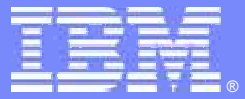

Tivoli Monitoring

Tips on Monitoring and Managing your z/VM and Linux on System z

Session 9152

Bob Neill (neillrd@us.ibm.com)

Tracy Dean (tld1@us.ibm.com)

Tivoli software

ON DEMAND BUSINESS

Acknowledgements

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

Agenda

- Why is it needed?
  - New work loads
  - New monitoring needs
- Brief Product Overview
  - VM Performance overview
  - A monitoring infrastructure – ITM
- System Tips
- Monitoring Scenario
- Additional Product Integration

Opportunity

- New workload
  - Linux on zSeries/System z
  - WAS, DB2, Oracle, Java
- Performance
  - Real and Virtual resources
  - Monitoring needs
  - Managing needs
- Need to be able to do Systems management
  - Suite of products
    - OMEGAMON XE on z/VM and Linux
    - Operations Manager for z/VM
    - Tape Manager for z/VM
    - Backup and Restore Manager for z/VM
    - Archive Manager for z/VM

z/VM Performance Overview - 5,000 Foot View

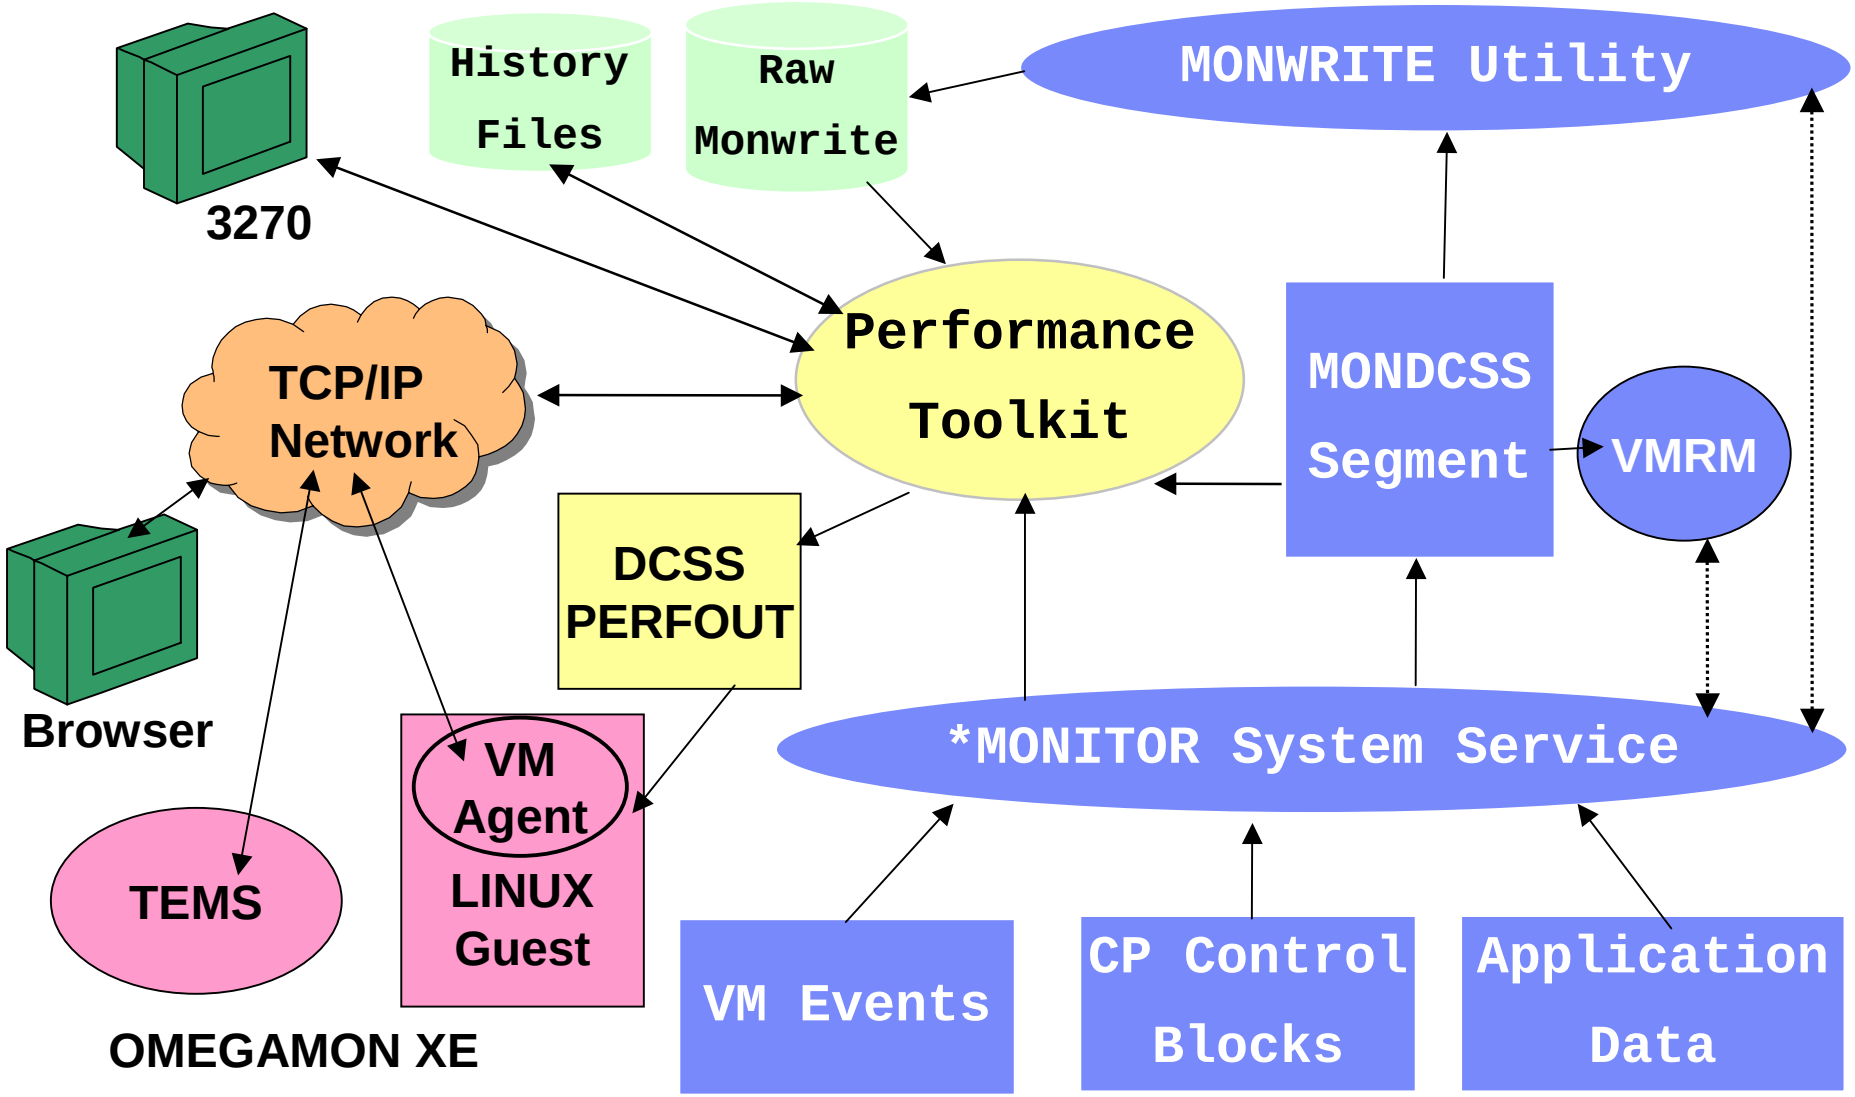

z/VM Performance Toolkit Overview

- Full-Screen operator console (FCON)
- Real Time performance monitor capability for z/VM.
  - CPU Performance
  - Storage Utilization
  - Channel, I/O Device Performance
  - Detailed I/O Analysis
  - Detailed User Performance data
  - TCP/IP Server Performance
  - Linux Performance data
- 3270 interface, with ability to exploit GDDM graphics.
- Web server capability.
- Ability to customize screens.
- Some integration with other platform monitors (Linux).

B - GDLVM7 - [32 x 80]

```
FCX124          Performance Screen Selection  (FL520 VM63967)    GDLVM7

General System Data          I/O Data                    History Data (by Time)
1. CPU load and trans.       11. Channel load            31. Graphics selection
2. Storage utilization       12. Control units           32. History data files*
3. Reserved                  13. I/O device load*        33. Benchmark displays*
4. Priv. operations          14. CP owned disks*         34. Correlation coeff.
5. System counters           15. Cache extend. func.*    35. System summary*
6. CP IUCV services          16. DASD I/O assist         36. Auxiliary storage
7. SPOOL file display*       17. DASD seek distance*     37. CP communications*
8. LPAR data                 18. I/O prior. queueing*    38. DASD load
9. Shared segments           19. I/O configuration       39. Minidisk cache*
A. Shared data spaces        1A. I/O config. changes     3A. Storage mgmt. data*
B. Virt. disks in stor.                                  3B. Proc. load & config*
C. Transact. statistics      User Data                   3C. Logical part. load
D. Monitor data              21. User resource usage*    3D. Response time (all)*
E. Monitor settings          22. User paging load*       3E. RSK data menu*
F. System settings           23. User wait states*       3F. Scheduler queues
G. System configuration      24. User response time*     3G. Scheduler data
H. VM Resource Manager       25. Resources/transact.*    3H. SFS/BFS logs menu*
                             26. User communication*     3I. System log
I. Exceptions                27. Multitasking users*     3K. TCP/IP data menu*
                             28. User configuration*     3L. User communication
K. User defined data*        29. Linux systems*          3M. User wait states

            Pointers to related or more detailed performance data
            can be found on displays marked with an asterisk (*).


Select performance screen with cursor and hit ENTER
Command ===>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return

MA    b                                                        31/015
```

