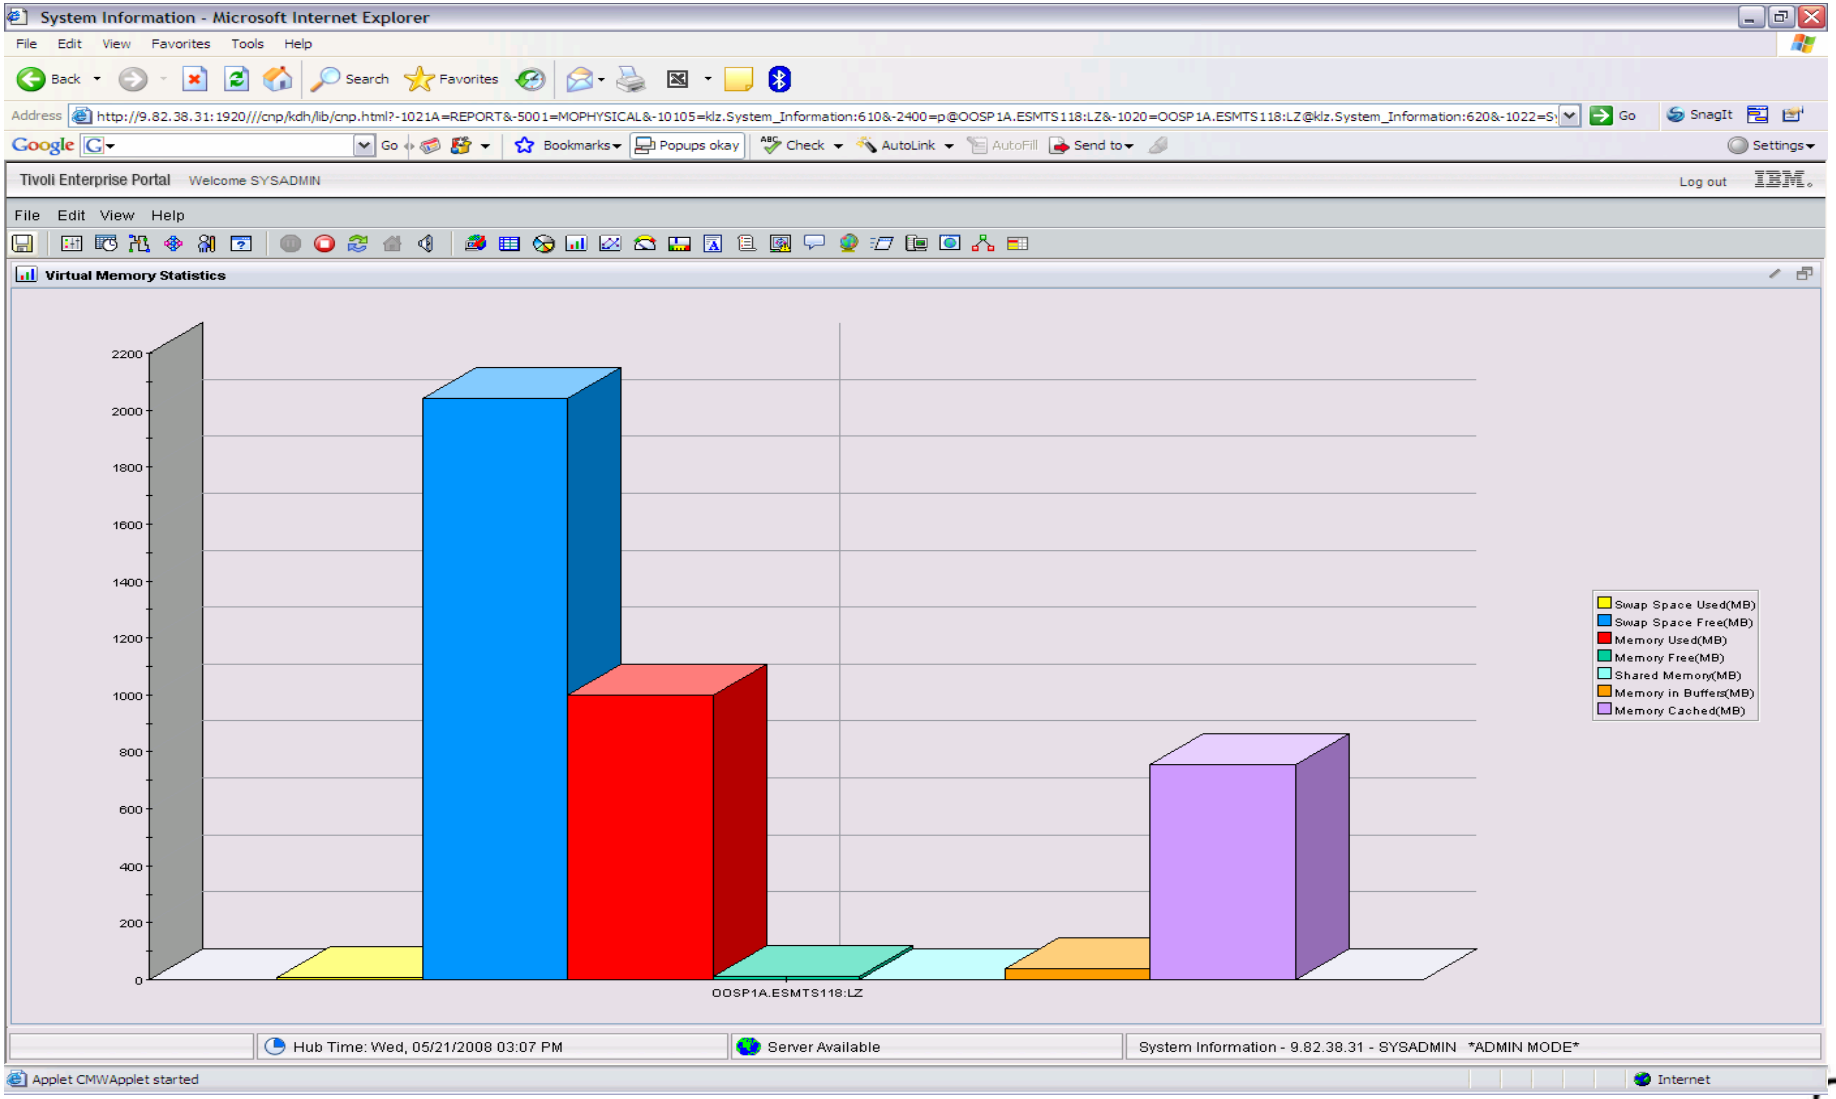

Working Set Size

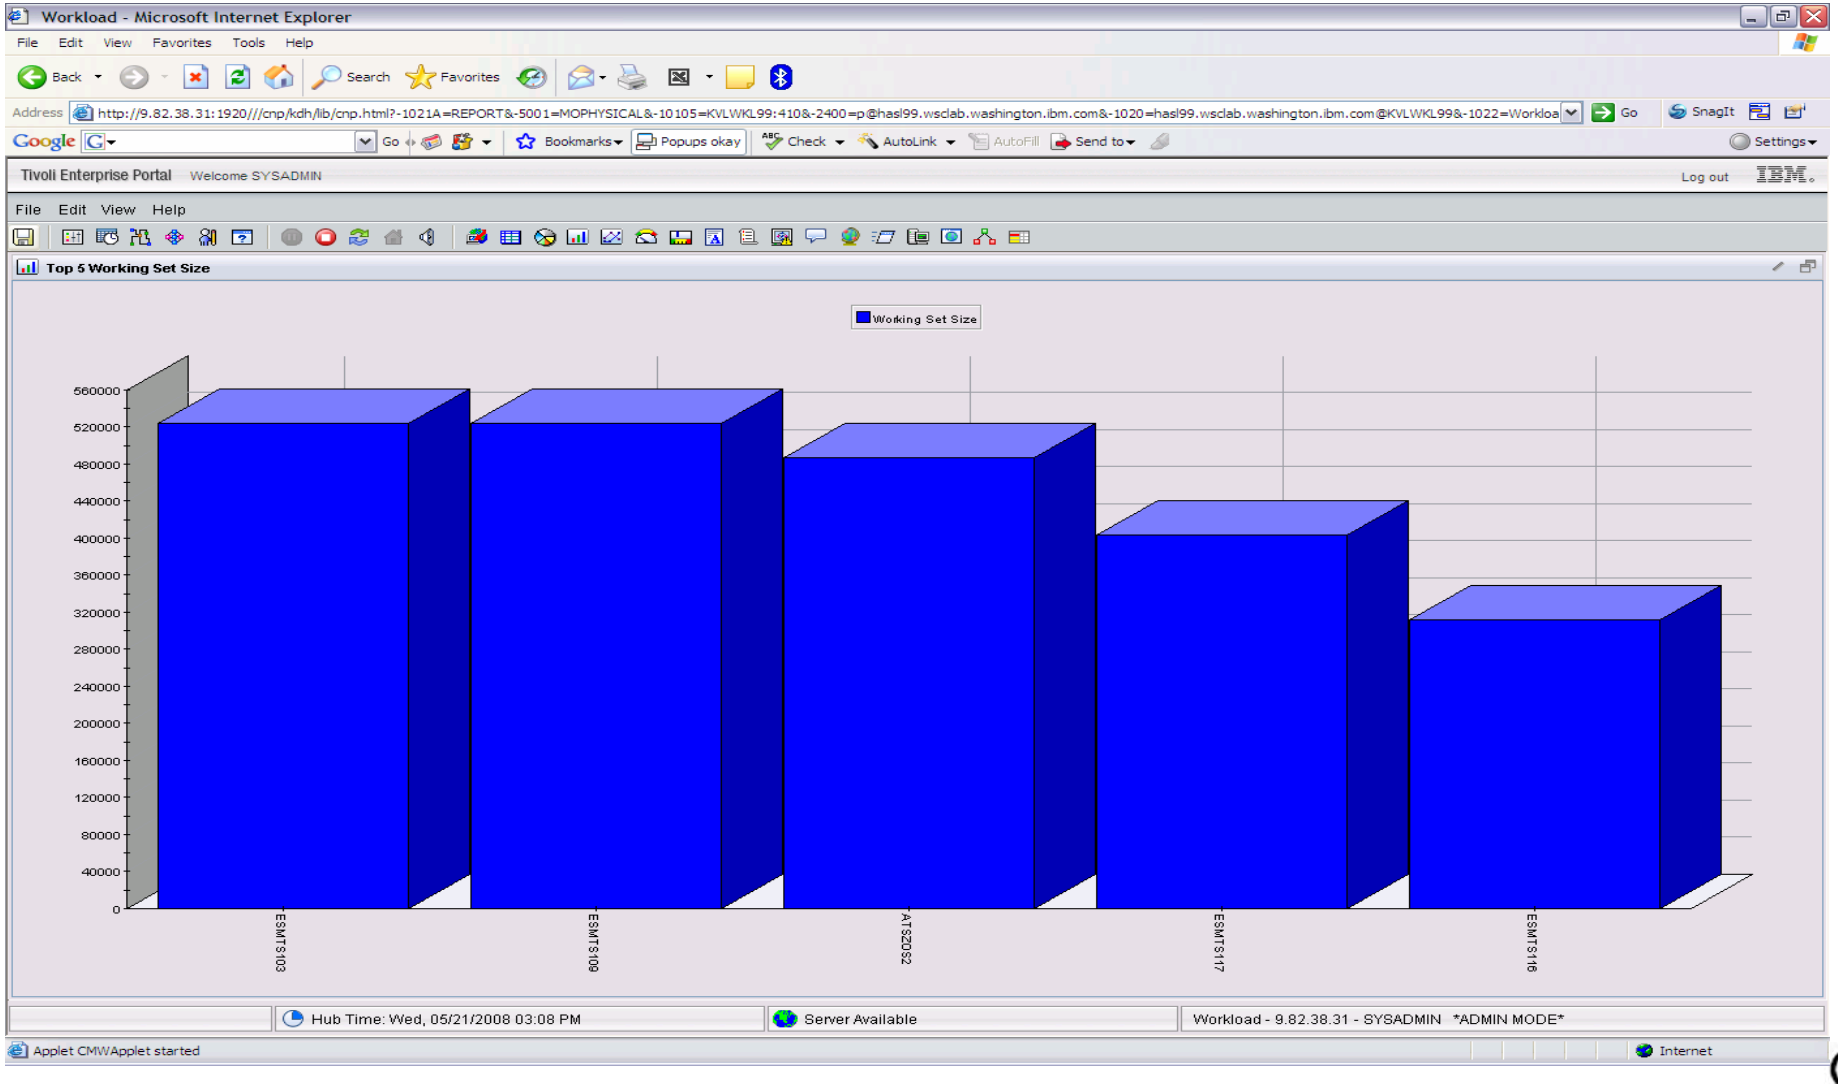

Page/Swap Attributes

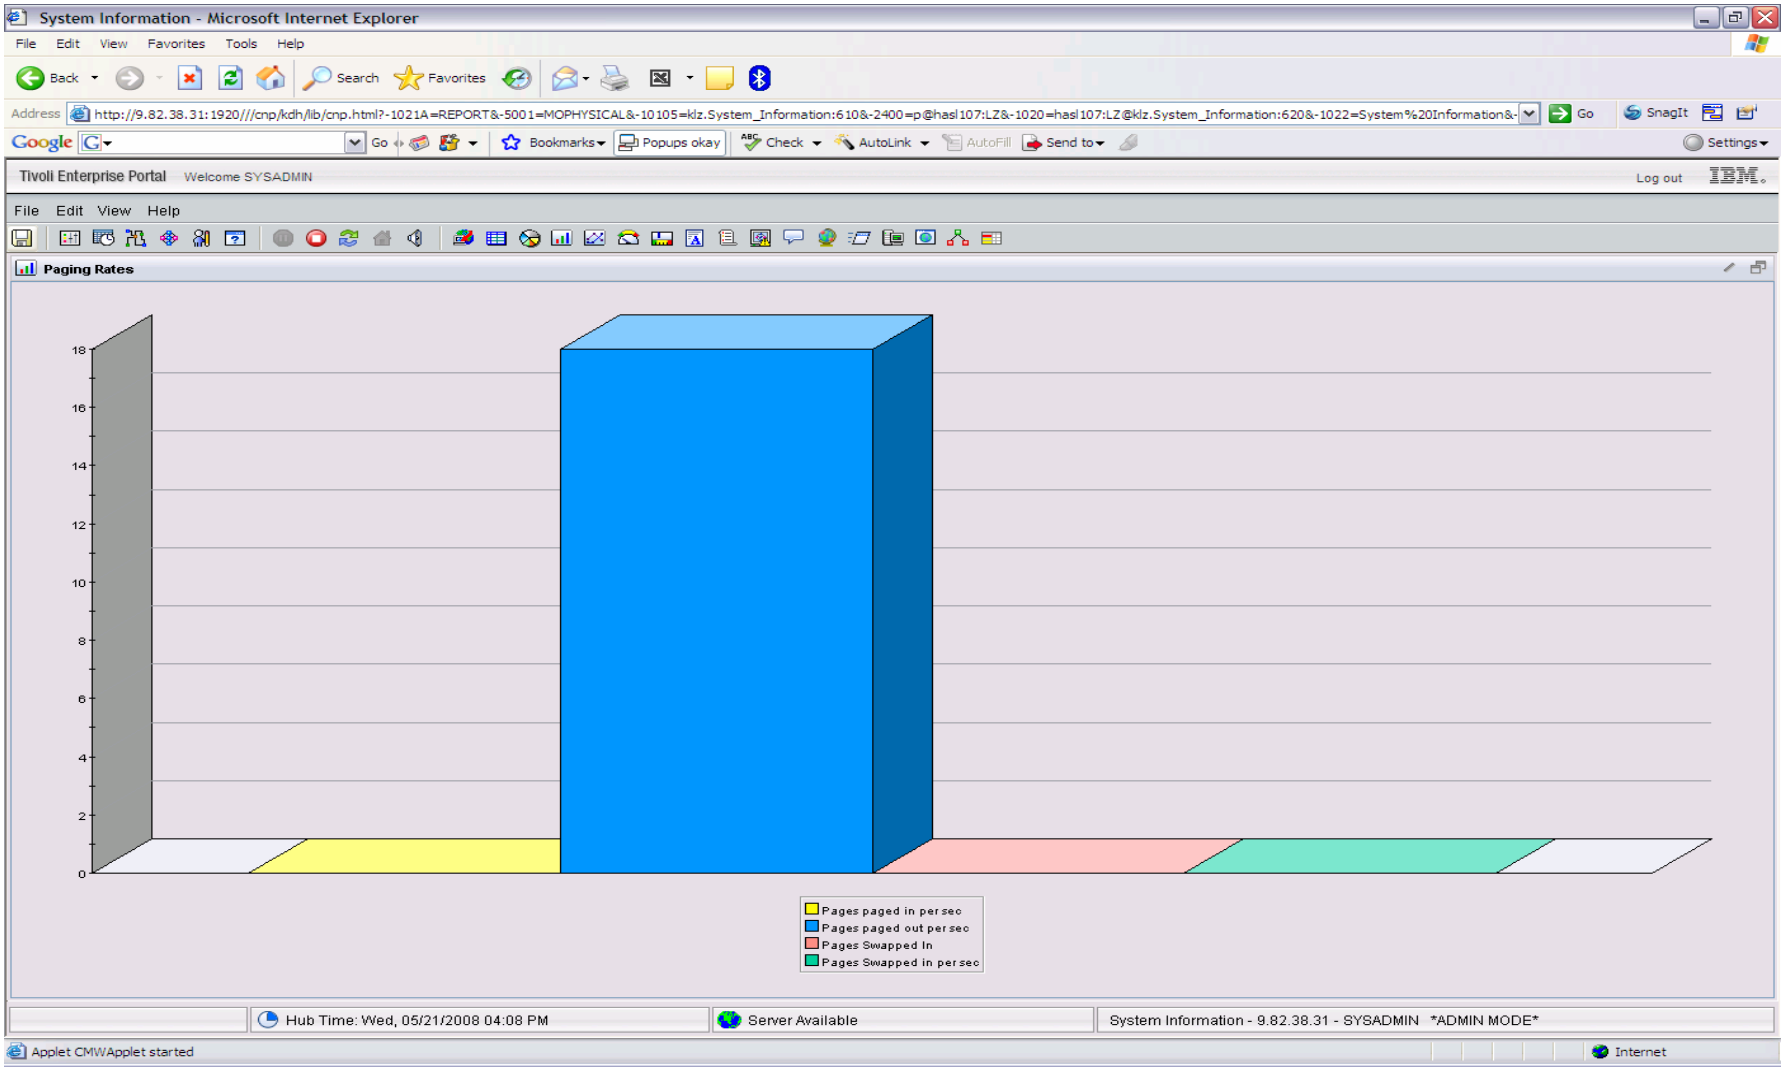

VDISK

- 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

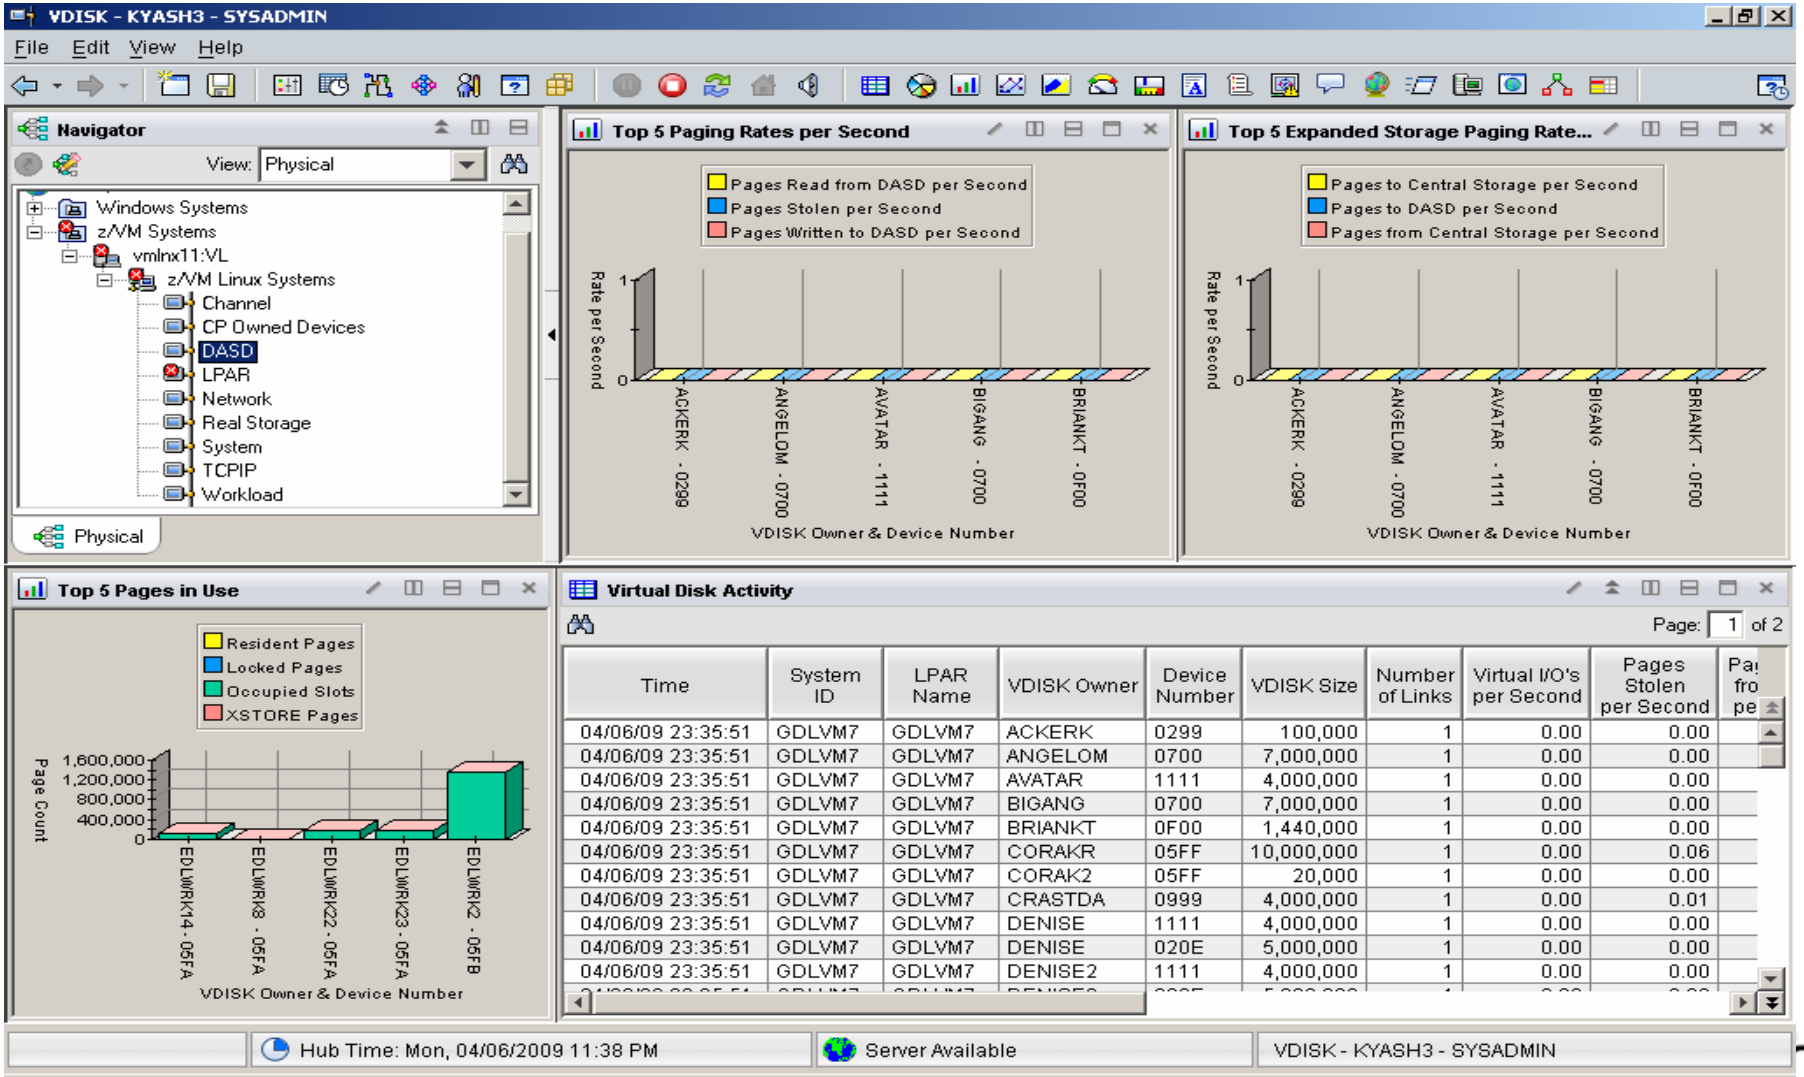

Memory Configuration

- Plan on a virtual to real (V:R) memory ratio in the range of 1.5:1 to 3:1.