OMEGAMON XE on z/VM and Linux agents

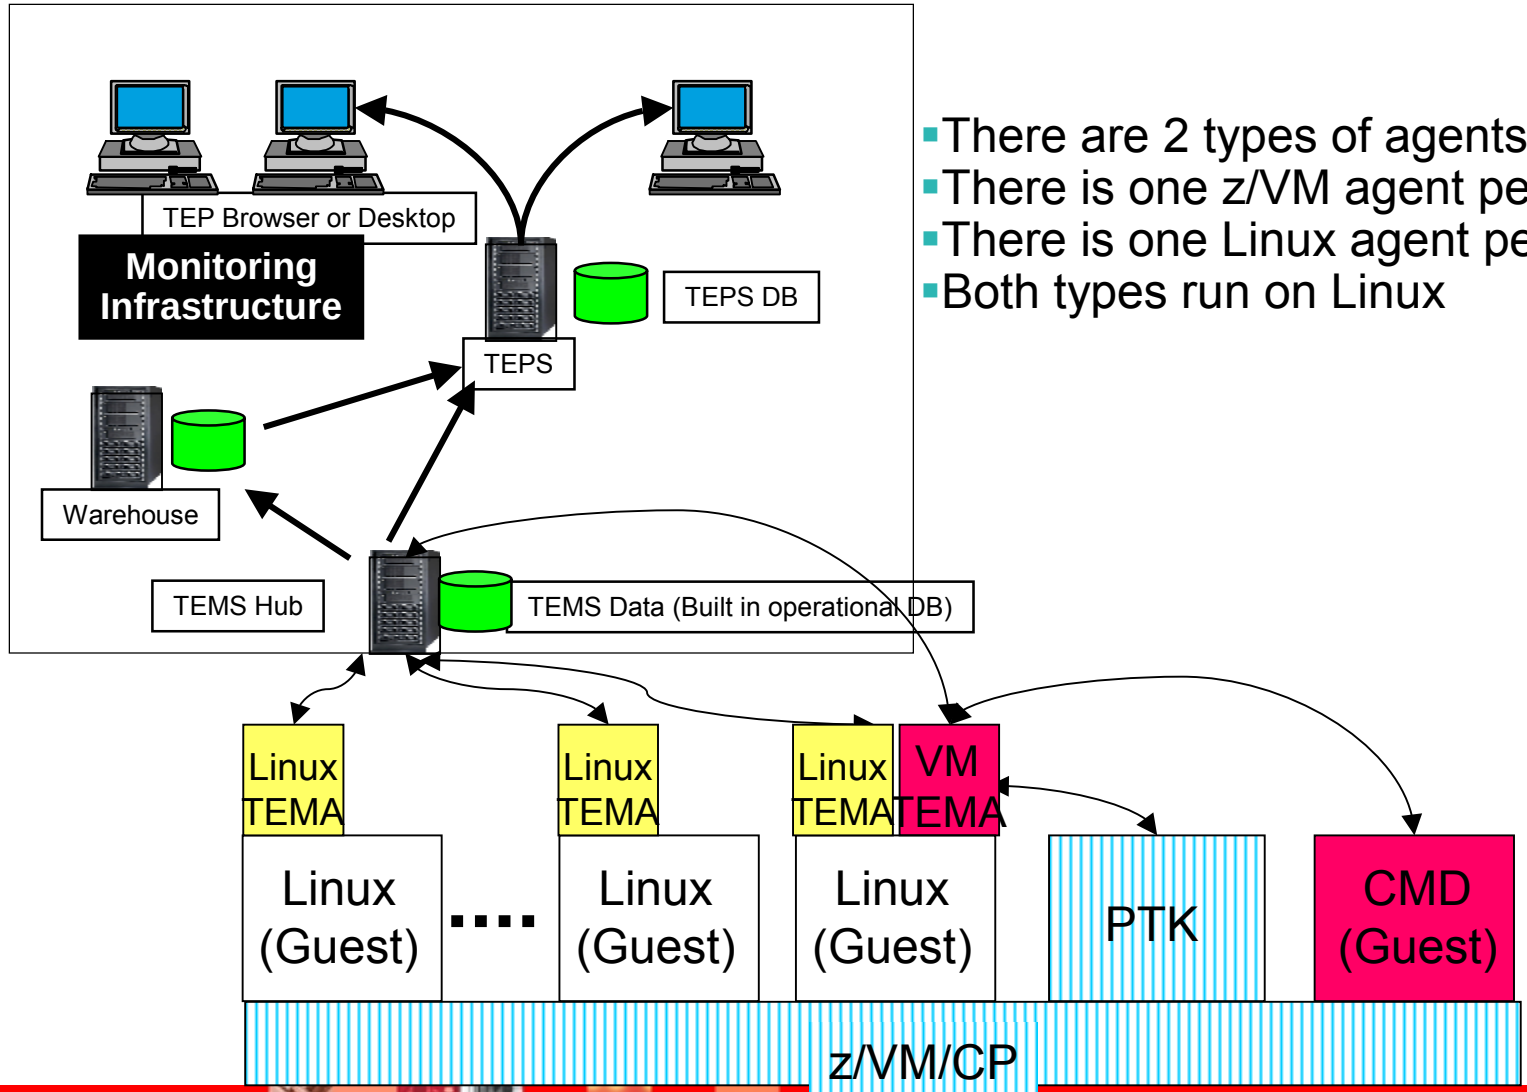

- There are 2 types of agents
- There is one z/VM agent per z/VM LPAR
- There is one Linux agent per Linux Guest
- Both types run on Linux

Tivoli Enterprise Portal

Improve your ability to Diagnose and Resolve Problems

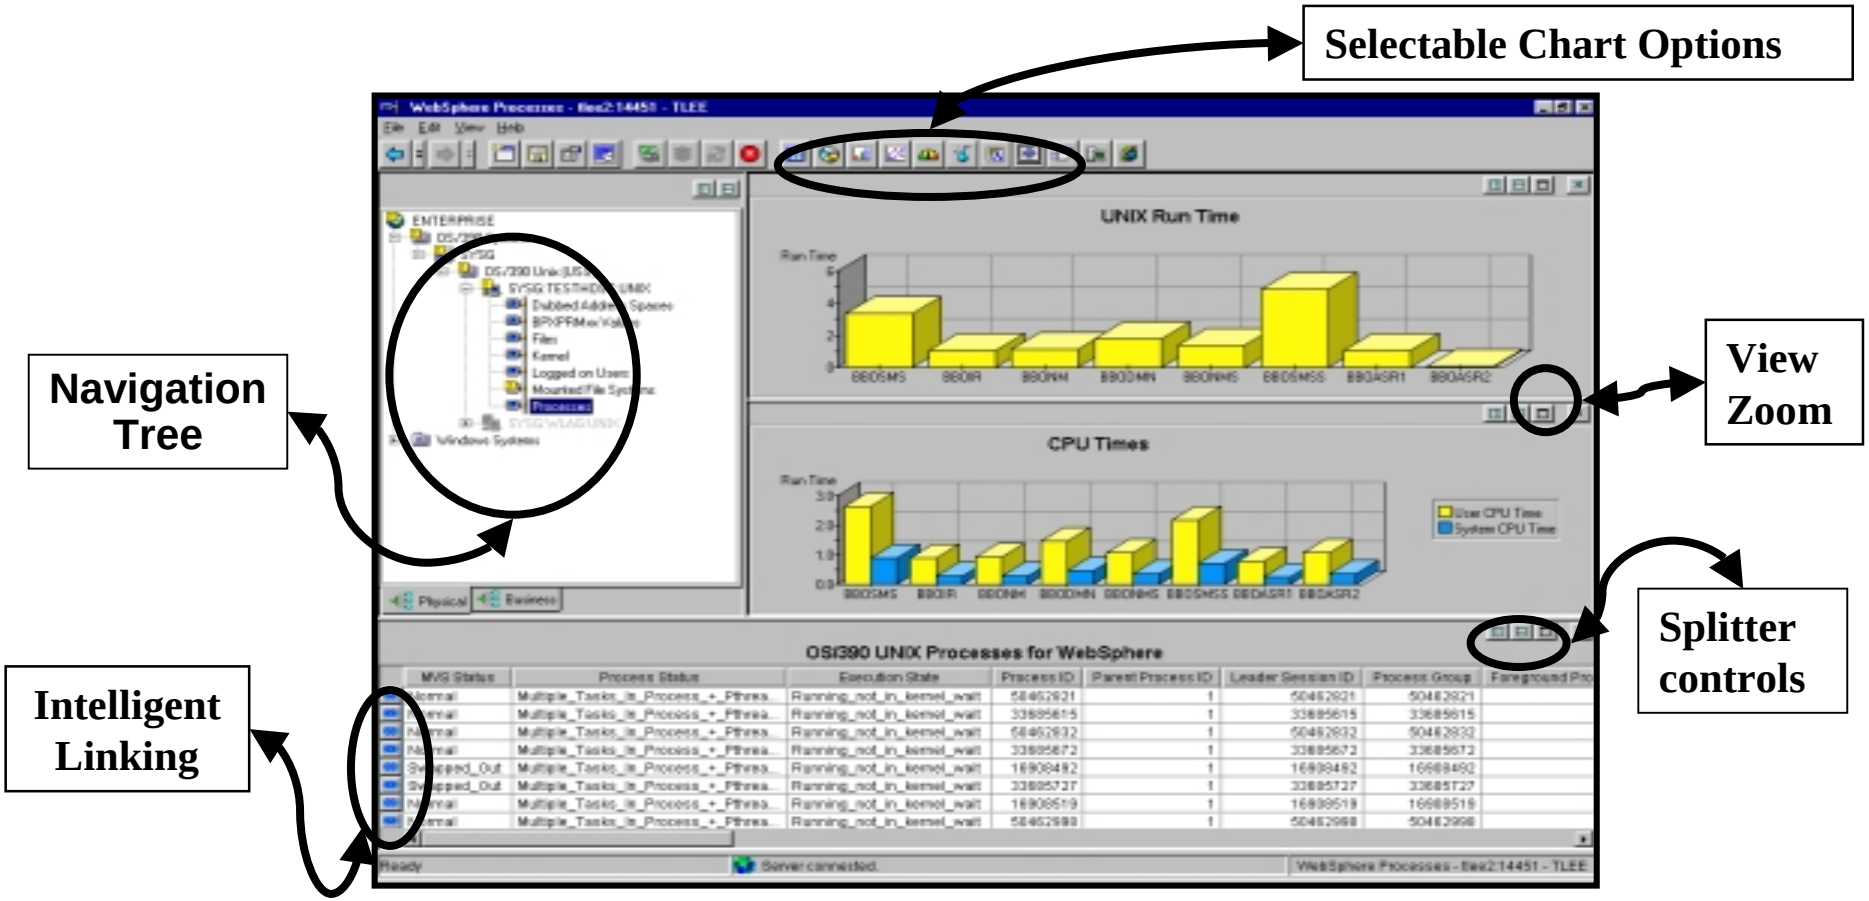

Situations

- A situation describes one or more conditions that you want to test
  - Each condition compares a user-specified value against attribute data collected from managed systems
- If all conditions are met, the situation evaluates to true and an alert indicator icon appears on the TEP to let you know that a problem exists
- When you create a situation, you can also specify automated responses to take place when the situation becomes true (Take Action)
- Each management agent comes with a set of pre-defined situations that can be set to start running as soon as the management agent is connected

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**

Linux

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

Tips and Scenarios

- Overall health of your z/VM systems
- Managing CP Owned Volumes
- Paging rate problems
- LPAR and processor usage
- Spin Locks
- Virtual Disks
- Minidisk Cache
- DASD
- I/O contention
- Sizing Linux guests
- System running slowly—how to debug