- Recommend configuring some processor memory as 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 25% of memory configured as expanded:
  - Typically 2–4GB of expanded storage is sufficient, 1GB minimum.
  - 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.

OMEGAMON Memory Configuration

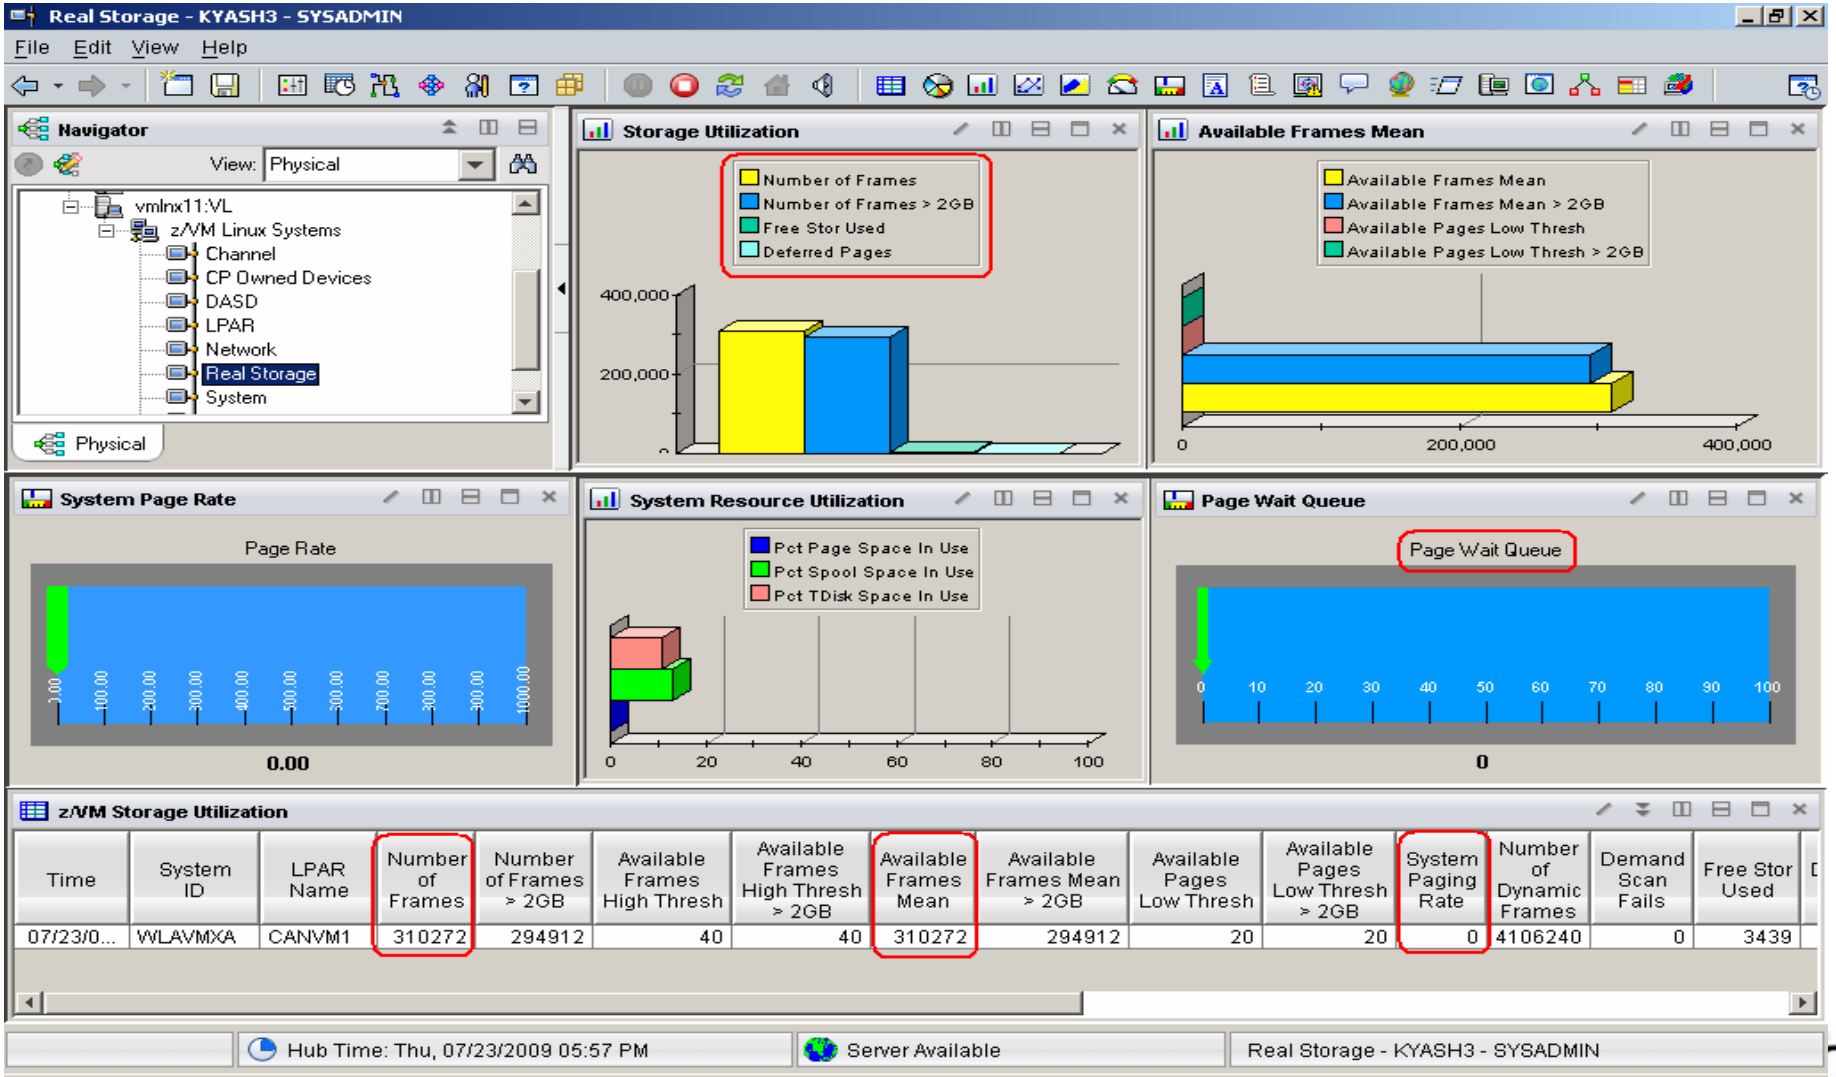

Paging Subsystem

- 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.

OMEGAMON CP Owned Devices – Paging Subsystem

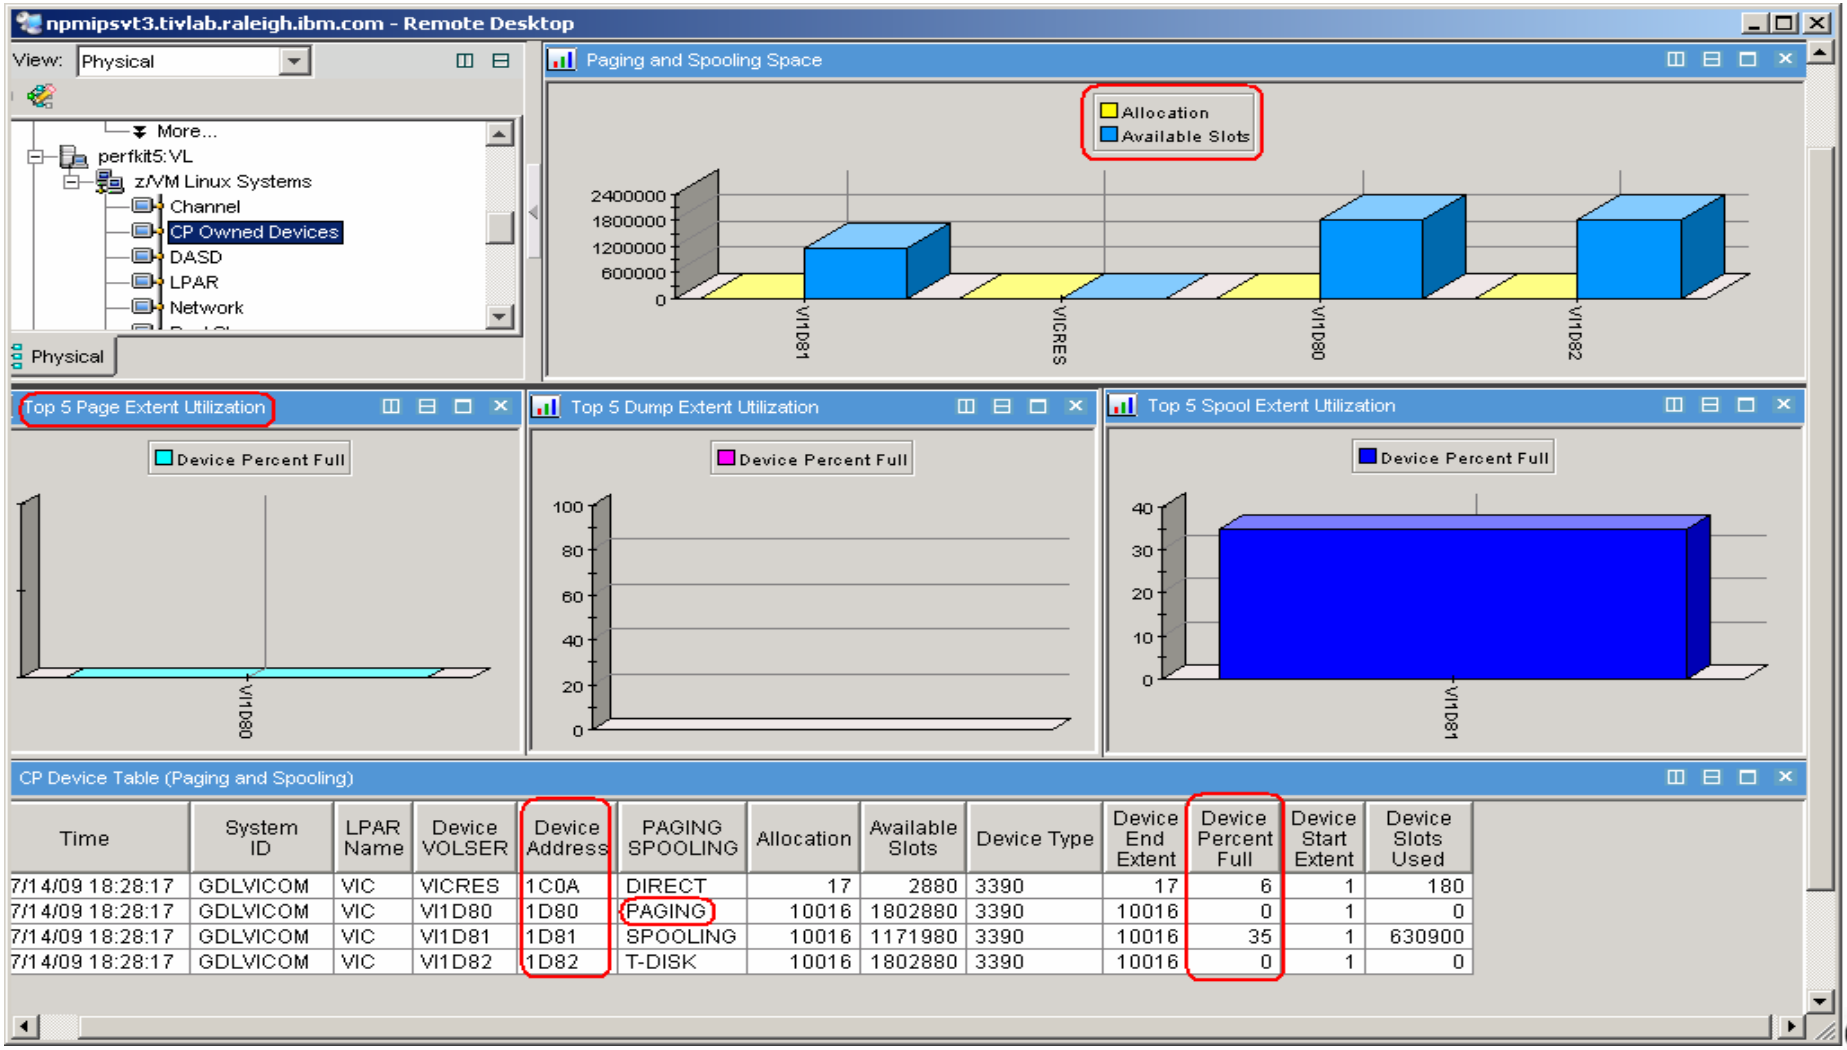

Time	System ID	LPAR Name	Device VOLSER	Device Address	PAGING SPOOLING	Allocation	Available Slots	Device Type	Device End Extent	Device Percent Full	Device Start Extent	Device Slots Used
7/14/09 18:28:17	GDLVICOM	VIC	VICRES	1C0A	DIRECT	17	2880	3390	17	6	1	180
7/14/09 18:28:17	GDLVICOM	VIC	VI1D80	1D80	PAGING	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

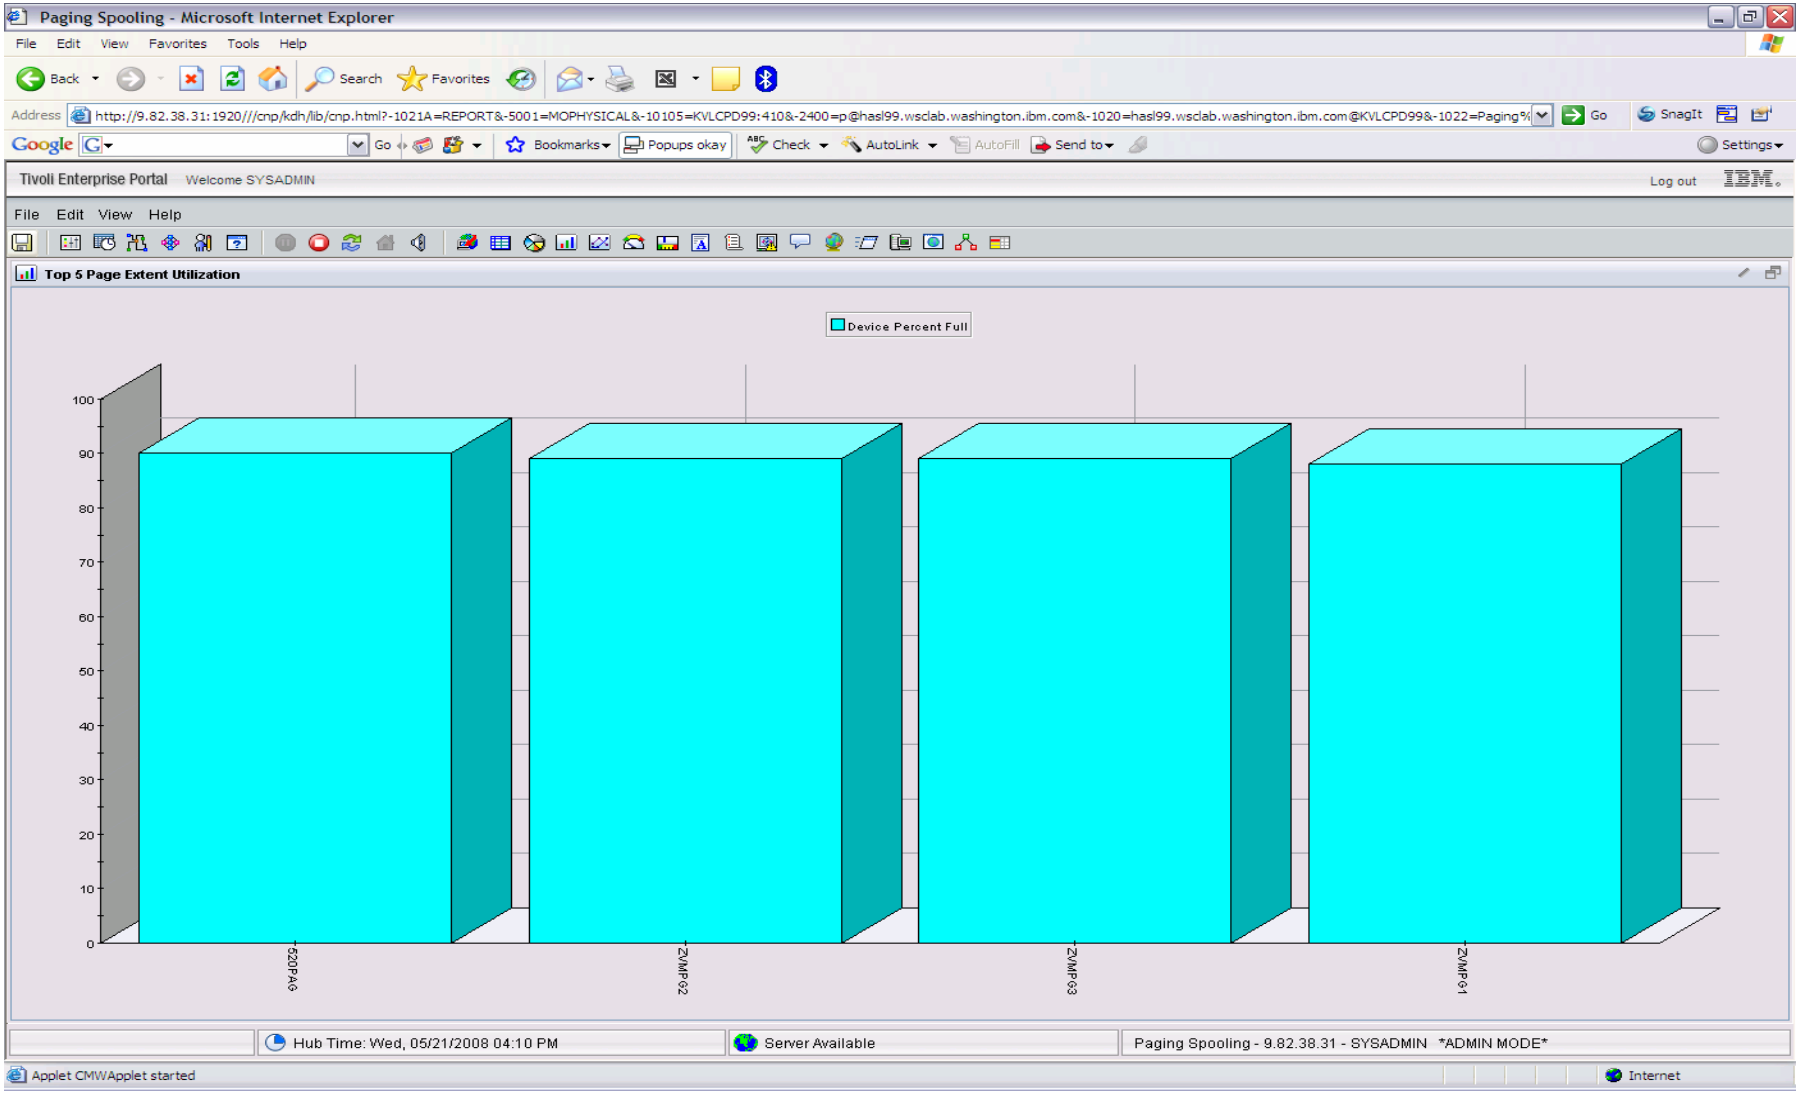

Minidisk Cache

- 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.
- 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 OMEGAMON XE and/or the Q MDC command.