Tips—Overall Health of Your System

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

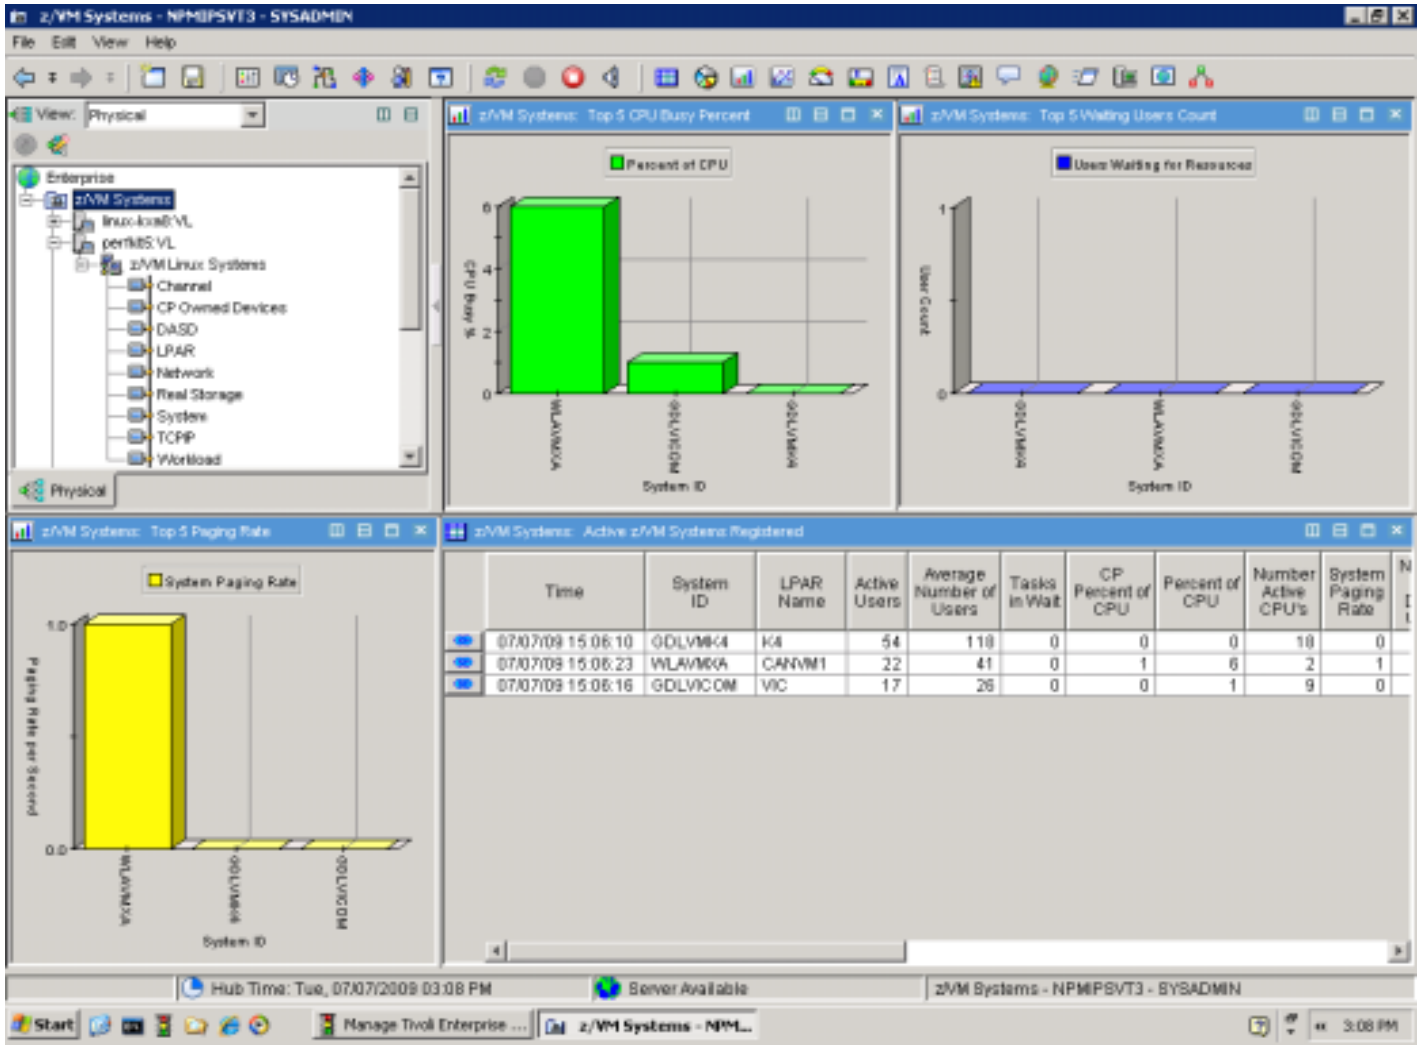

Tips—Overall Health of Your System

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

Tips—Overall Health of Your System

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

Tips—CP Owned Volumes

One place to check to see if you need to add more space on your CP-Owned volumes

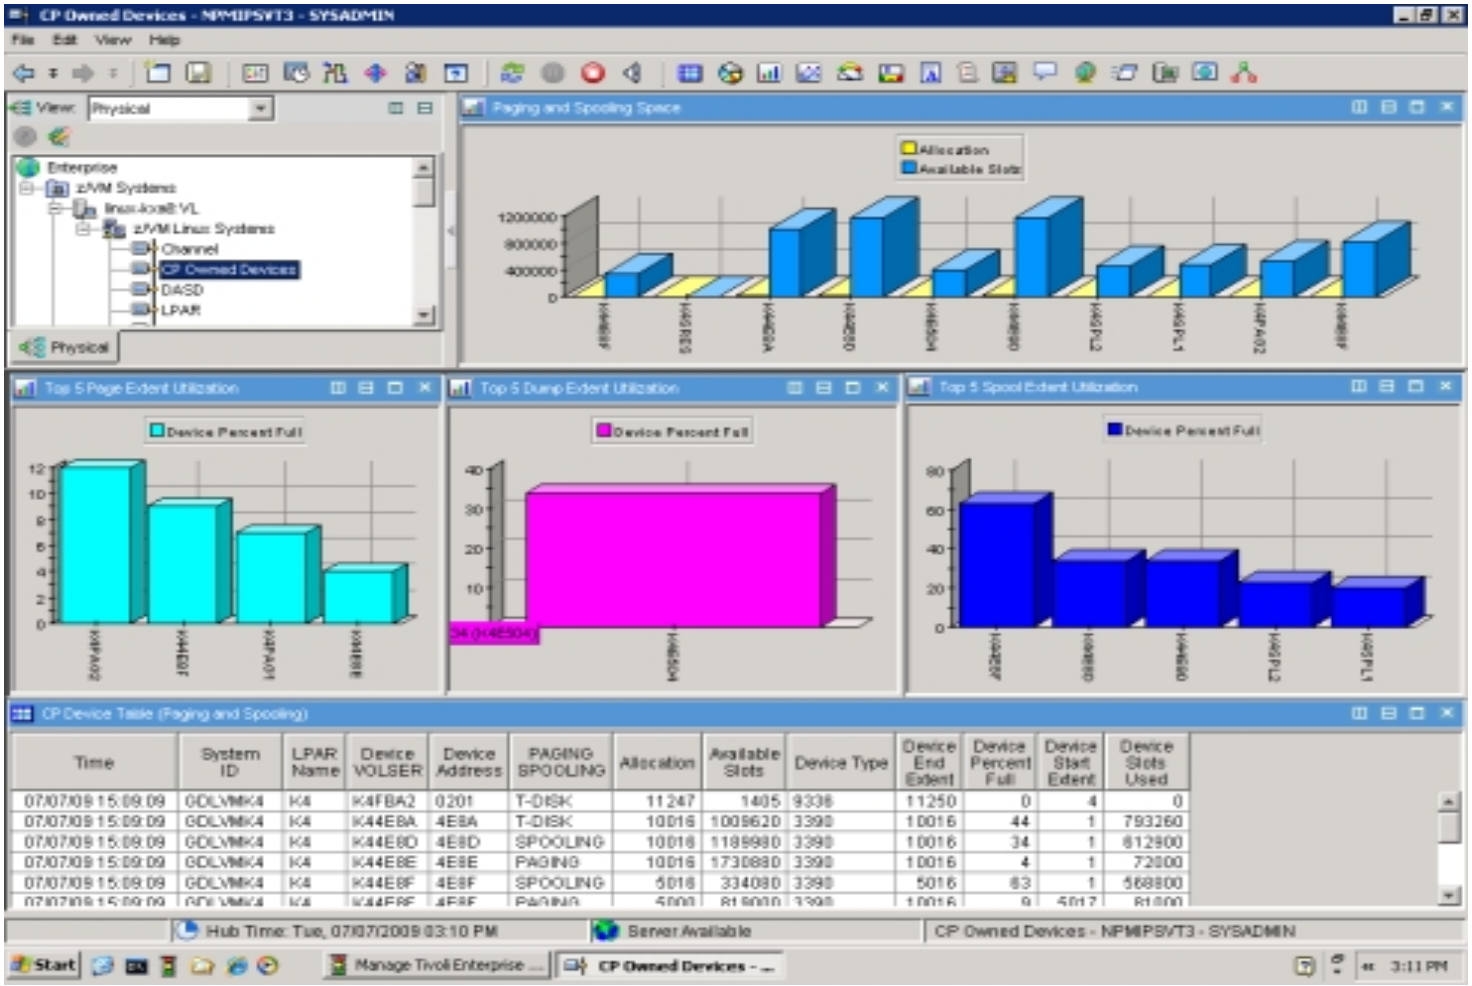

Tips—CP Owned Volumes

- CP-Owned Volumes
  - Consist of
    - Page, Spool, Tdisk, Directory and Dump (subset of spool)
  - Considerations
    - Page
      - Are my volumes getting full—do I need to add more
      - Is my paging spread out sufficiently
    - Spool
      - Used mainly for Reader and Print files, Dumps and Named Saved Systems
      - If full, can I delete old spool files, or do I need more?
      - Good to have an automatic cleanup program, perhaps based on age of spool file (SFPURGER utility)

Tips—CP Owned Volumes (paging)

- General tips
  - Page space utilization should always be < 50%
  - Never put Paging and Spool space on the same volume
  - Allocate Spool and Page volumes to try and reduce I/O contention by separating them as much as possible (control unit, channel, etc)
  - Dedicated paging devices reduce contention for paging
  - Don’t put highly used files on the same volume as paging and spool space, such as the CMS system disk
  - Use devices of the same size and geometry
  - Use your fastest devices for Paging
  - Multiple Paging devices allow more overlap of paging operations
  - Large contiguous free space allows for greater paging efficiency
  - Expanded storage can be used for paging