OMEGAMON MDISK Cache Allocations

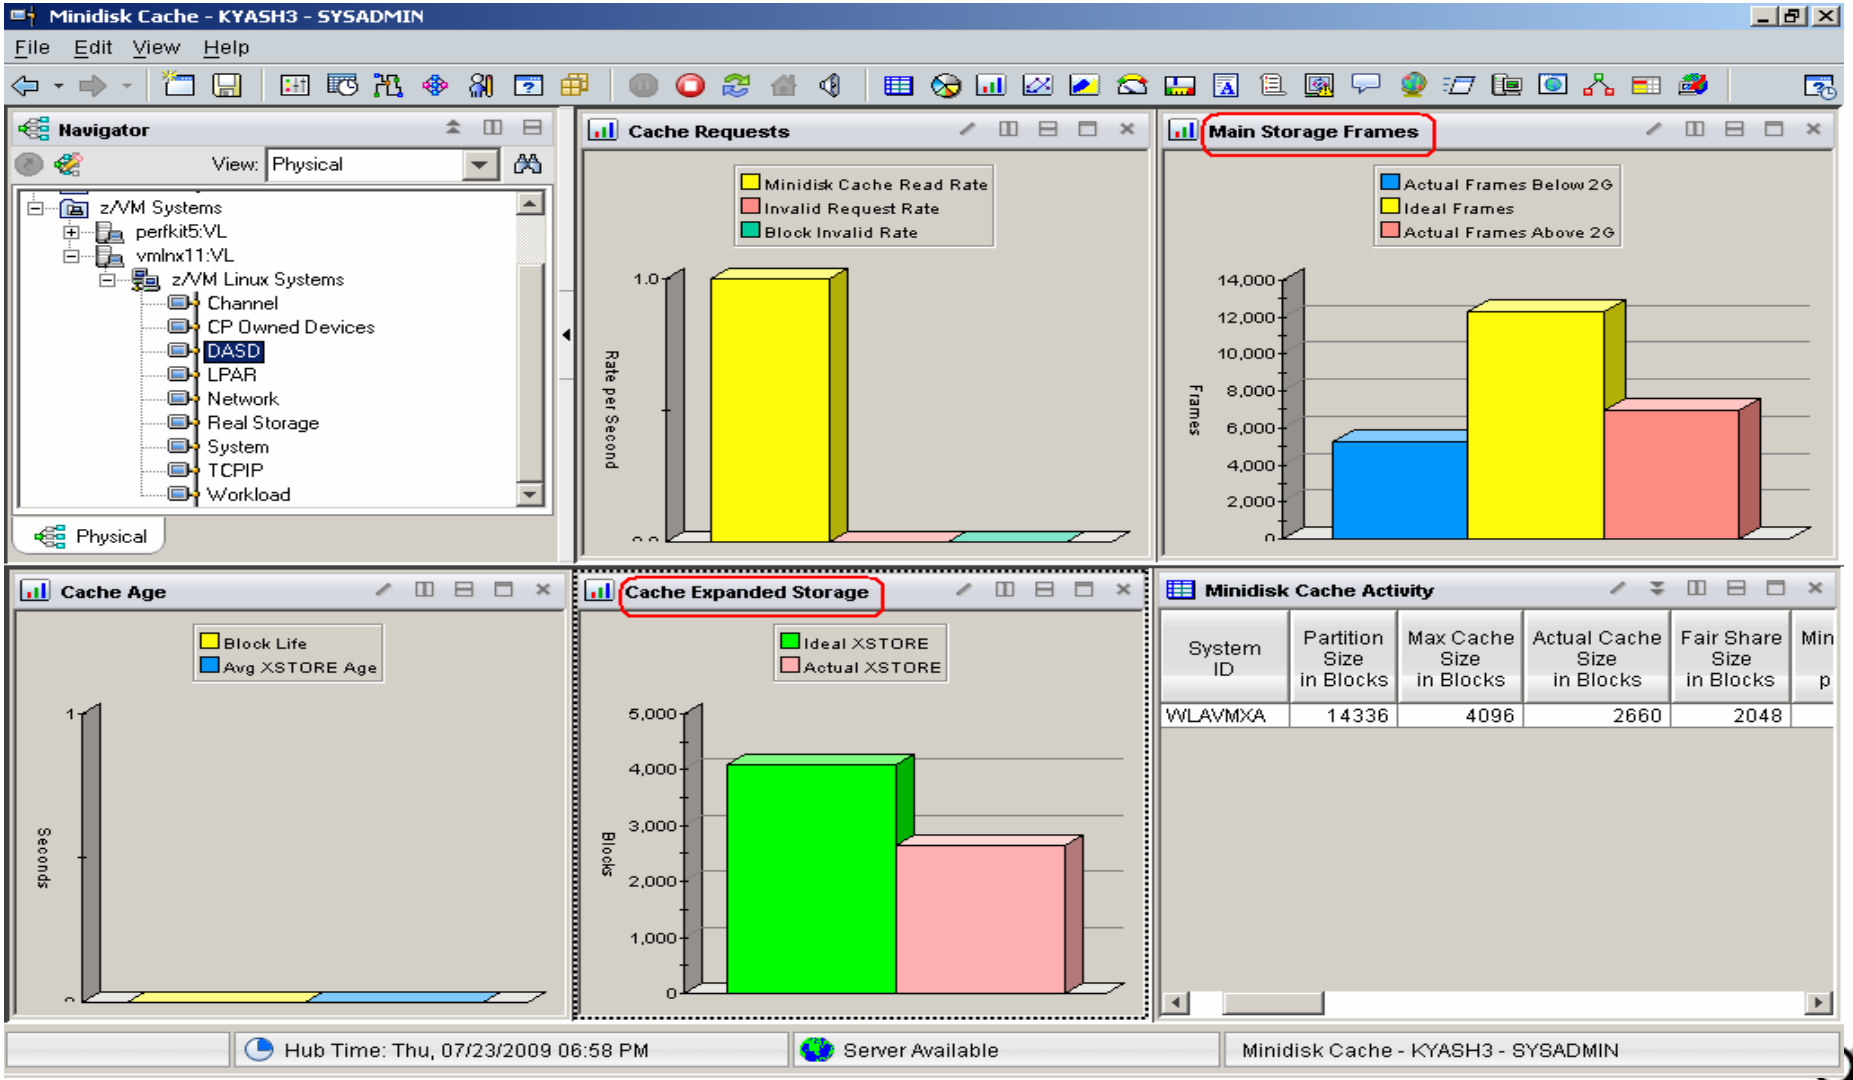

OMEGAMON MDISK Cache Allocations – p. 2

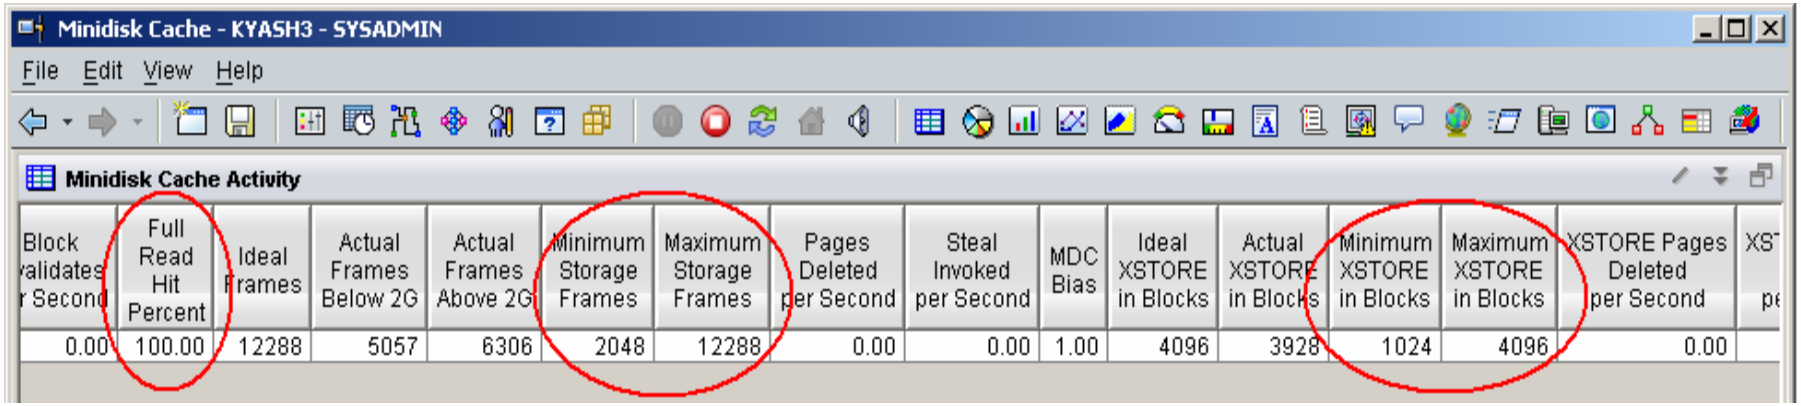

Block validates r Second	Full Read Hit Percent	Ideal Frames	Actual Frames Below 2G	Actual Frames Above 2G	Minimum Storage Frames	Maximum Storage Frames	Pages Deleted per Second	Steal Invoked per Second	MDC Bias	Ideal XSTORE in Blocks	Actual XSTORE in Blocks	Minimum XSTORE in Blocks	Maximum XSTORE in Blocks	XSTORE Pages Deleted per Second	XS 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

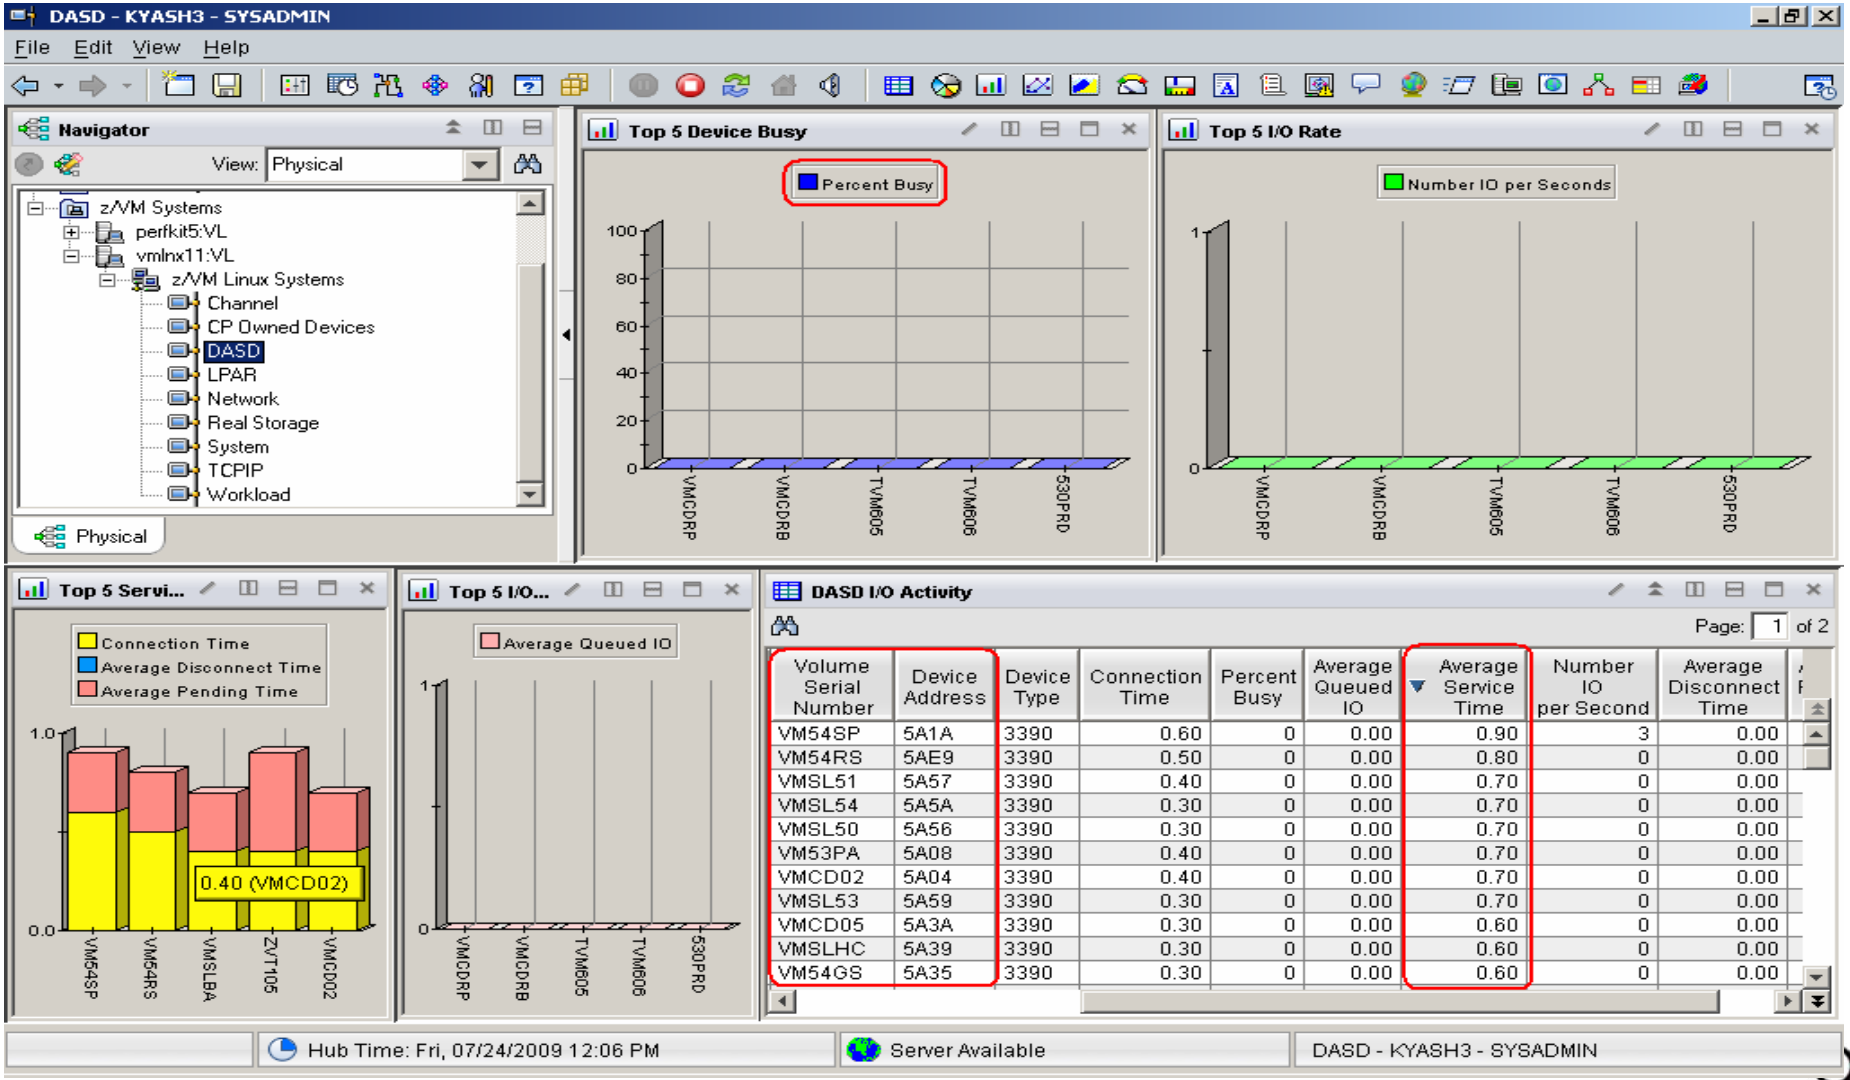

System Dump & Spool Space

- 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 Q ALLOC SPOOL command.
    - SFPURGER utility:
      - Rule based tool to clean up spool space.
      - Included in the no charge CMS Utilities Feature (CUF).

System Dump & Spool Space

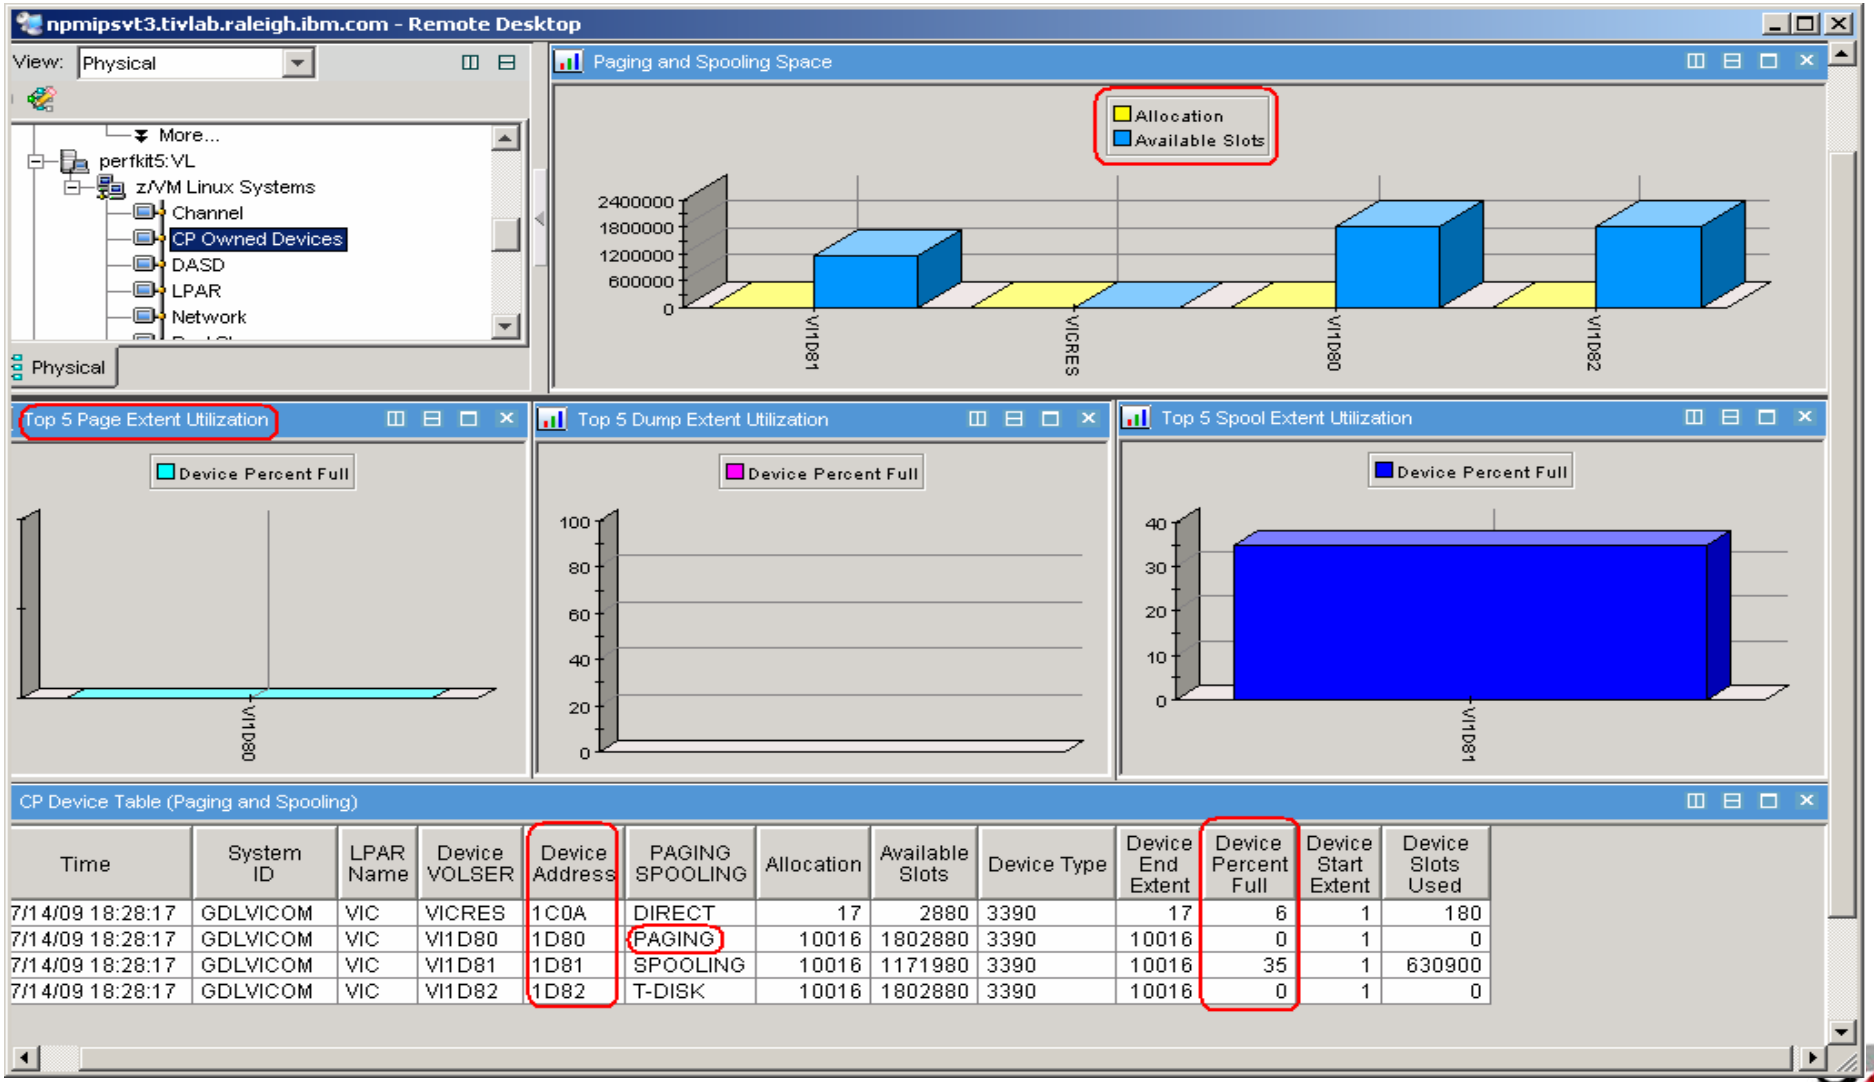

Time	System ID	LPAR Name	Device VOLSER	Device Address	PAGING SPOOLING	Allocation	Available Slots	Device Type	Device End Extent	Device Percent Full	Device Start Extent	Device Slots Used
7/14/09 18:28:17	GDLVICOM	VIC	VICRES	1C0A	DIRECT	17	2880	3390	17	6	1	180
7/14/09 18:28:17	GDLVICOM	VIC	VI1D80	1D80	PAGING	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