Tips—LPAR Usage

LPAR workspace allows you to look at all your LPARs across the CEC

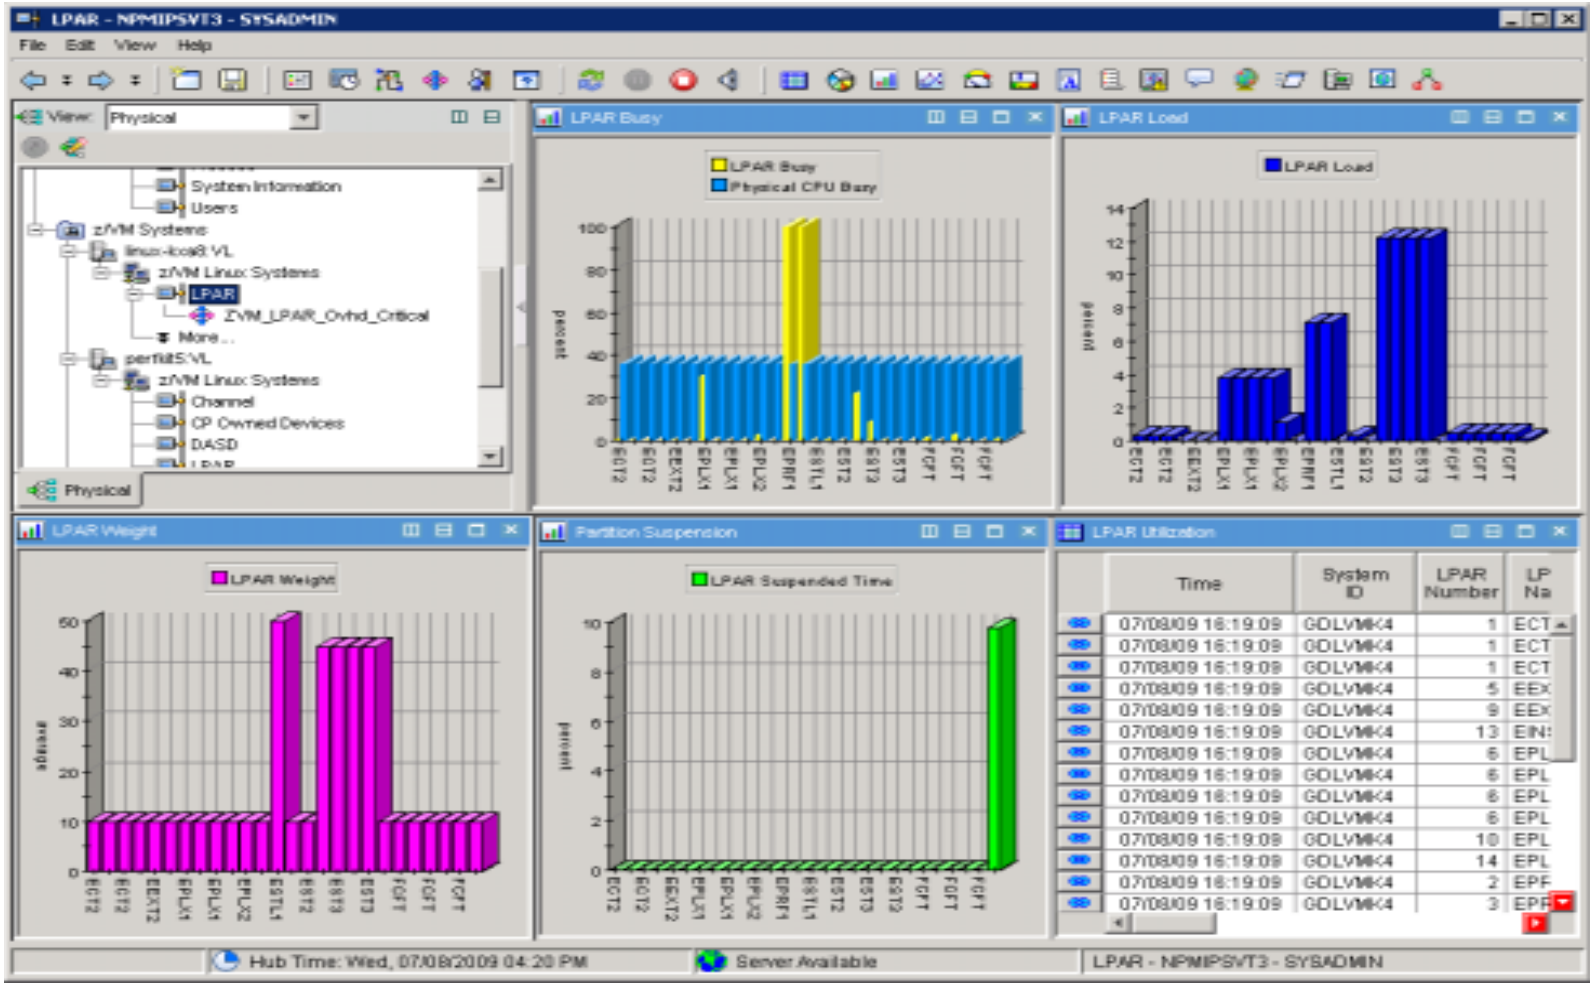

Tips—LPAR Usage

- Considerations
  - Are the LPAR weights balanced as you expected
  - Do you have an dedicated processors, and should you?
  - Do you have the right number of processors per LPAR?
    - Too few means tasks must wait for a processor
      - Multi-threaded applications can use multiple processors concurrently
    - Too many means extra overhead
  - Look at the %Busy
    - Are there some processors being underutilized?
      - Are they needed?
      - Can you shift work (virtual machines) to them
  - Look at multiple intervals, or use historical data before jumping to conclusions
  - May want to drill down to individual processors for more details

System Processor Utilization Workspace

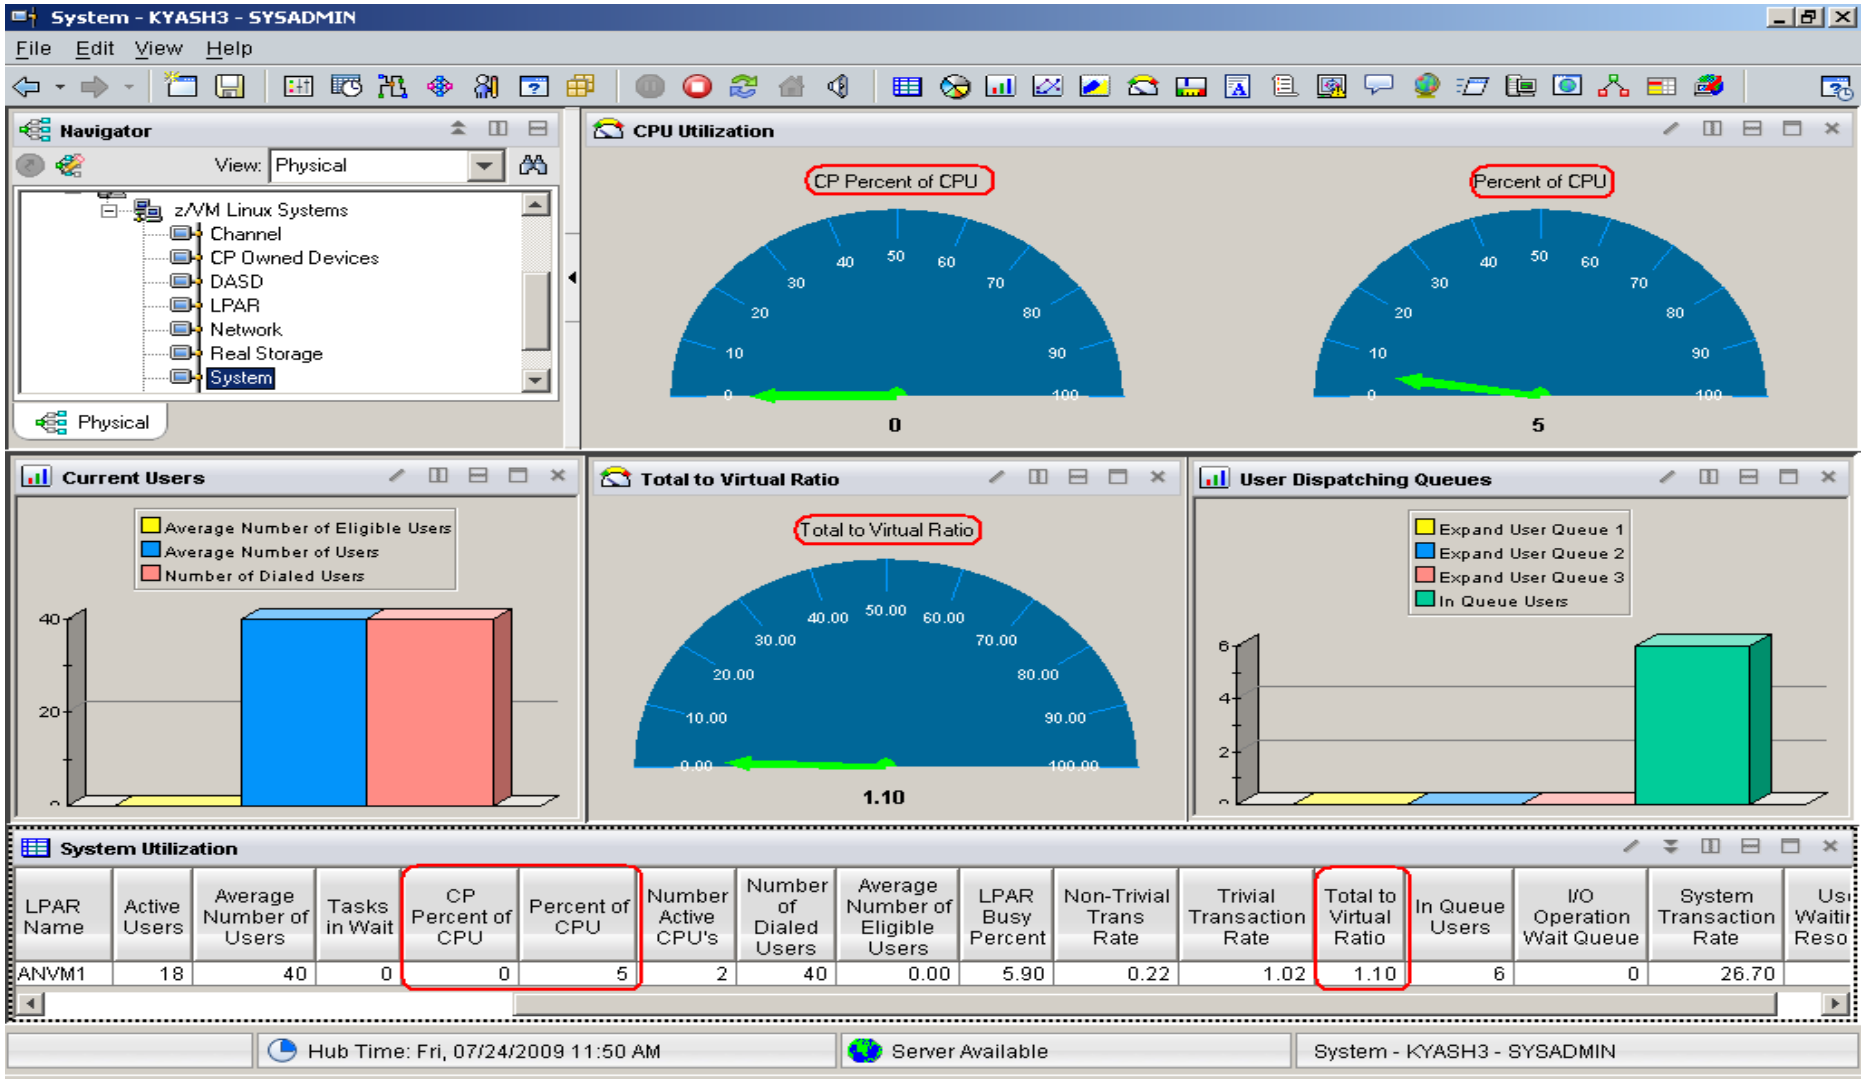

Processor Utilization

- **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.