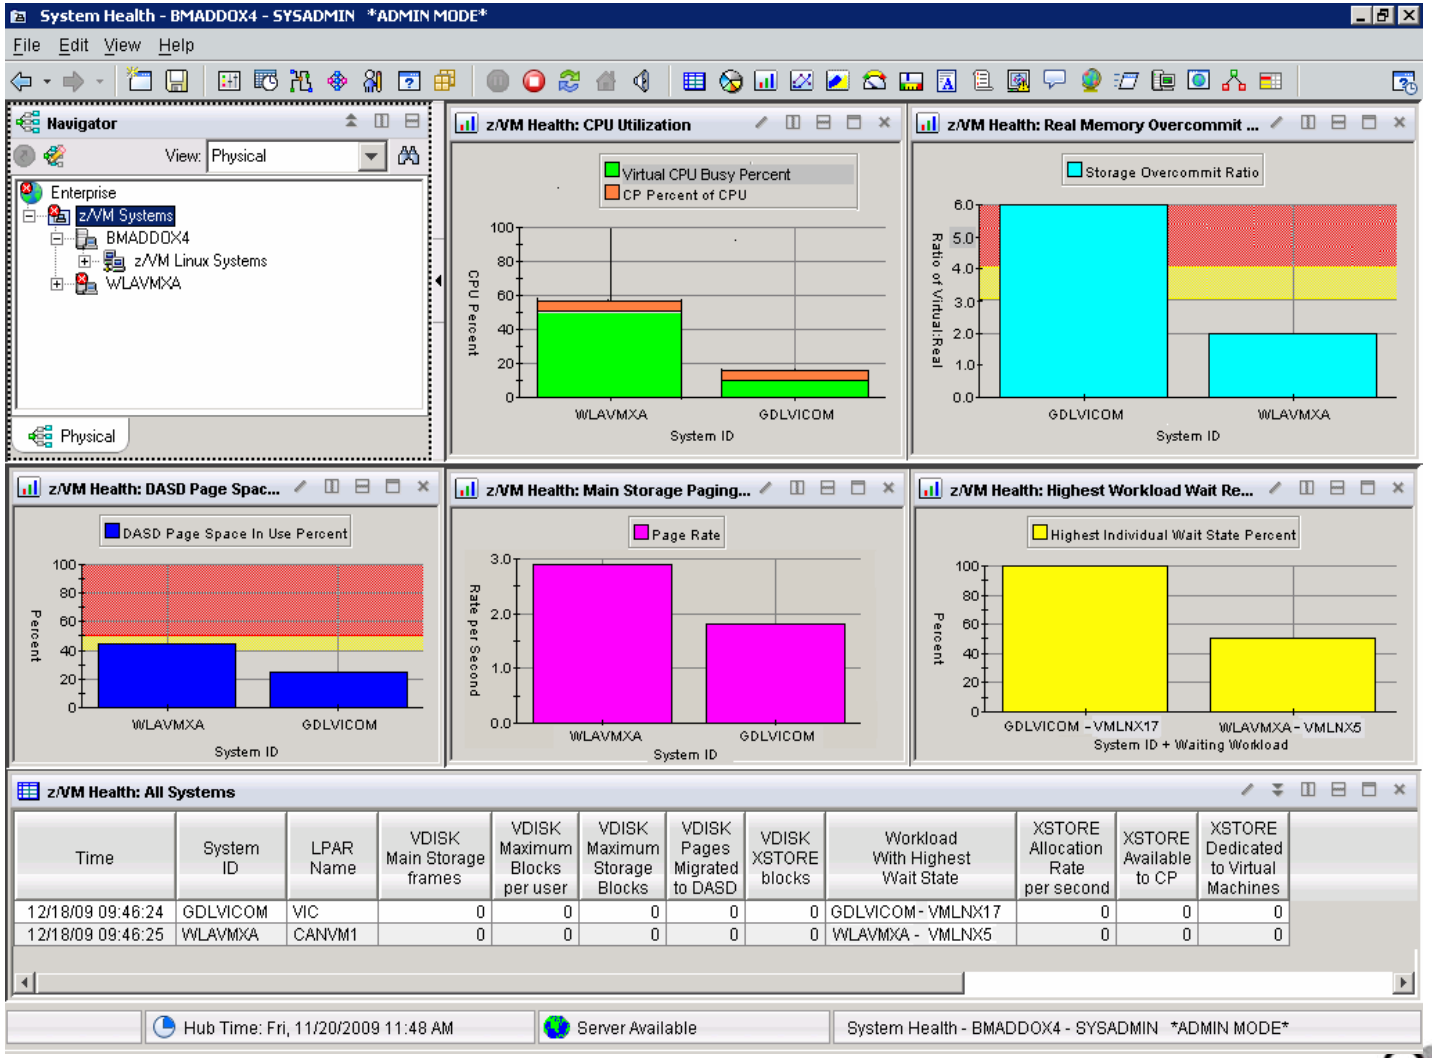

V4.1.2 IF 1: Resource Constraint Analysis (Waits)

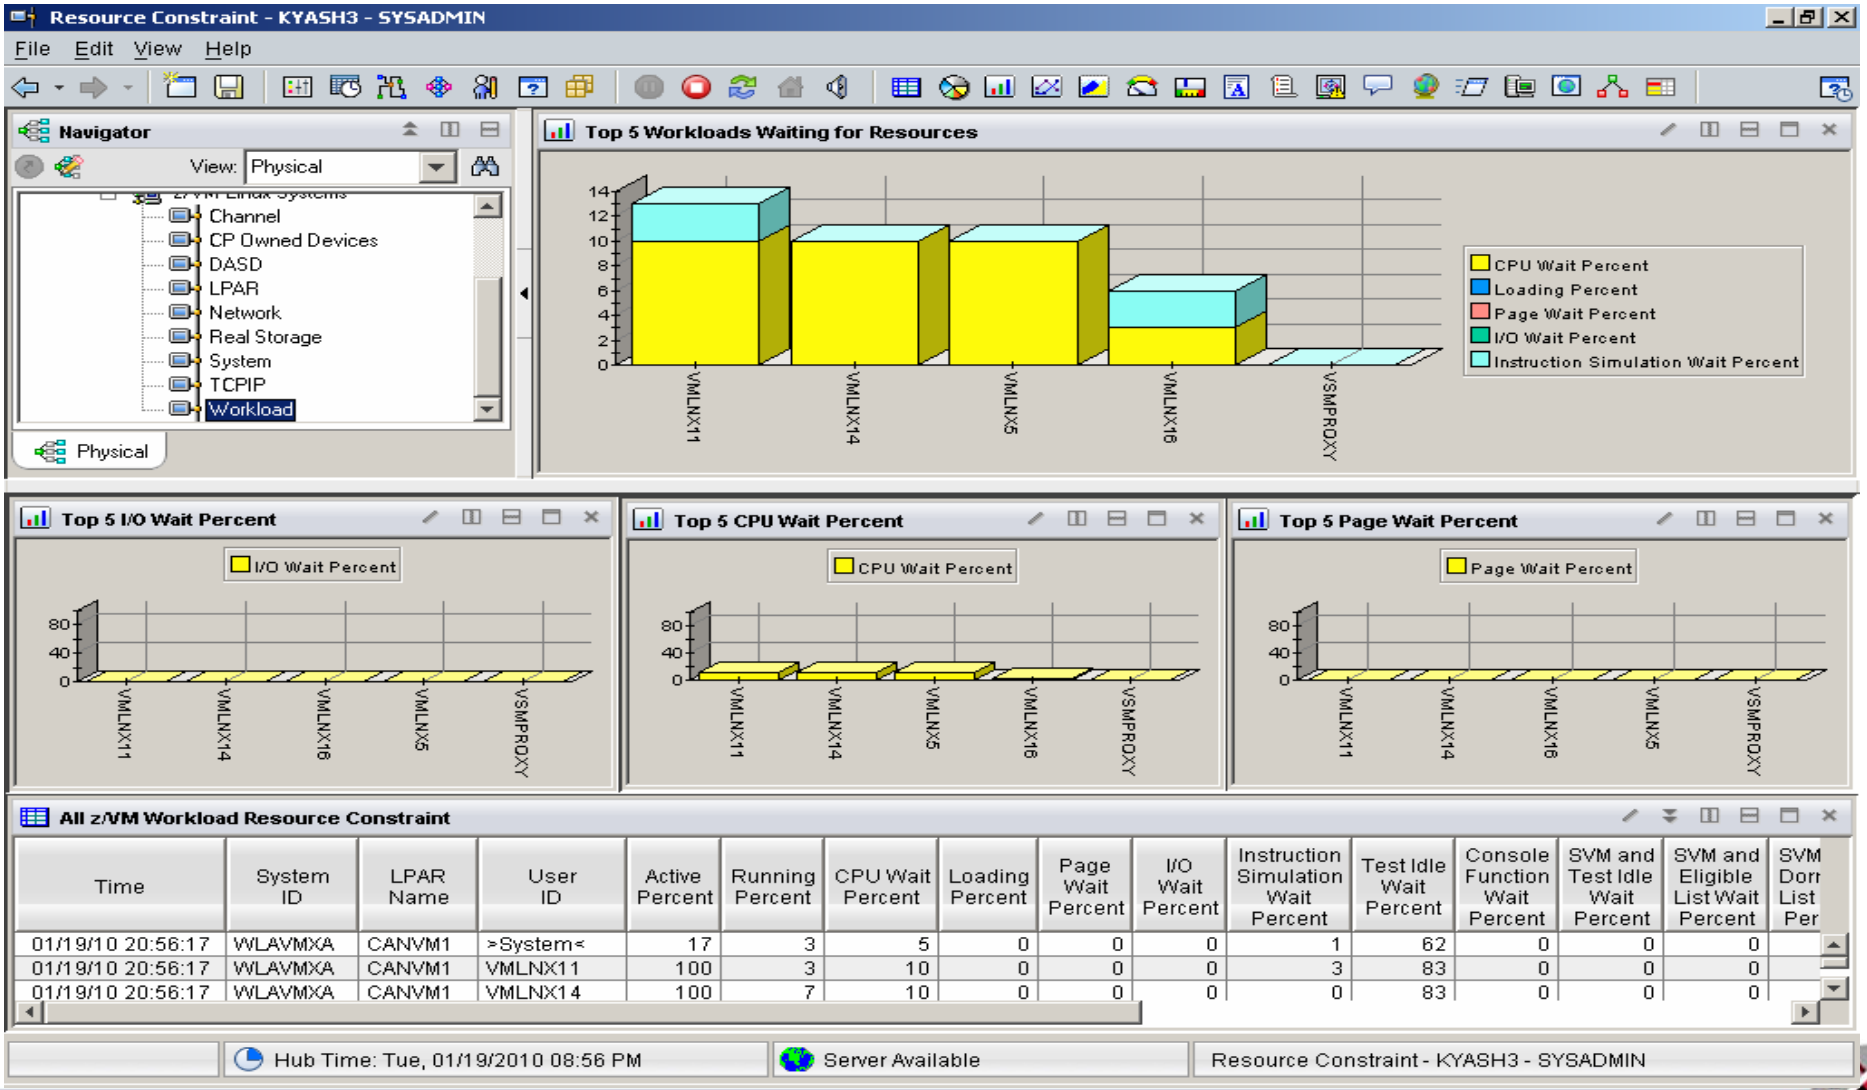

Do not ignore the hardware!

- 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.

- Logical Processors
  - LPAR recommendation – no greater than a 4:1 logical to real ratio.
  - z/VM 5.1 and z/VM 5.2 support up to 24 processors.
  - z/VM 5.3 and z/VM 5.4 support up to 32 processors.

LPAR Utilization Workspace

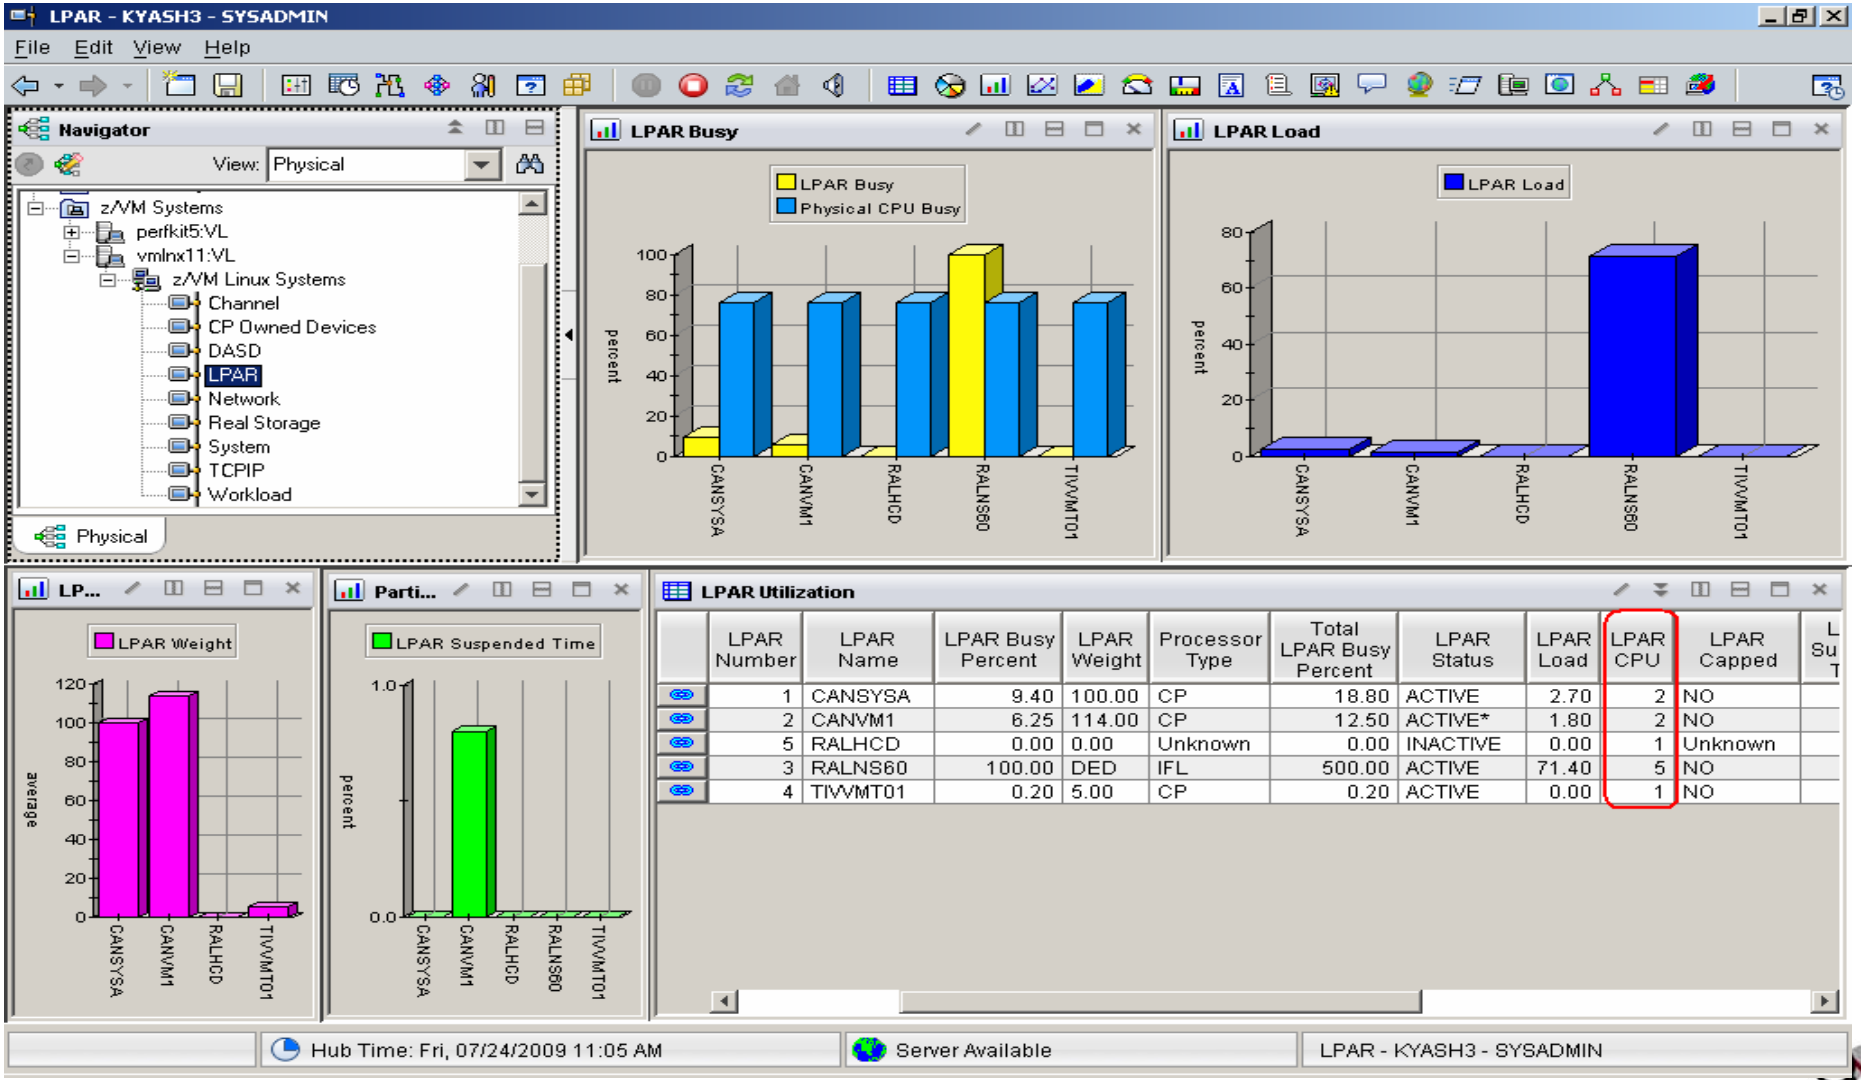

LPAR Number	LPAR Name	LPAR Busy Percent	LPAR Weight	Processor Type	Total LPAR Busy Percent	LPAR Status	LPAR Load	LPAR CPU	LPAR Capped
1	CANSYSA	9.40	100.00	CP	18.80	ACTIVE	2.70	2	NO
2	CANVM1	6.25	114.00	CP	12.50	ACTIVE*	1.80	2	NO
5	RALHCD	0.00	0.00	Unknown	0.00	INACTIVE	0.00	1	Unknown
3	RALNS60	100.00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