Spinlock Workspace

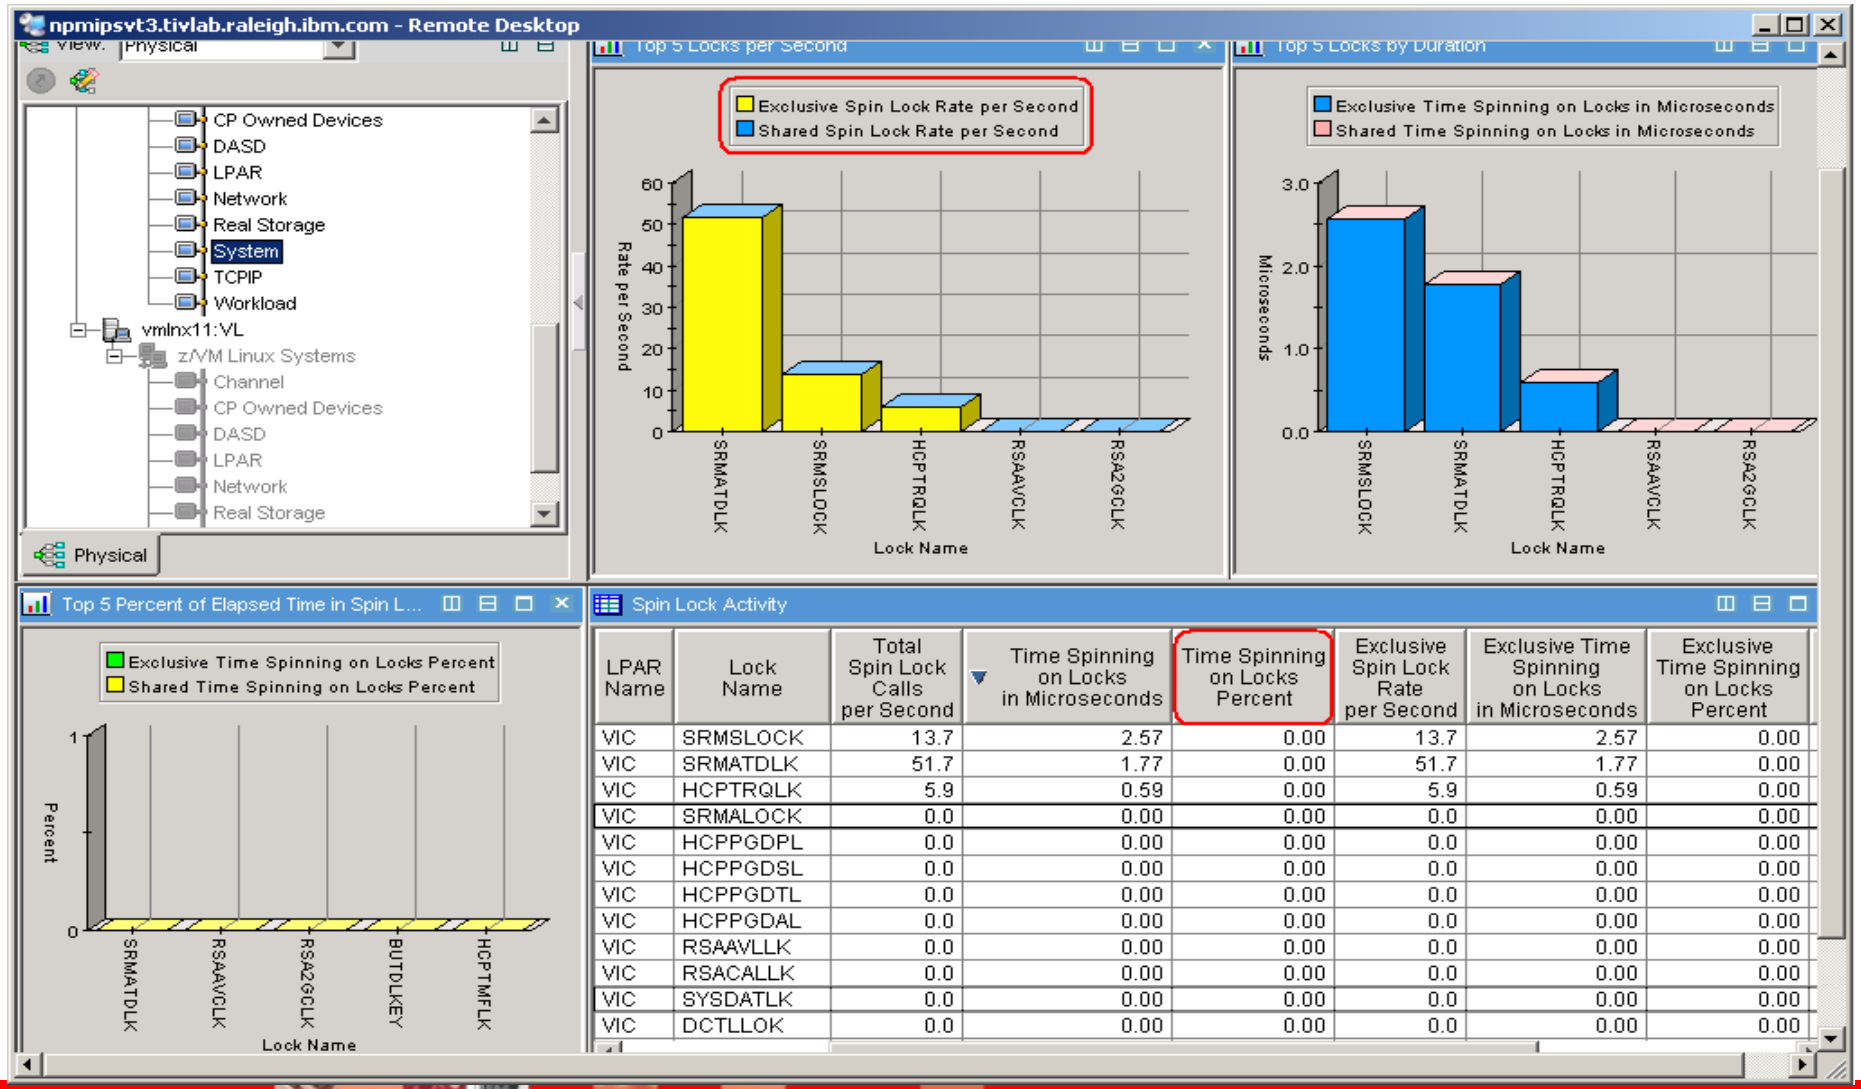

LPAR Name	Lock Name	Total Spin Lock Calls per Second	Time Spinning on Locks in Microseconds	Time Spinning on Locks Percent	Exclusive Spin Lock Rate per Second	Exclusive Time Spinning on Locks in Microseconds	Exclusive Time Spinning on Locks Percent
VIC	SRMSLOCK	13.7	2.57	0.00	13.7	2.57	0.00
VIC	SRMATDLK	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

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.

VDISK Workspace

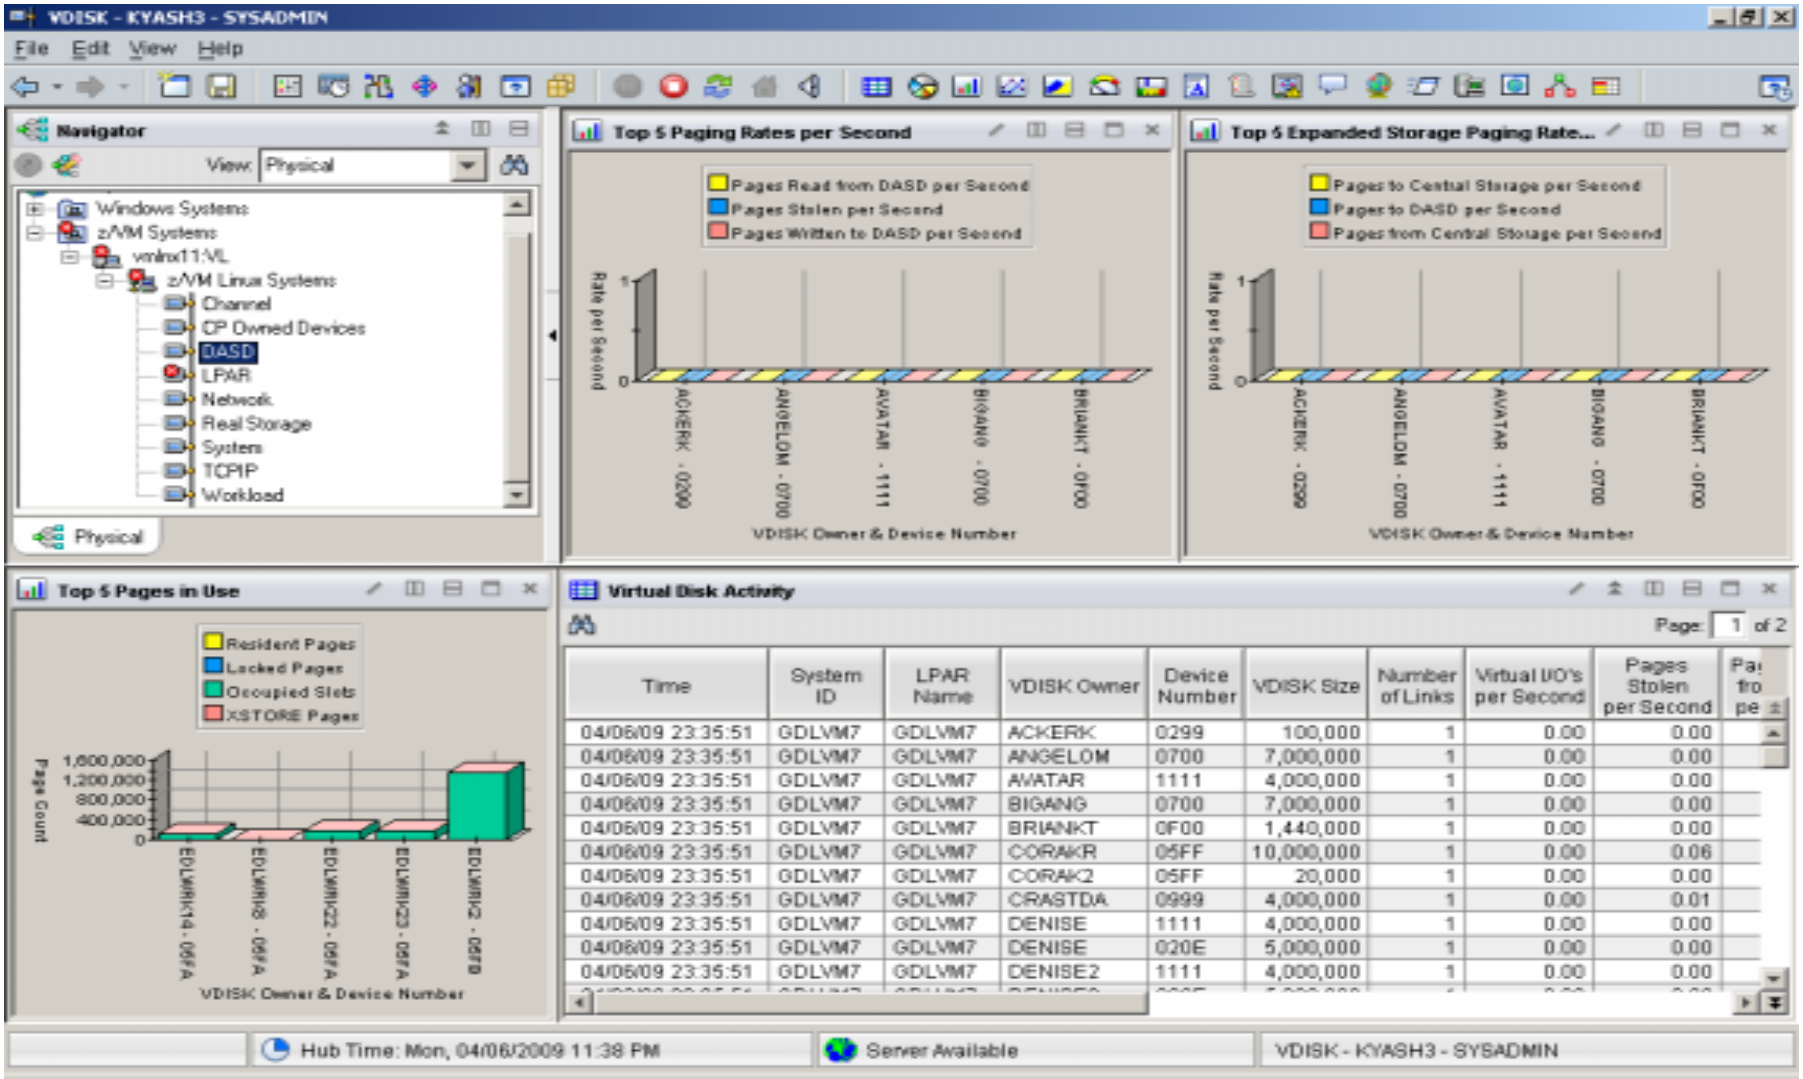

VDISK

- 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.