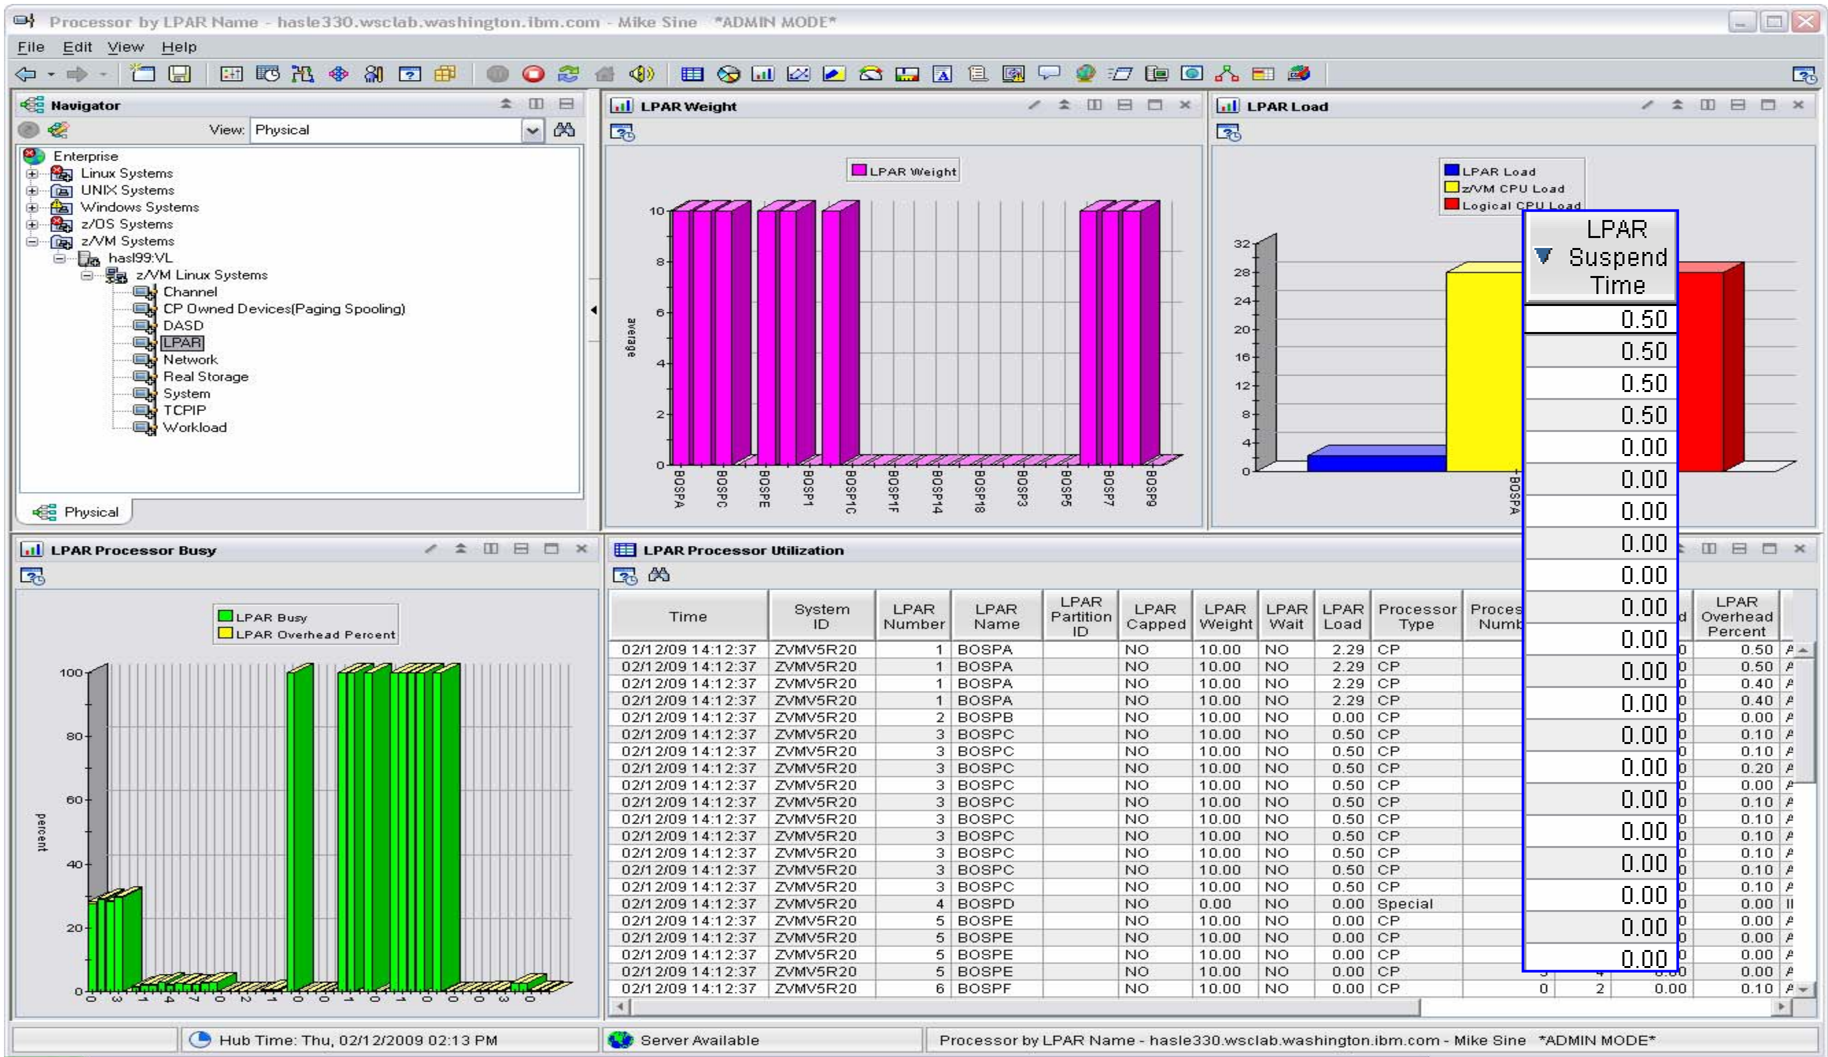

LPAR Utilization Workspace – Tabular View

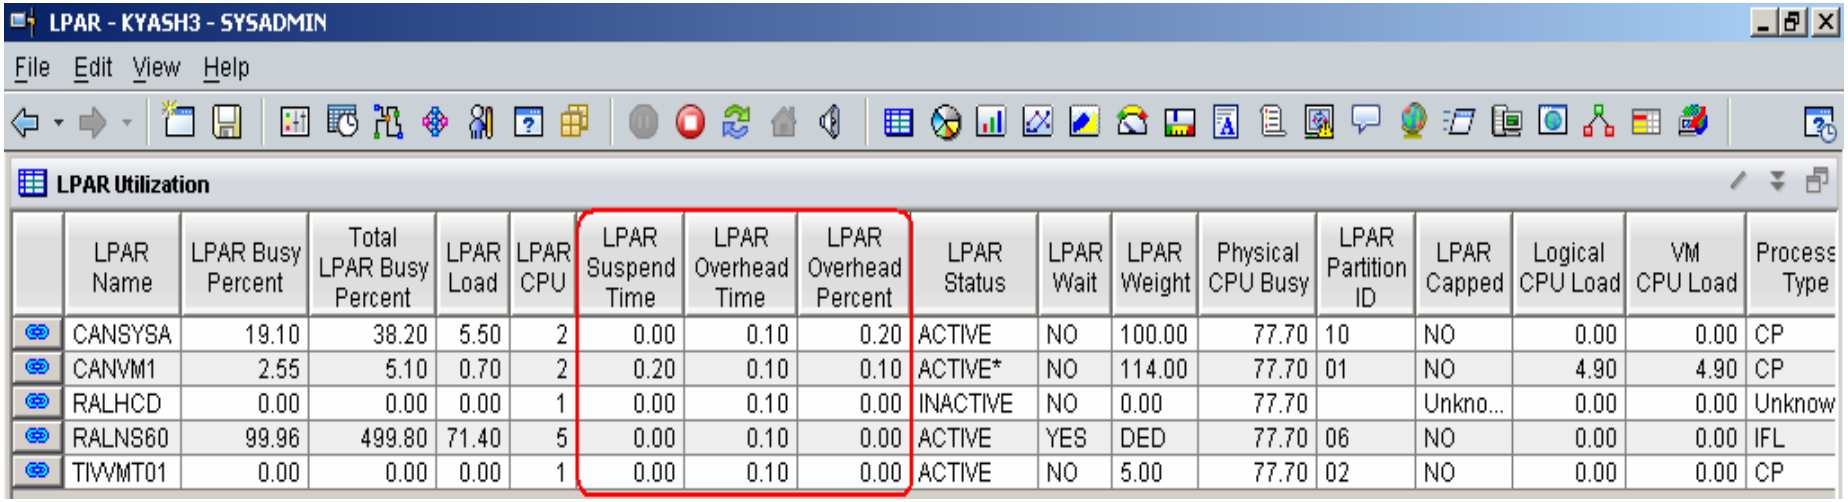

LPAR Name	LPAR Busy Percent	Total LPAR Busy Percent	LPAR Load	LPAR CPU	LPAR Suspend Time	LPAR Overhead Time	LPAR Overhead Percent	LPAR Status	LPAR Wait	LPAR Weight	Physical CPU Busy	LPAR Partition ID	LPAR Capped	Logical CPU Load	VM CPU Load	Process Type
CANSYSA	19.10	38.20	5.50	2	0.00	0.10	0.20	ACTIVE	NO	100.00	77.70	10	NO	0.00	0.00	CP
CANVM1	2.55	5.10	0.70	2	0.20	0.10	0.10	ACTIVE*	NO	114.00	77.70	01	NO	4.90	4.90	CP
RALHCD	0.00	0.00	0.00	1	0.00	0.10	0.00	INACTIVE	NO	0.00	77.70		Unkno...	0.00	0.00	Unknow
RALNS60	99.96	499.80	71.40	5	0.00	0.10	0.00	ACTIVE	YES	DED	77.70	06	NO	0.00	0.00	IFL
TIVVMT01	0.00	0.00	0.00	1	0.00	0.10	0.00	ACTIVE	NO	5.00	77.70	02	NO	0.00	0.00	CP

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.