OMEGAMON MDISK Cache Allocations

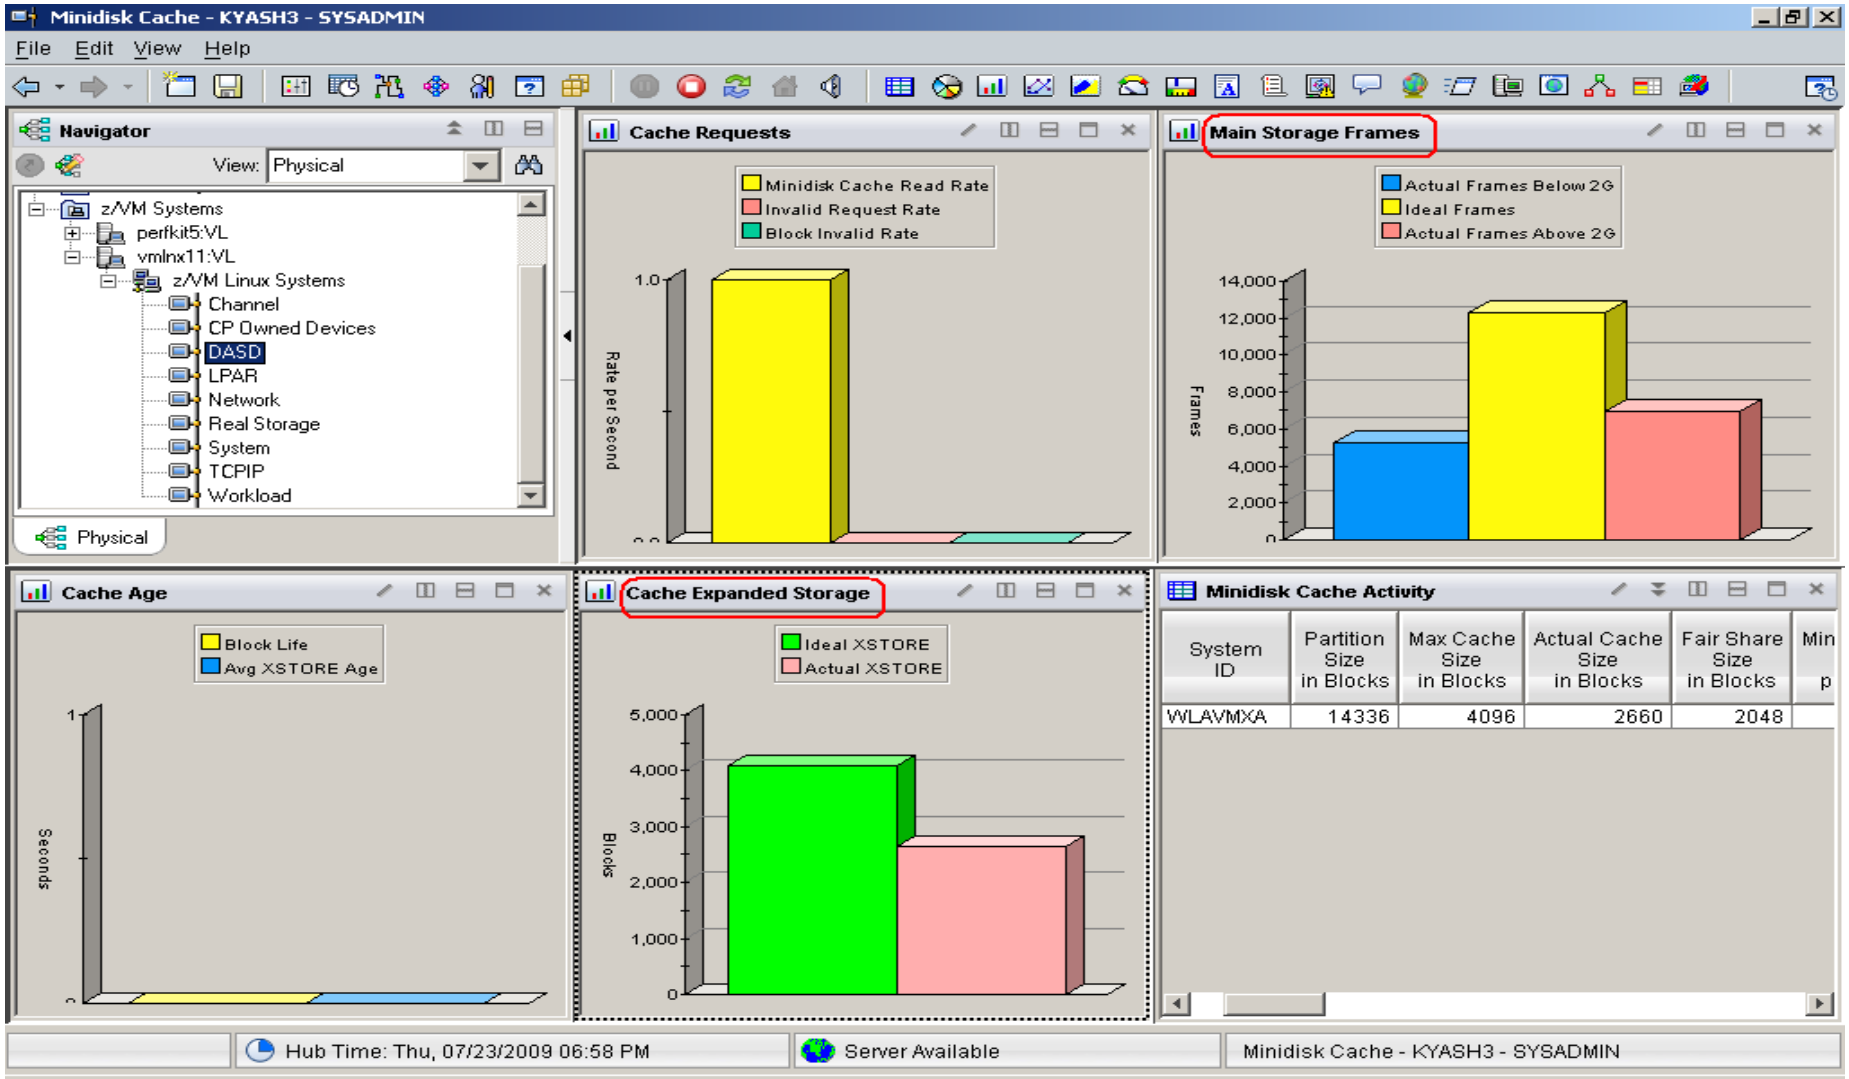

Minidisk Cache

- 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.

DASD I/O Workspace

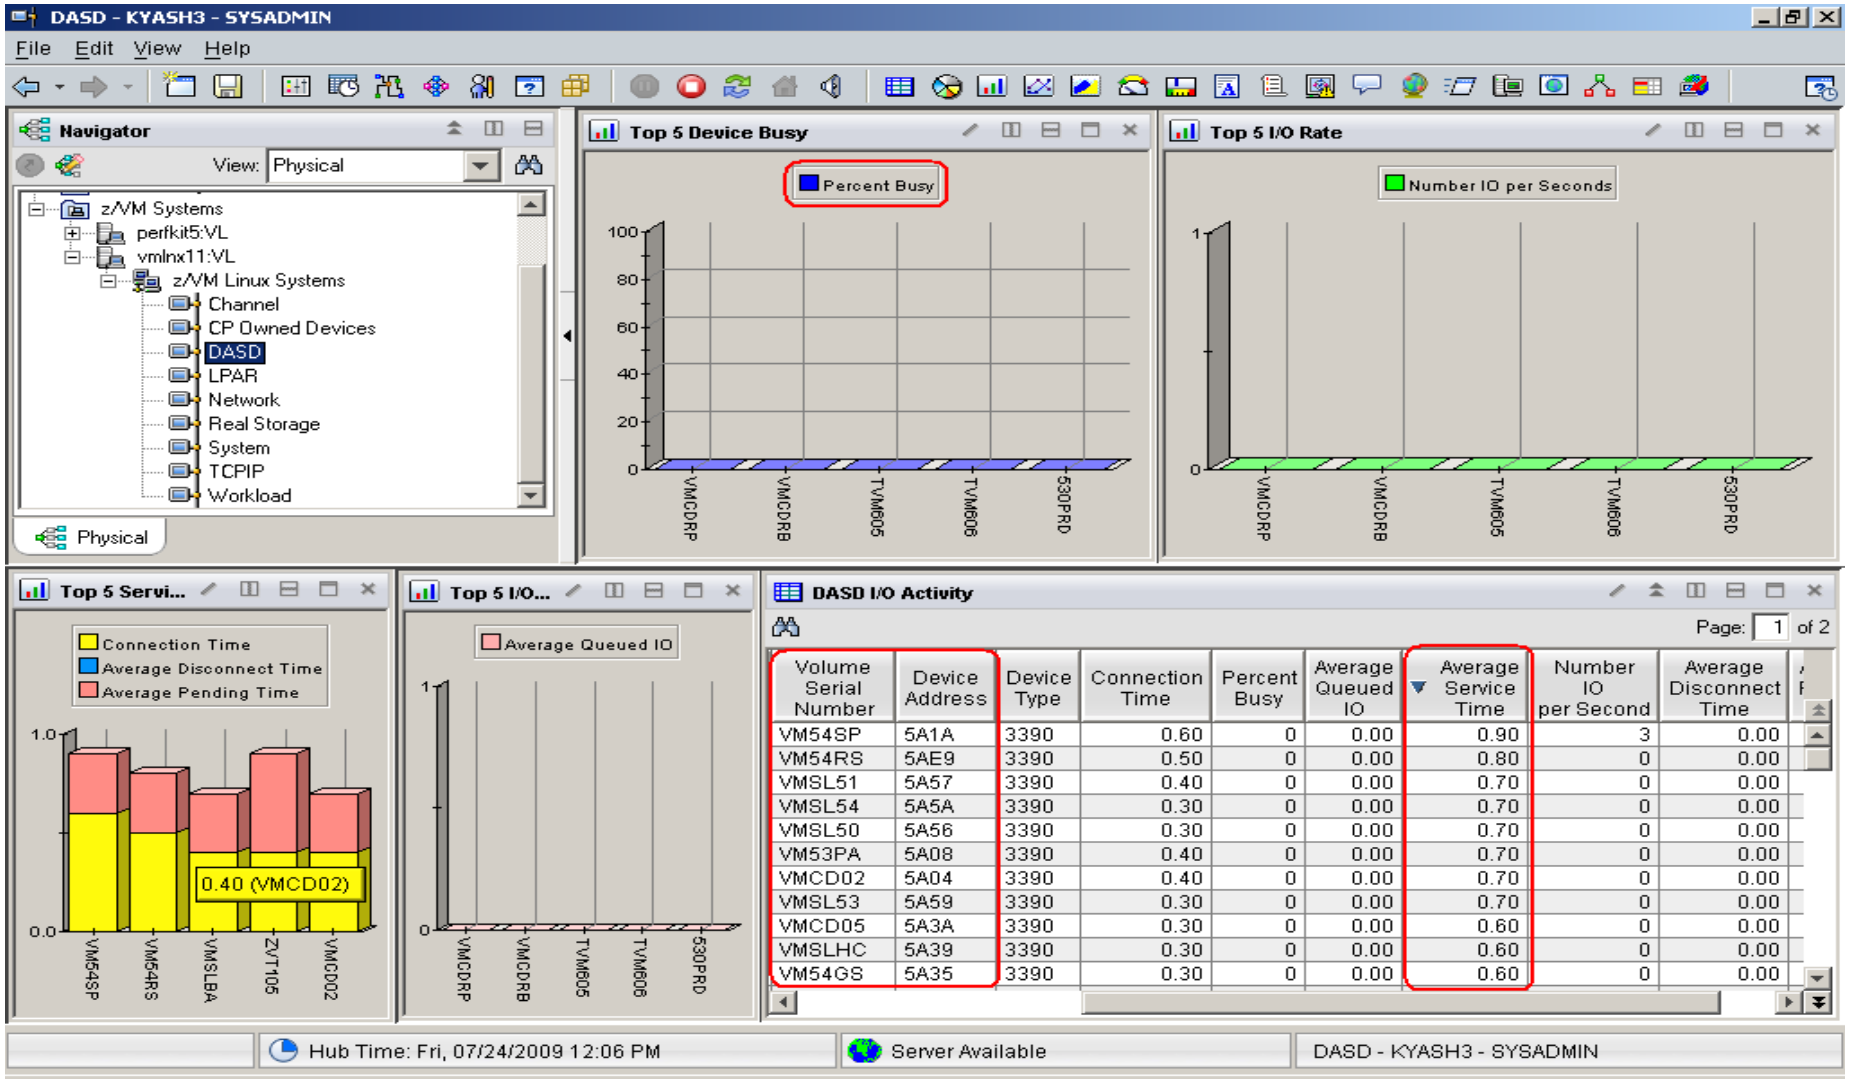

Direct Access Storage Devices (DASD)

- **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.

Sizing Linux Guests

Memory usage of a particular Linux virtual machine

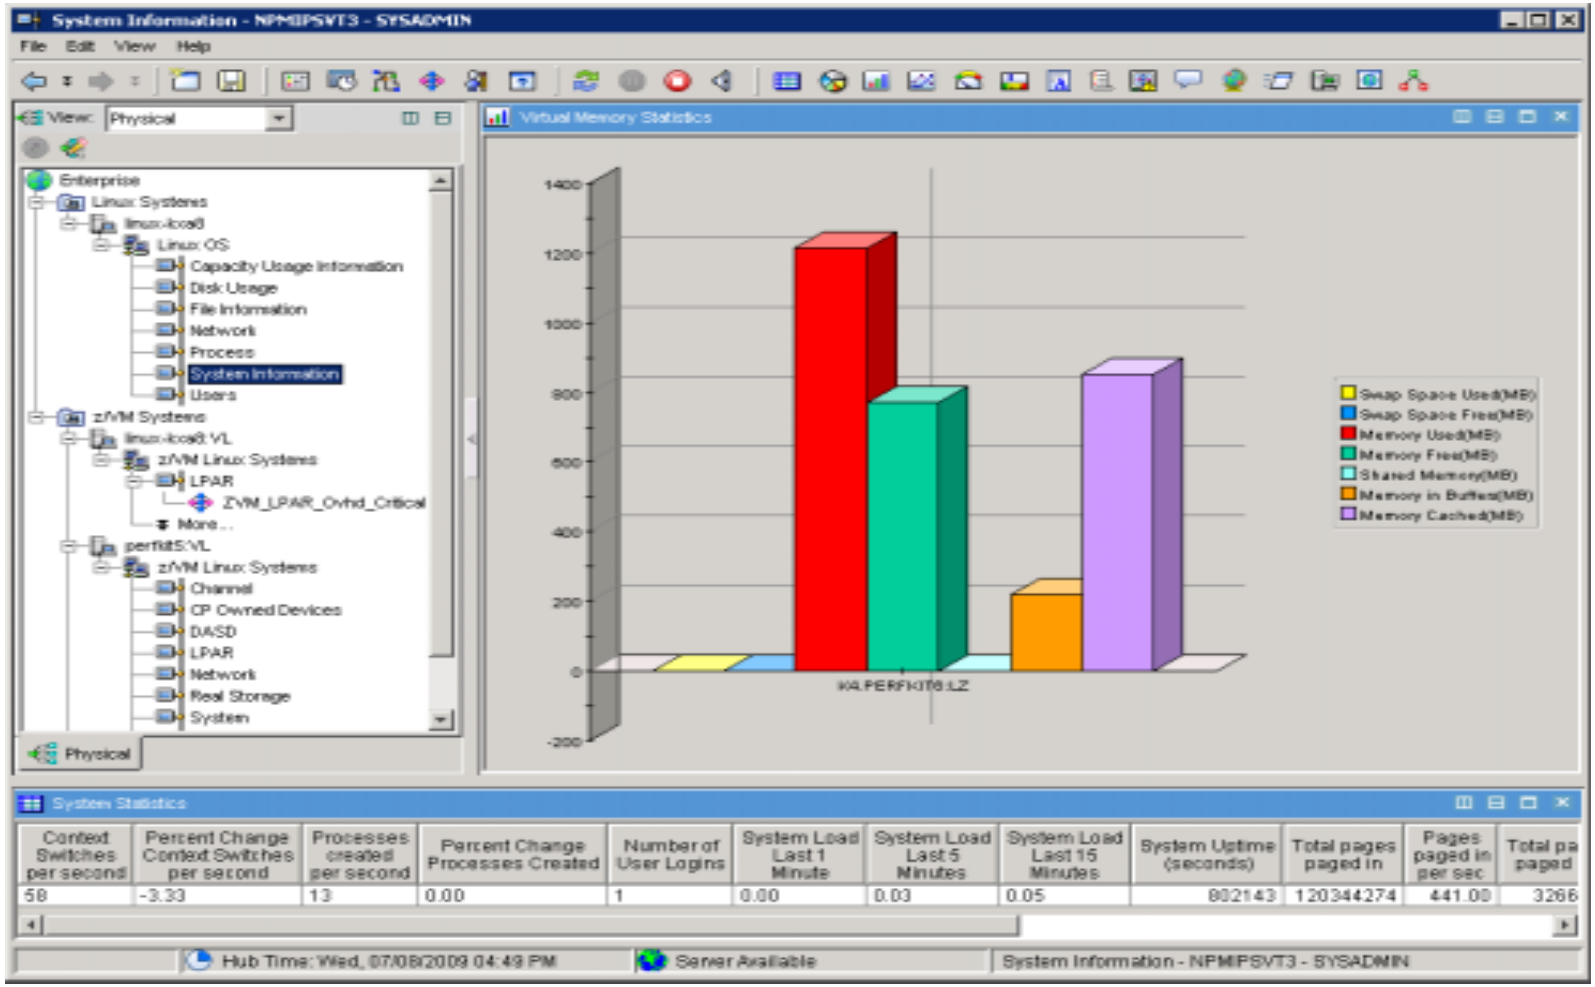

Sizing Linux Guests

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

Sizing Linux Guests

Working Set Size can be found on the Workload workspace of the z/VM agent

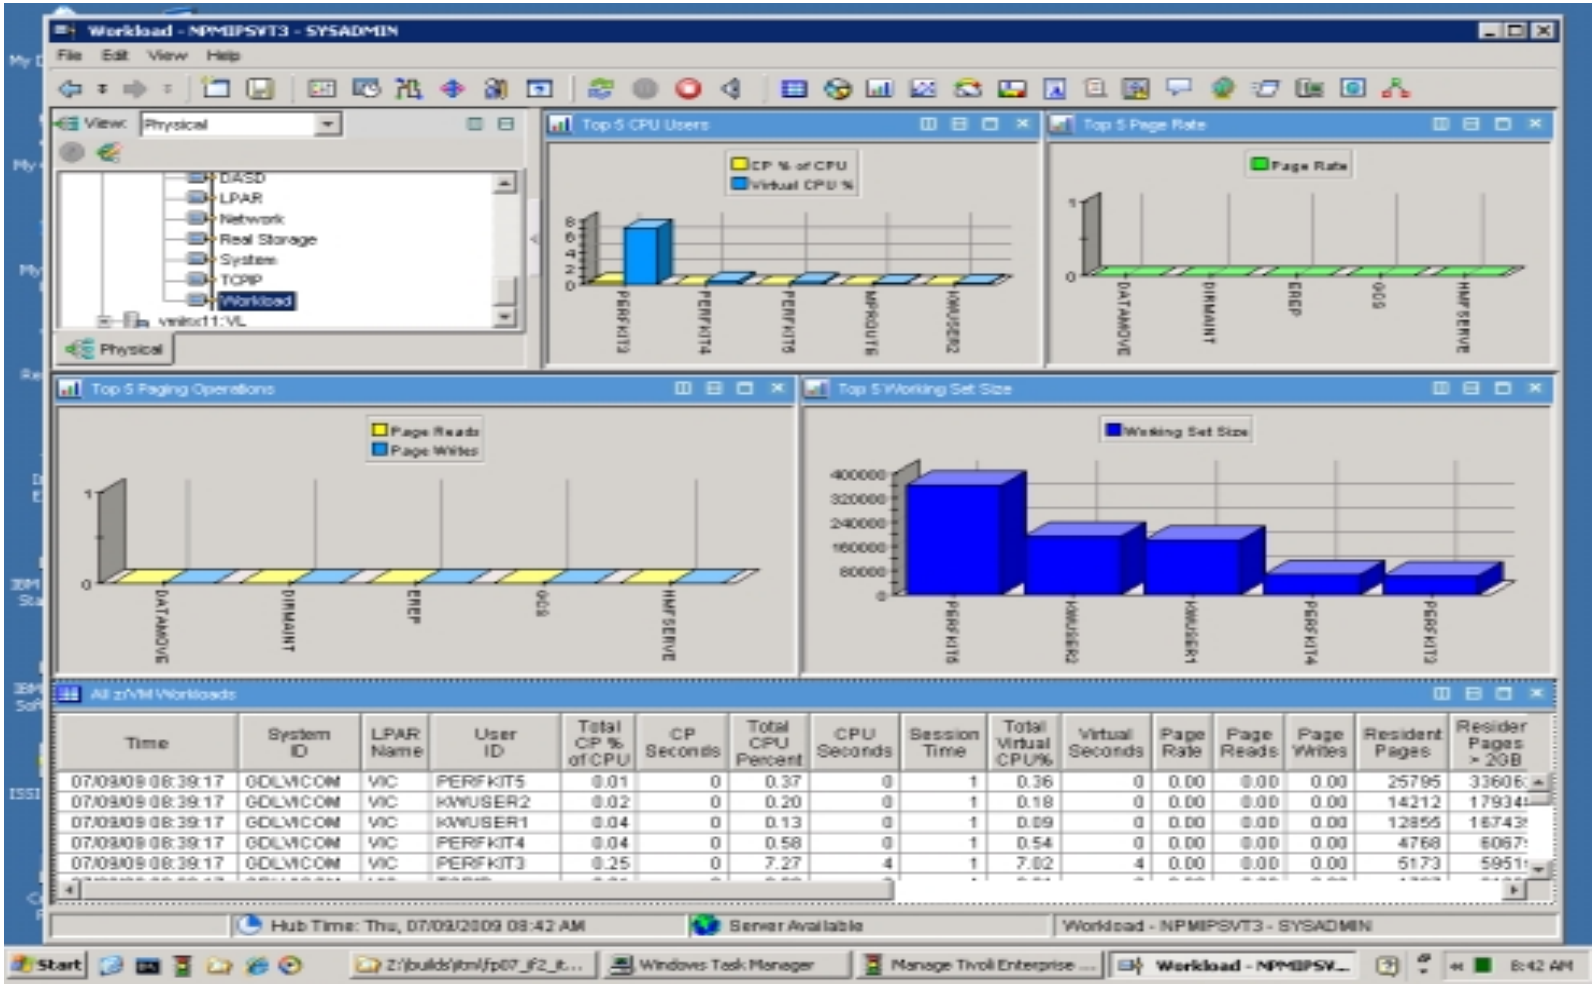

Scenario—System Running Slowly

System is running slowly. Check Workload workspace to see if any particular user is hogging the CPU.

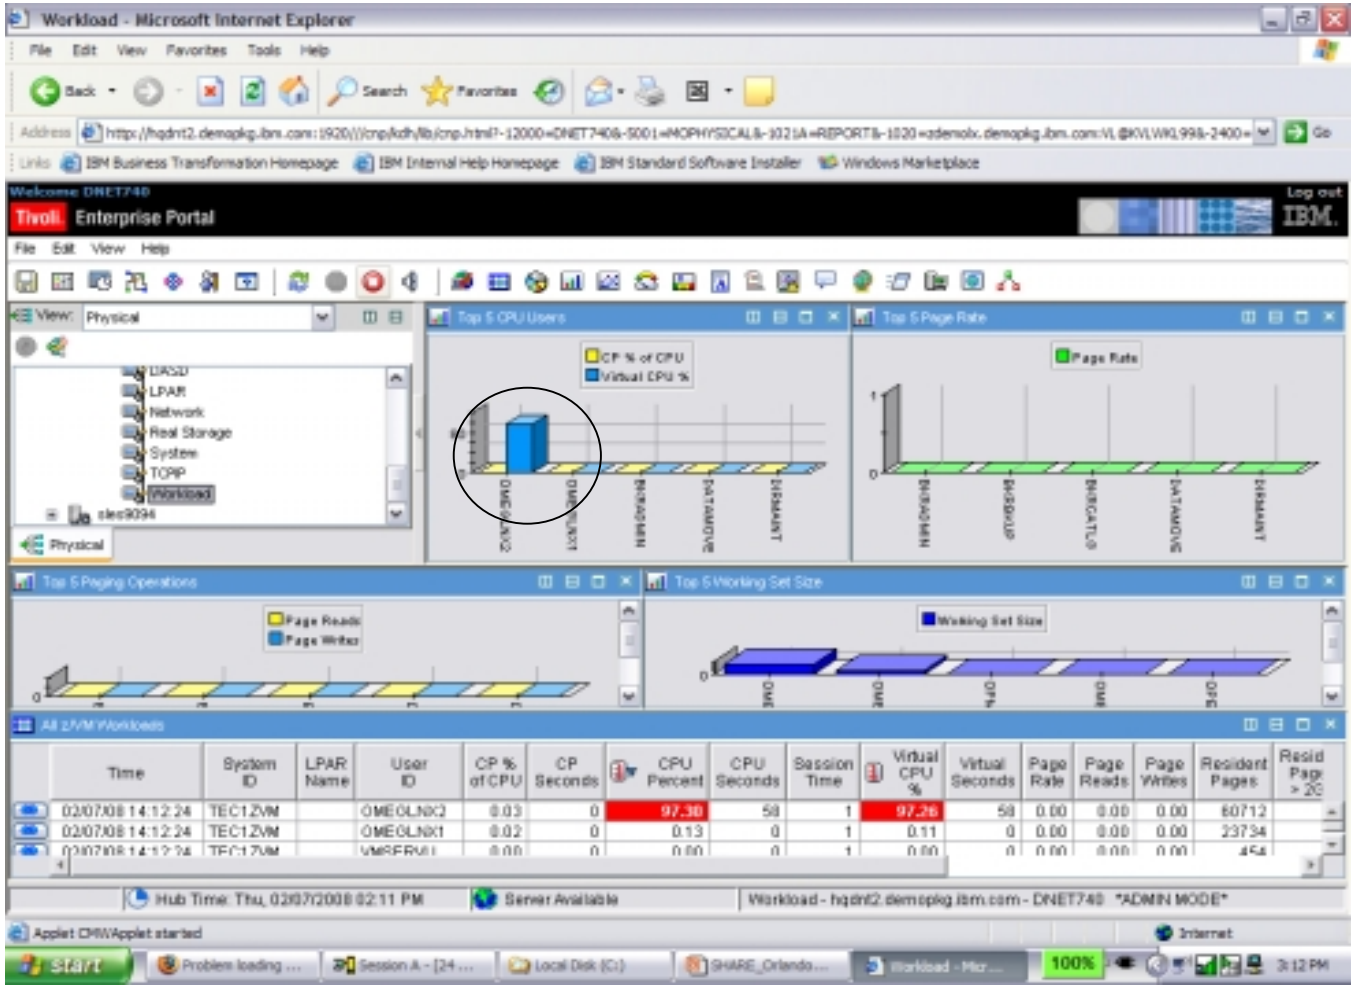

Scenario—System Running Slowly (cont)