Persistent Historical Views

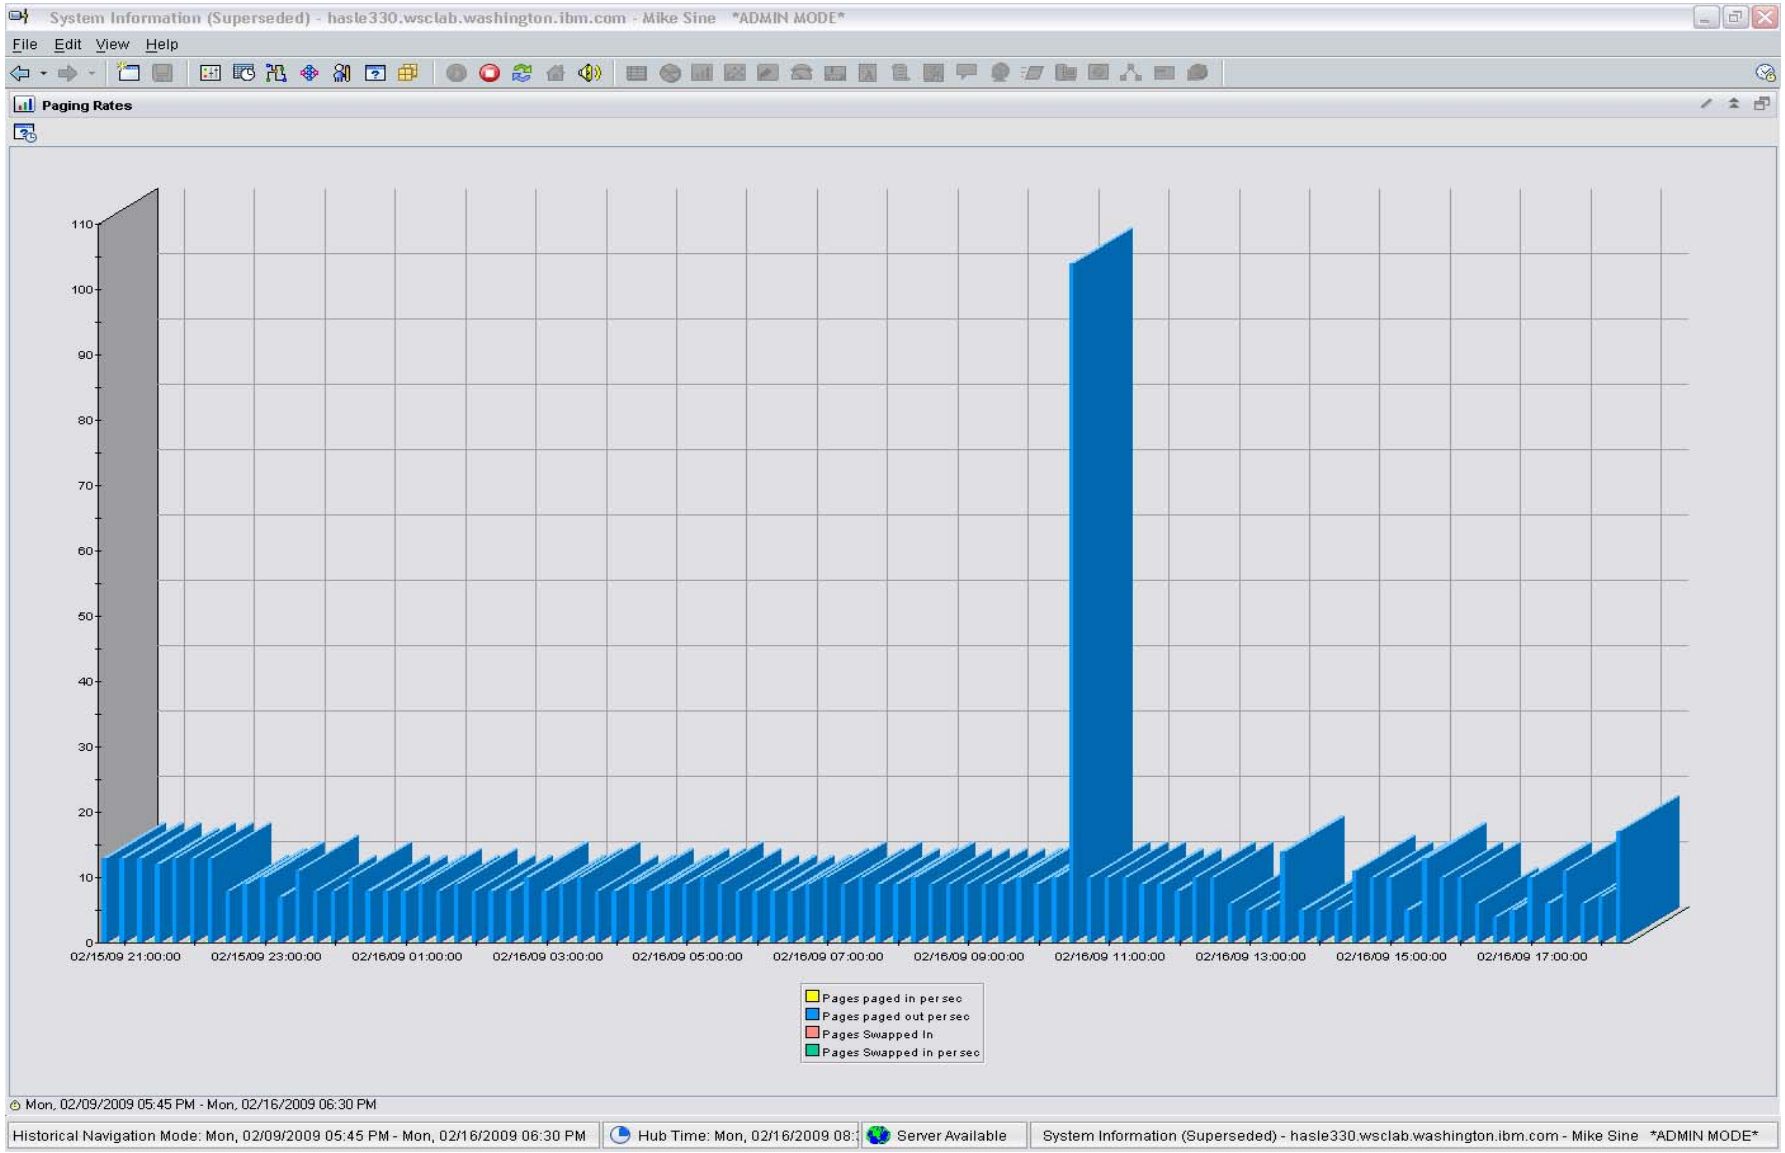

Persistent Historical Views

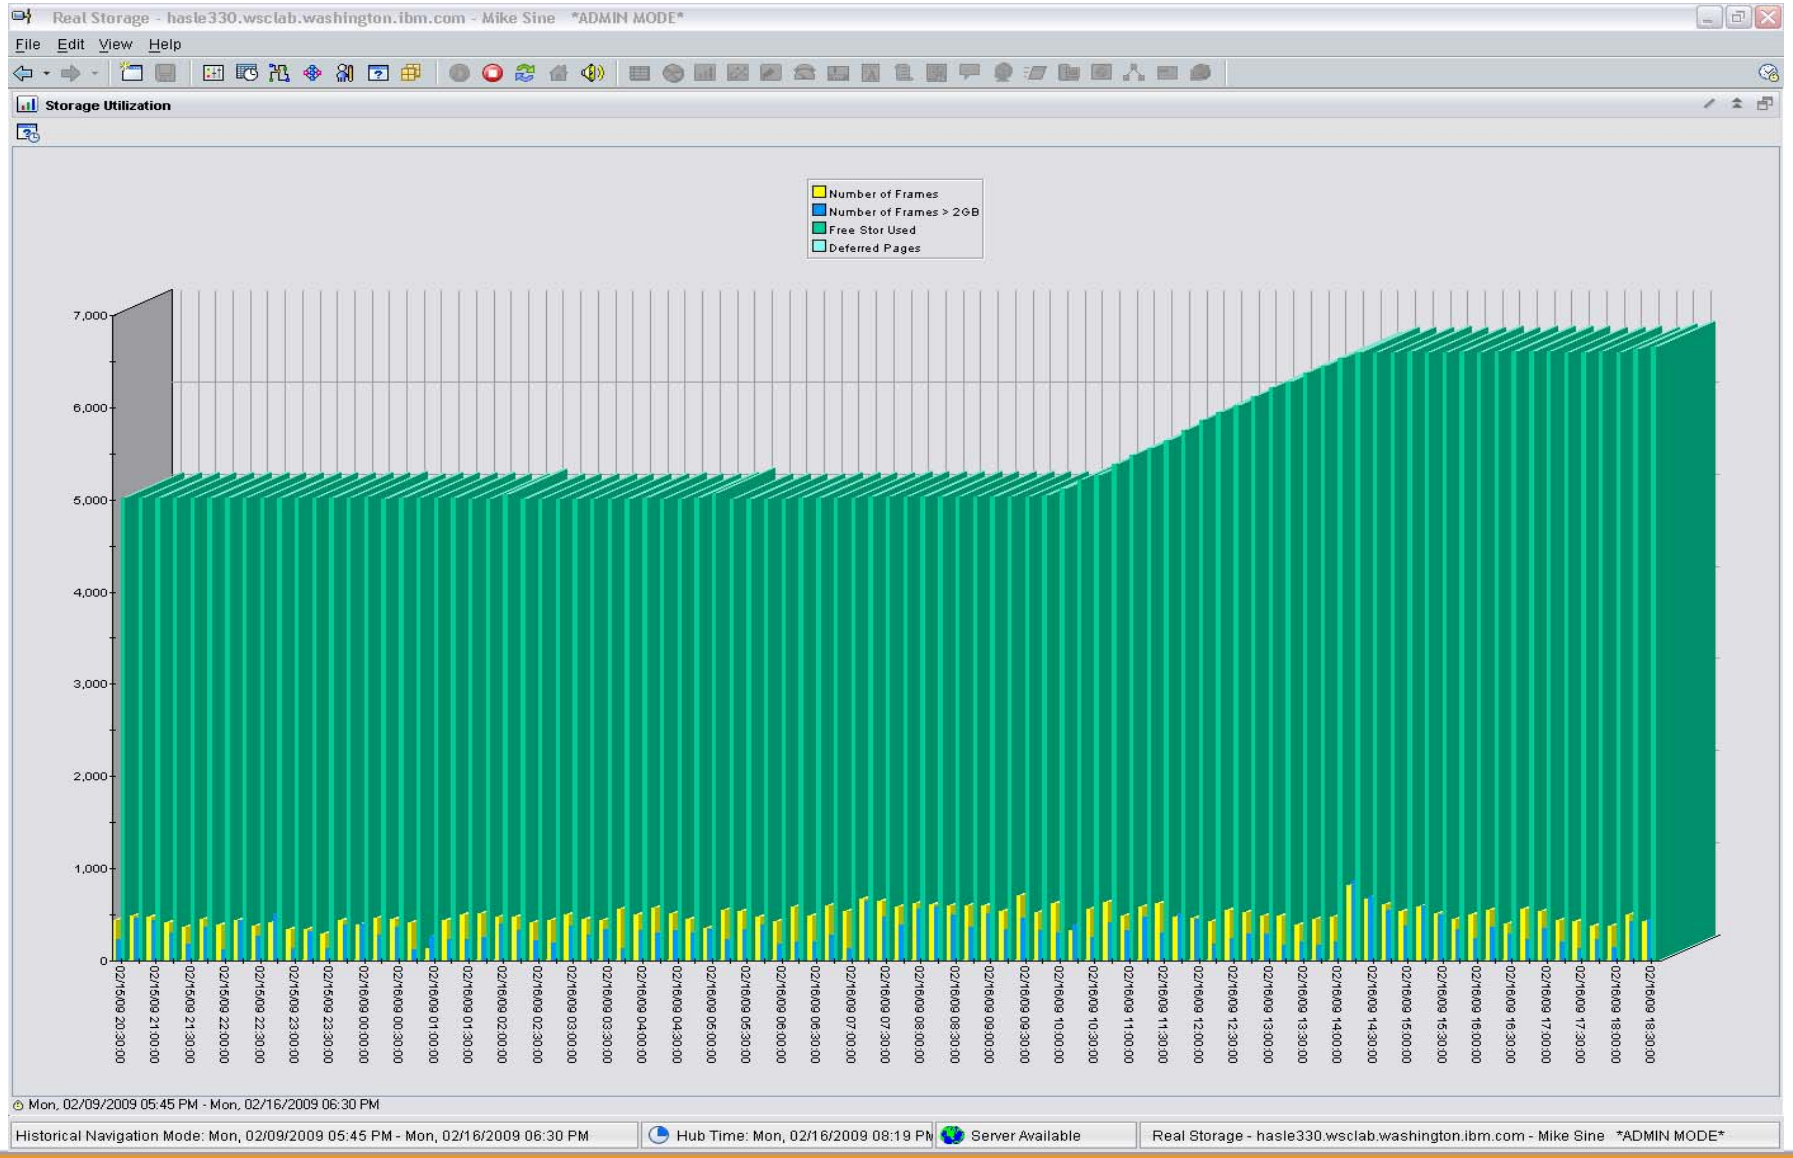

New Tivoli Common Reporting (TCR)

- 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.
  - Built on top of open source reporting tool called: BIRT.
  - 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

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

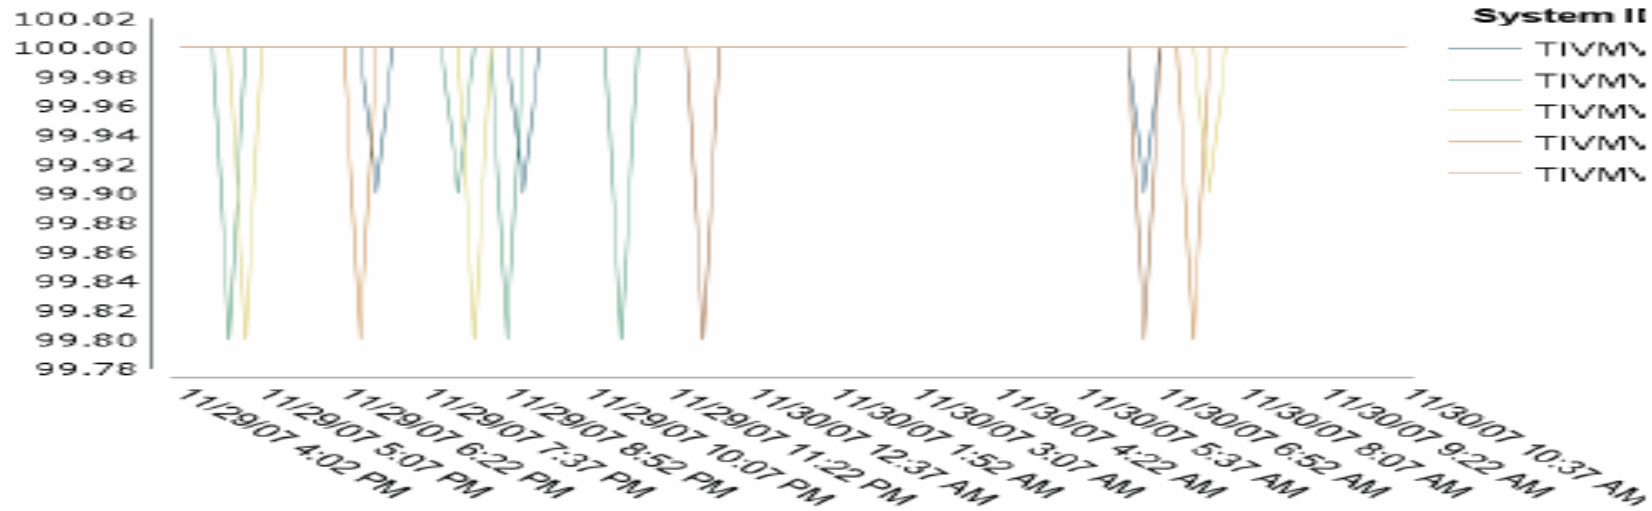

Available Summarization Time Periods:

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
- Why IBM
- Bringing it all together

What differentiates the IBM solution from the competition

- End to End Management and Seamless integration with other Tivoli Monitoring products through the Tivoli Enterprise Portal
  - Other vendors have multiple inconsistent user interfaces and cannot provide an end to end view spanning cross platform applications.
  - Some vendors are silo oriented and do not have the breadth and depth to manage the applications that are running on z/VM and Linux
  - If you are considering WebSphere, SAP, Oracle Financials, UDB,Oracle DB, etc for this platform only IBM has a fully integrated suite of monitoring tools across distributed and zSeries environments. Other vendors tend to only have consolidated alert consoles, with minimal launch in context capabilities
- Full suite of z/VM and Linux performance and management tools

IBM Management Portfolio for z/VM and Linux on z

IBM System z Virtualization Infrastructure
- IBM System z hardware (including LPAR hypervisor)
- IBM z/VM Version 5

Monitoring for Virtualization Infrastructure
- z/VM Virtual Machine Resource Manager (included with z/VM)
- IBM z/VM Performance Toolkit for VM (z/VM priced feature)
- IBM Director
- IBM Tivoli OMEGAMON XE on z/VM and Linux
- IBM Tivoli Monitoring
- IBM Tivoli Composite Application Manager for SOA
- IBM Tivoli Usage and Accounting Manager

Automation for Virtualization Infrastructure
- IBM Operations Manager for z/VM
- IBM Tivoli Enterprise Console
- IBM Tivoli Workload Scheduler

Provisioning Management
- IBM z/VM DirMaint (z/VM priced feature)
- z/VM Center task of IBM Director
- IBM Tivoli Provisioning Manager

Resiliency Management
- IBM Tivoli System Automation for Multiplatforms

Application Layer Management
- IBM Tivoli Application Dependency Discovery Manager
- IBM Tivoli OMEGAMON XE for Messaging
- IBM Tivoli Composite Application Manager for Response Time
- IBM Tivoli Composite Application Manager for Web Resources
- IBM Tivoli Composite Application Manager for Transactions
- IBM Tivoli License Compliance Manager

Extended Infrastructure Management *(Security)*
- IBM z/VM RACF Security Server (z/VM priced feature)
- IBM Tivoli zSecure
- IBM Tivoli Access Manager for e-business
- IBM Tivoli Access Manager for OS
- IBM Tivoli Federated Identity Manager
- IBM Tivoli Identity Manager
- IBM Directory Server
- IBM Directory Integrator
- IBM Tivoli Risk Manager

Extended Infrastructure Management *(Storage)*
- IBM SAN Volume Controller (SVC)
- IBM Tivoli Storage Manager
- IBM TotalStorage Productivity Center
- IBM Backup and Restore Manager for z/VM
- IBM Tape Manager for z/VM
- IBM Archive Manager for z/VM

Extended Infrastructure Management *(Network)*
- IBM z/VM RSCS (z/VM priced feature)

Business Services Management
- IBM Tivoli Business Service Manager
- IBM Tivoli Service Request Manager
- IBM Change and Configuration Management Database (CCMDB)

For specific releases, refer to Tivoli Platform Support Matrix at: **ibm.com**/software/sysmgmt/products/support/Tivoli_Supported_Platforms.html

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

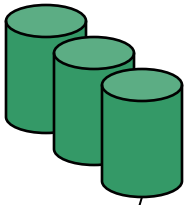

Monitor spool usage

Console monitoring

Service Virtual Machine being monitored

Take action

Operations Manager for z/VM

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

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

AGENDA

- Introduction
- Monitoring requirements
  - Virtual Linux and z/VM performance considerations
  - Don't forget the hardware
  - Integration from hardware – systems – applications
    Persistent historical views
- Why IBM
- Bringing 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

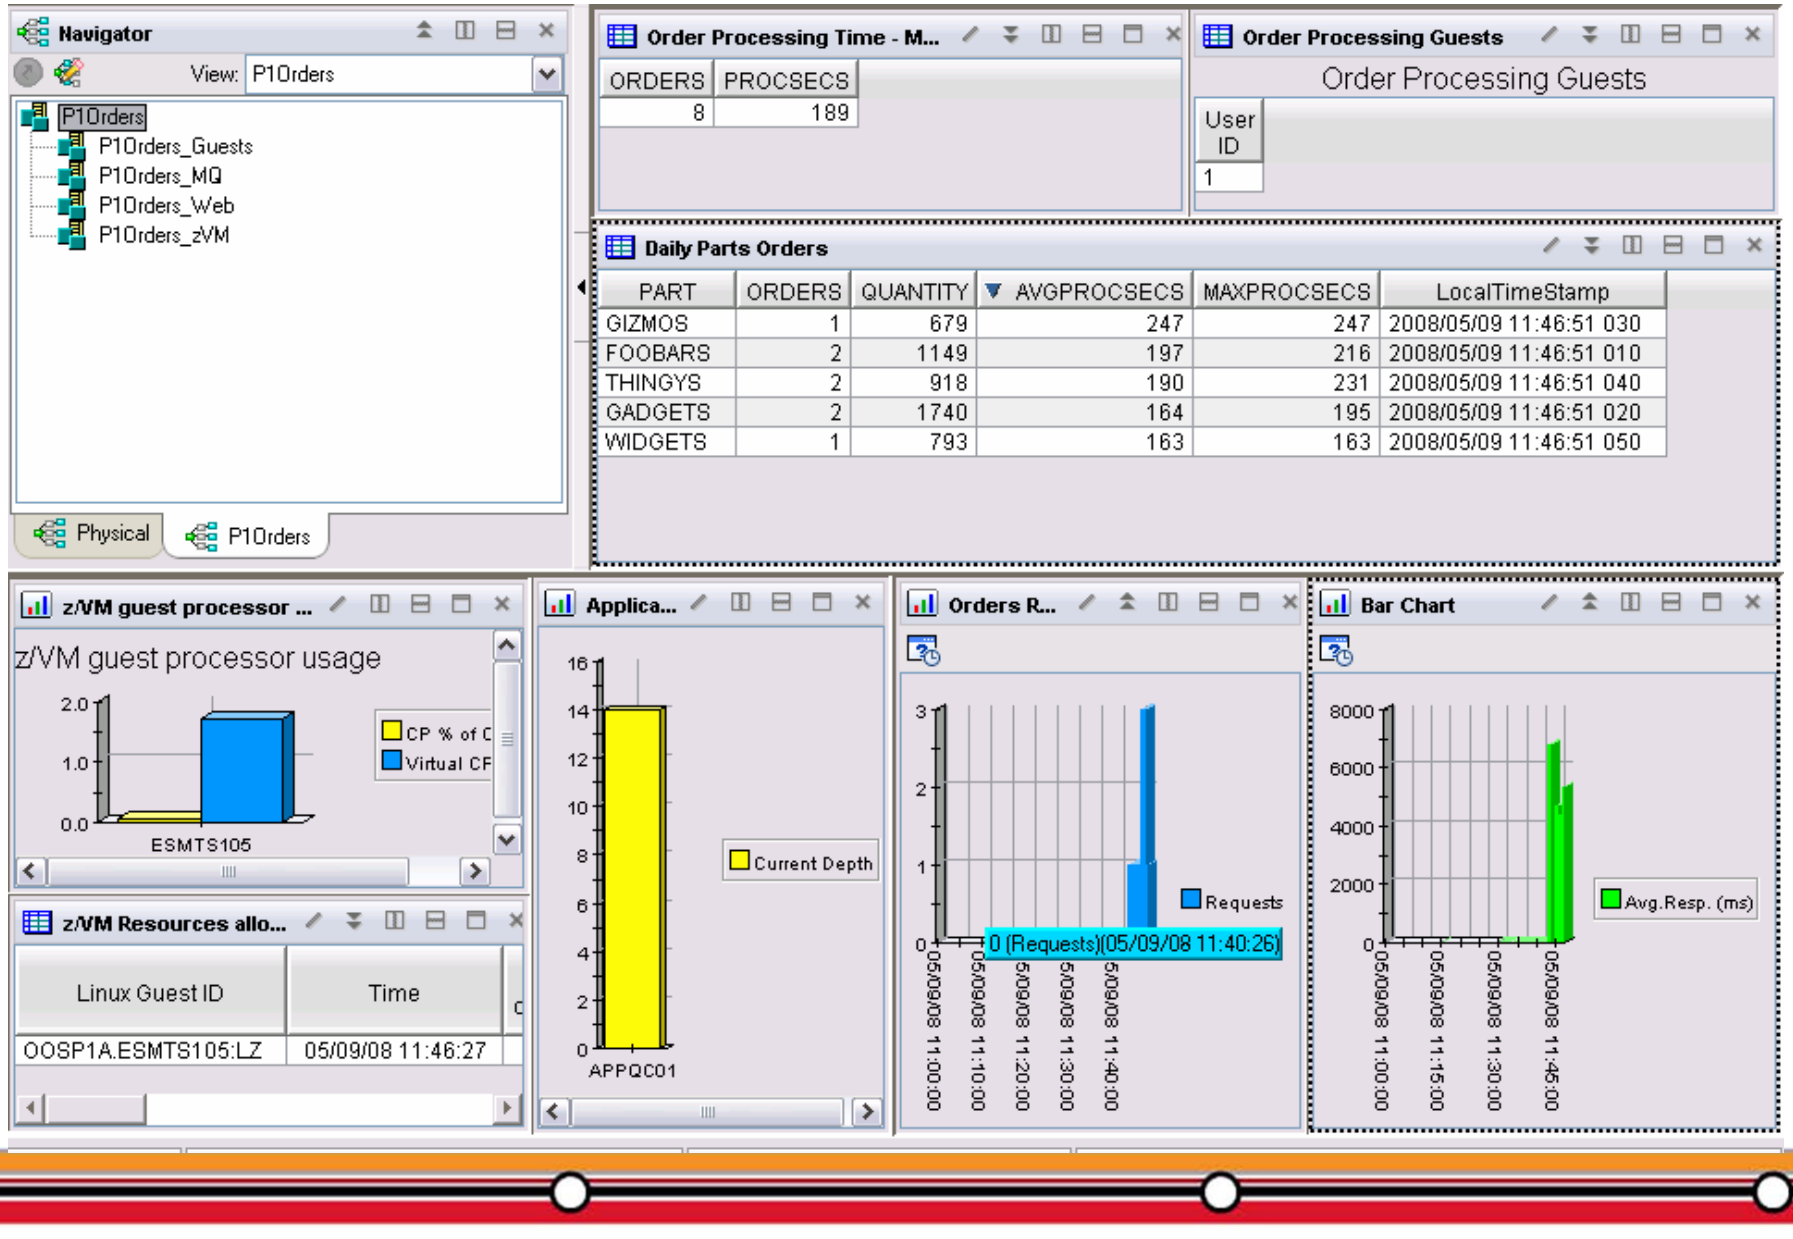

Scaling Scenario

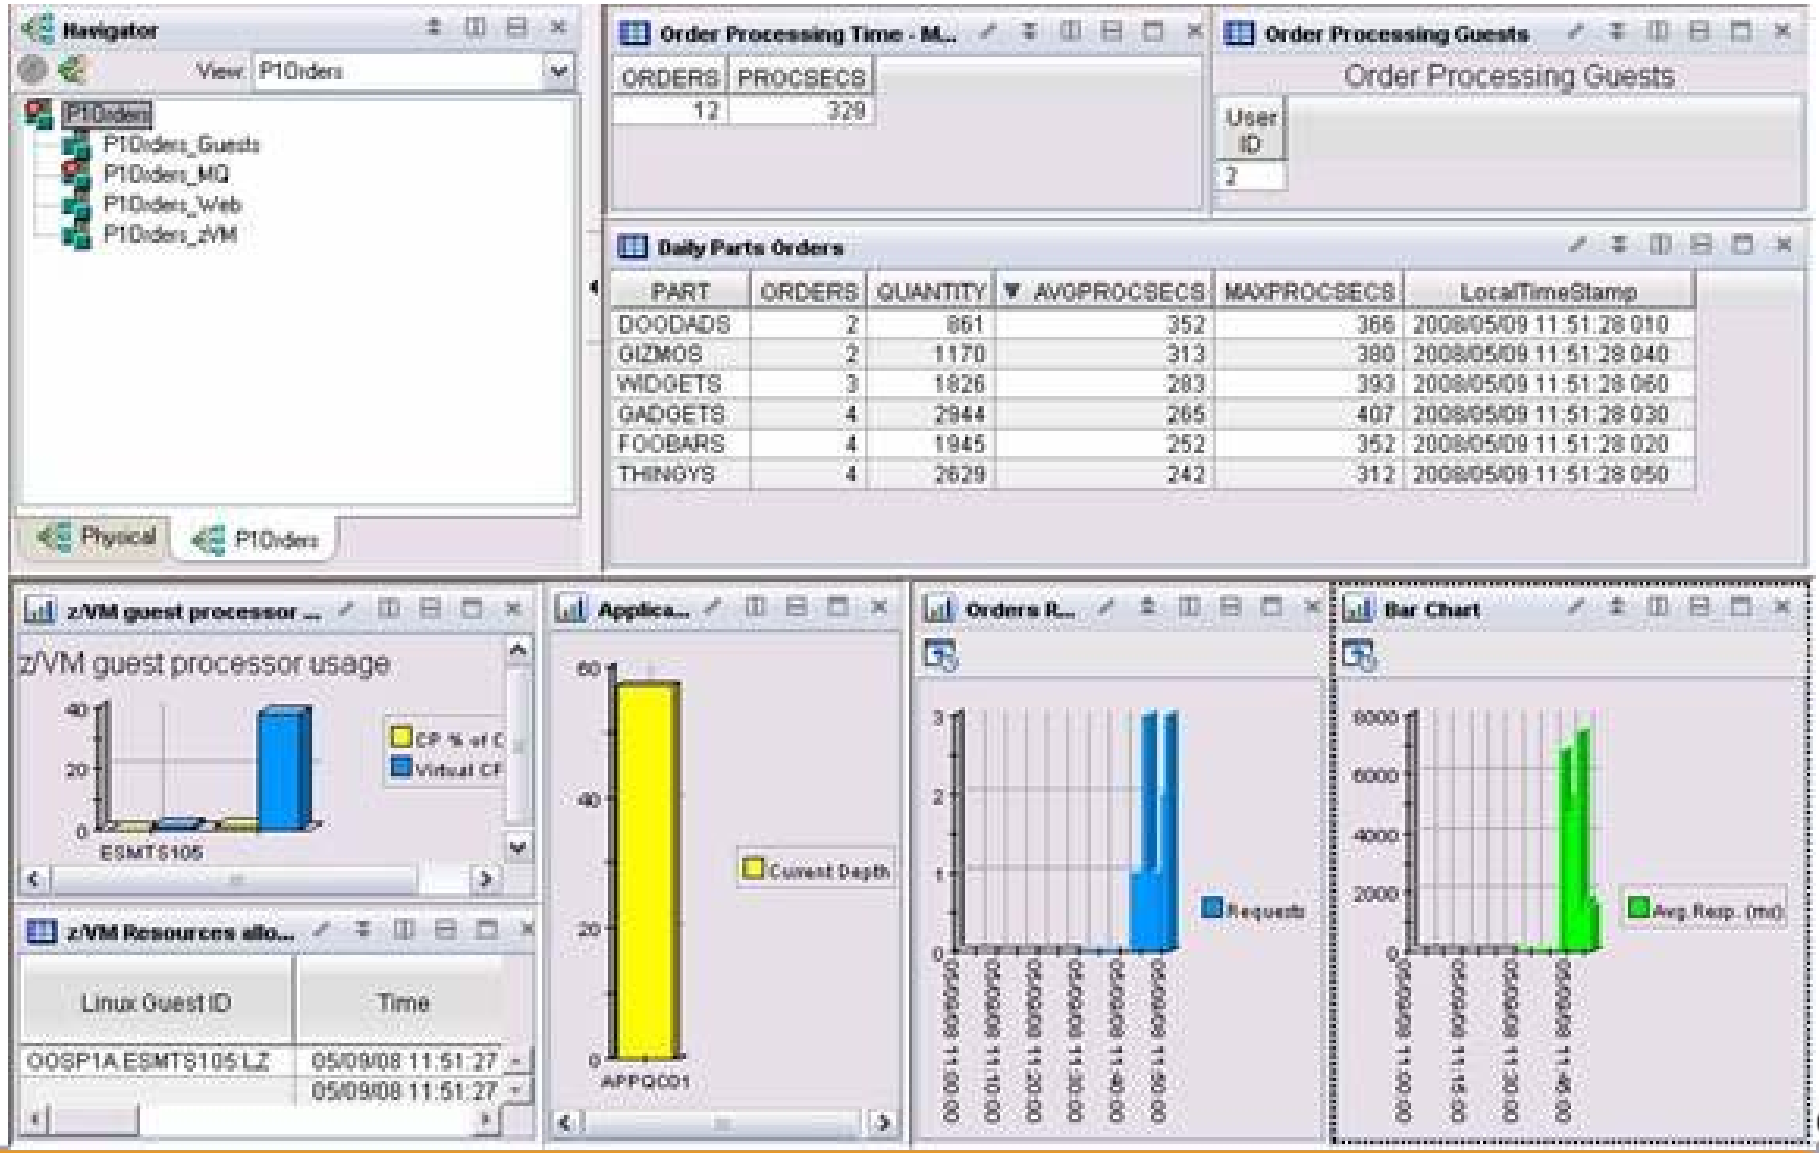

Adjusting Resources for a Linux Guest

- 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

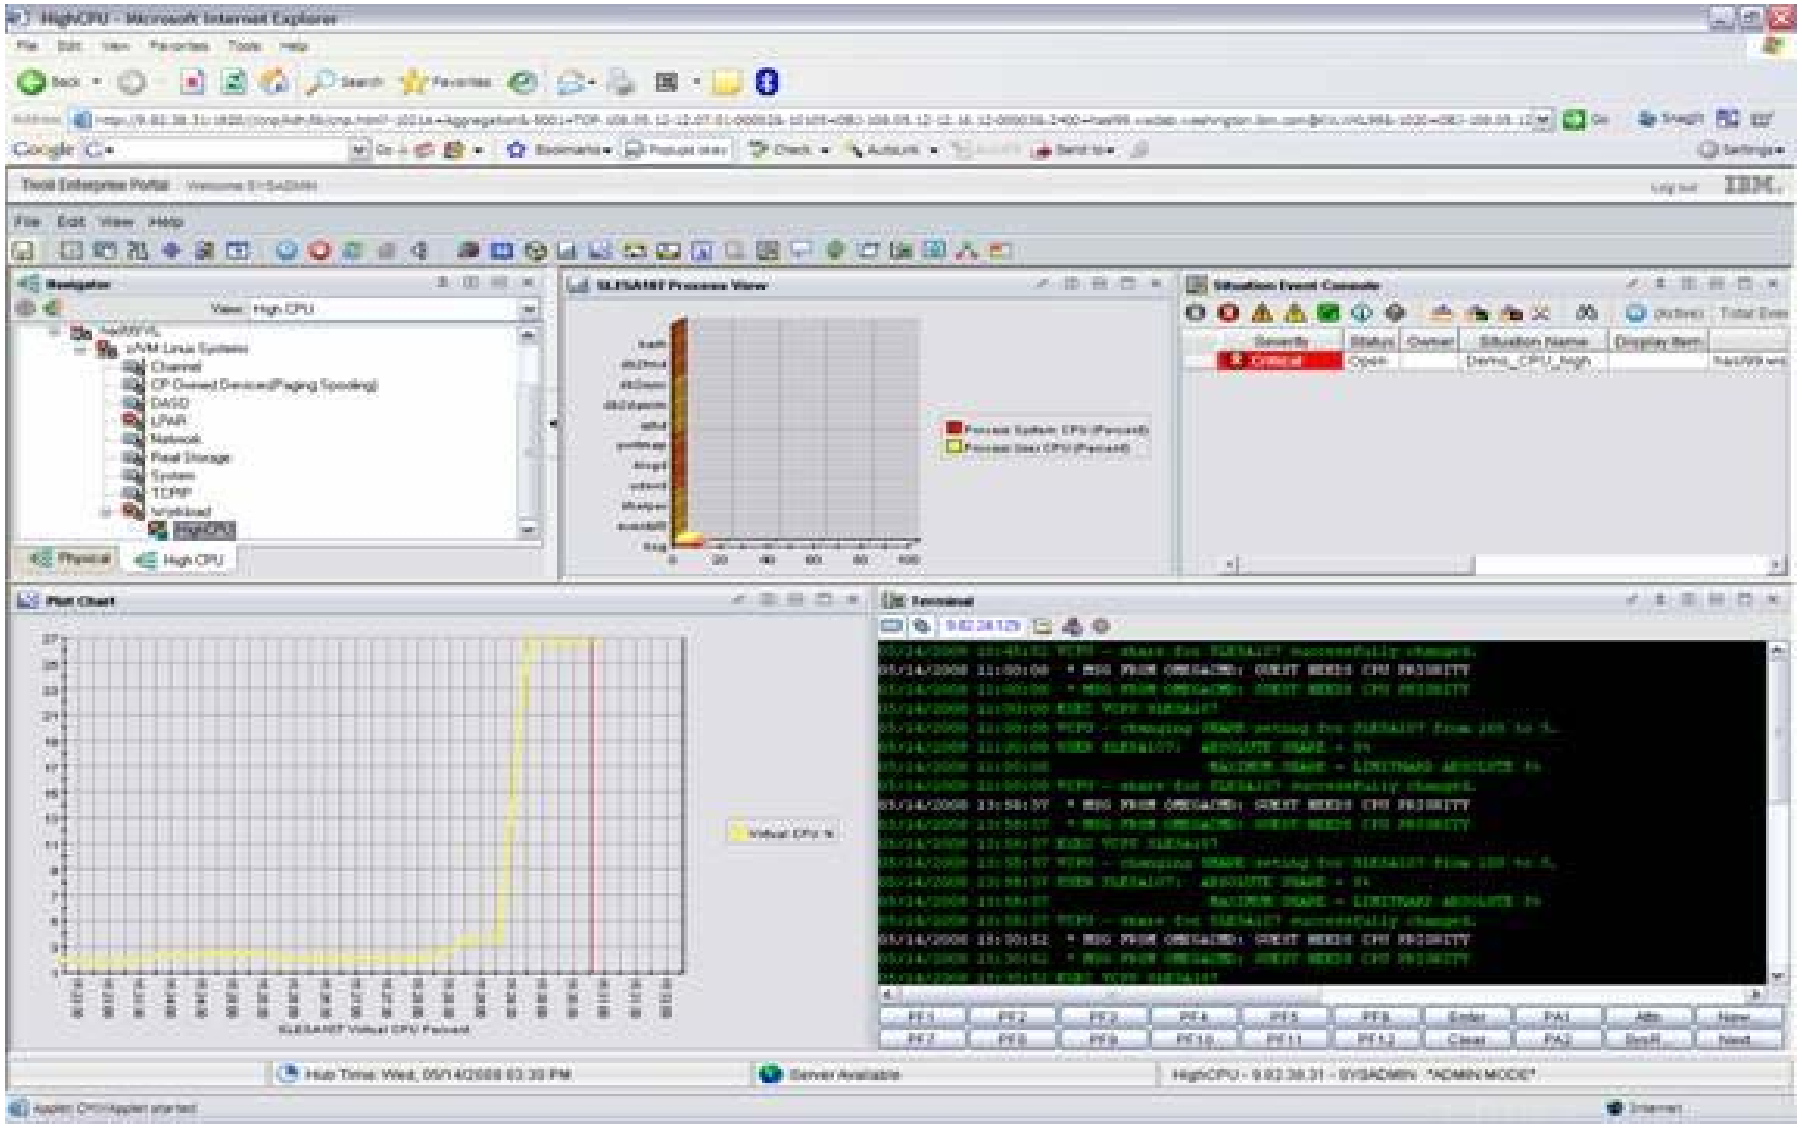

- Define a situation (alert) to detect high CPU consumption for Linux virtual machines.
- Define the automated “Take Action” to:
  - Direct a message to console monitored by Operations Manager.
  - Include in the message keywords to trigger Operations Manager rule.
    - Guest Name
    - Guest need CPU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

धन्यवाद

Hindi

多謝

Traditional Chinese

감사합니다

Korean

Спасибо

Russian

Gracias

Spanish

Thank

English

You

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