Predefined Link to take You directly To the Process workspace

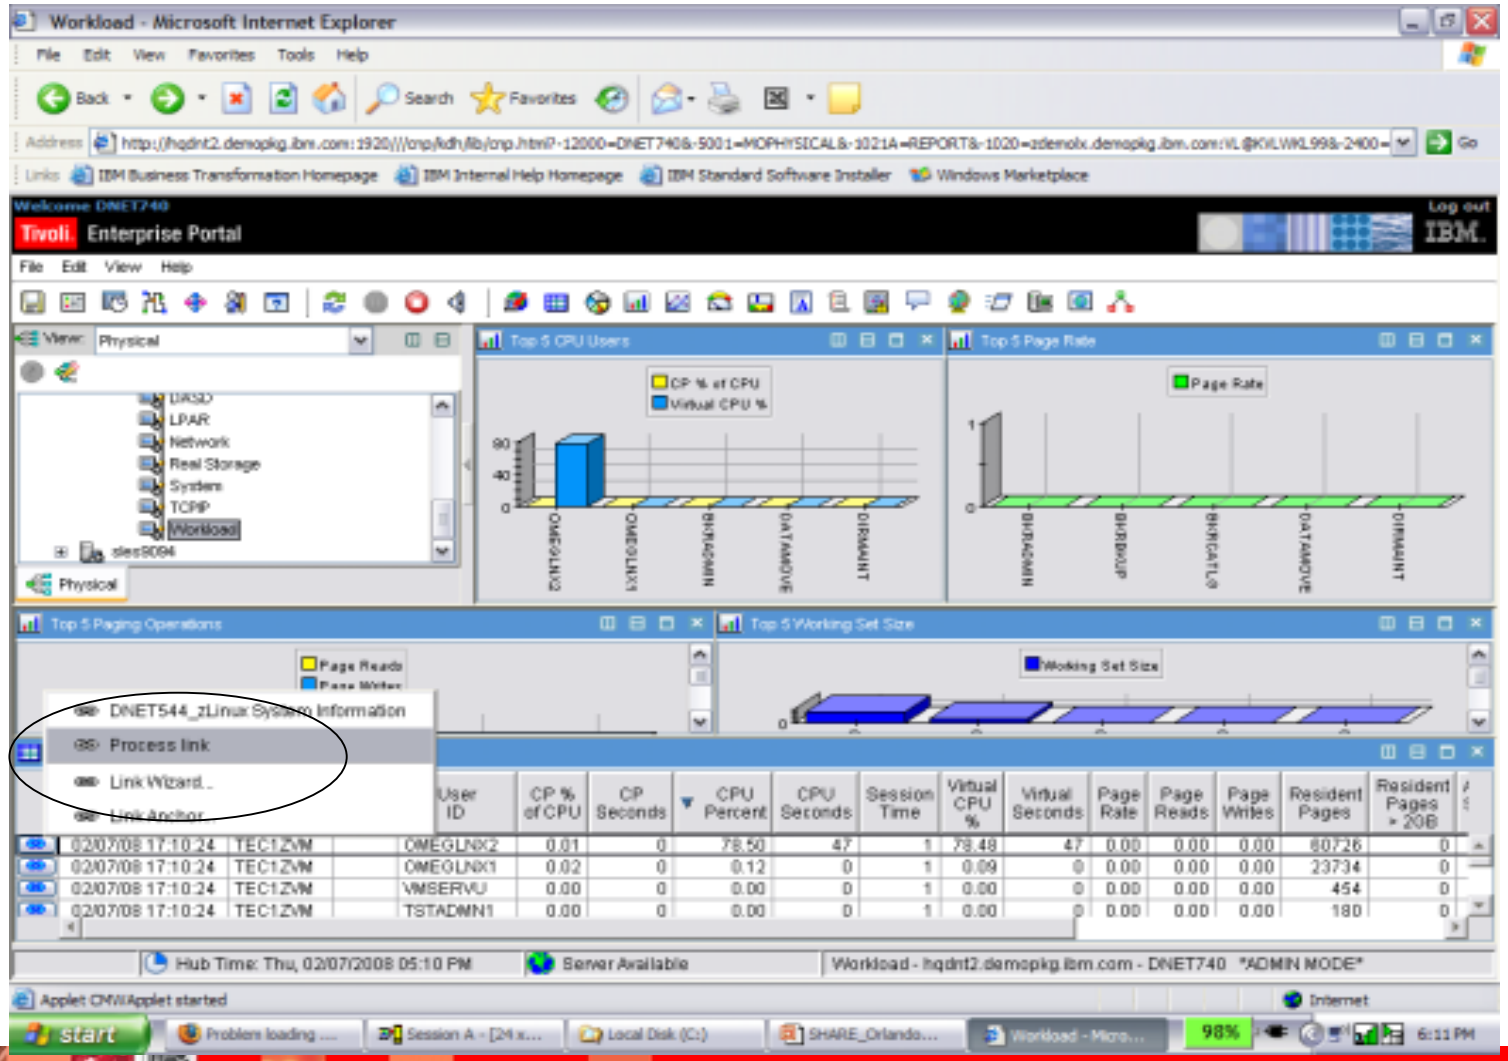

Scenario—System Running Slowly (cont)

See if there is a process which is using too much CPU

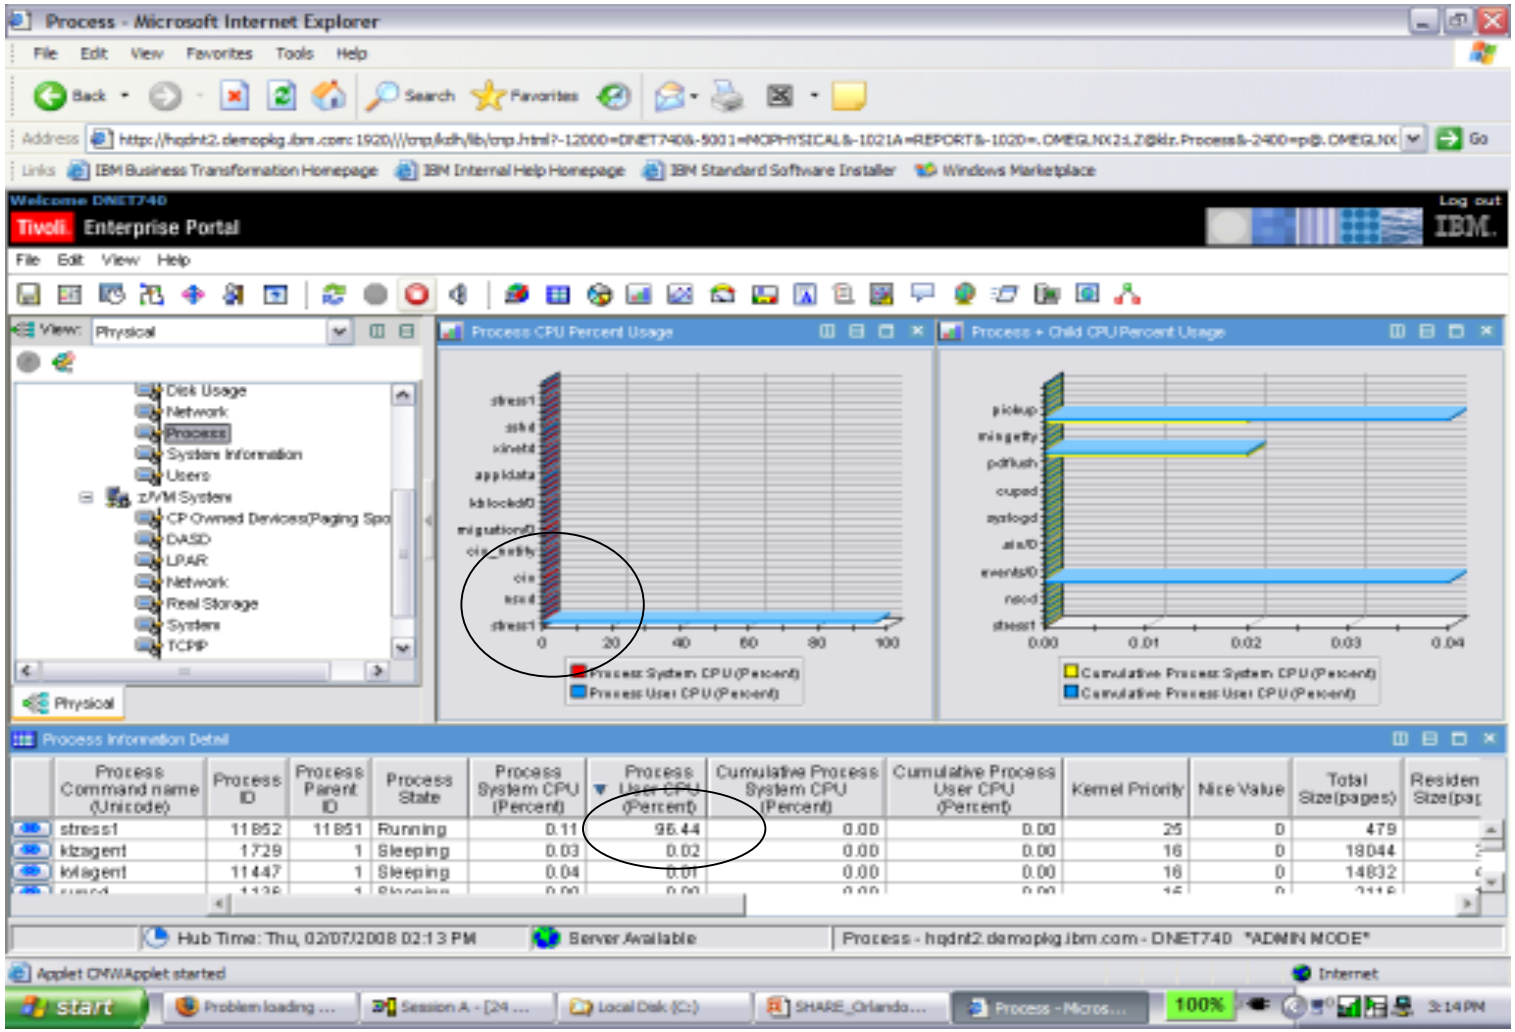

Scenario—System Running Slowly (cont)

You can issue a Take Action command to stop the offending process

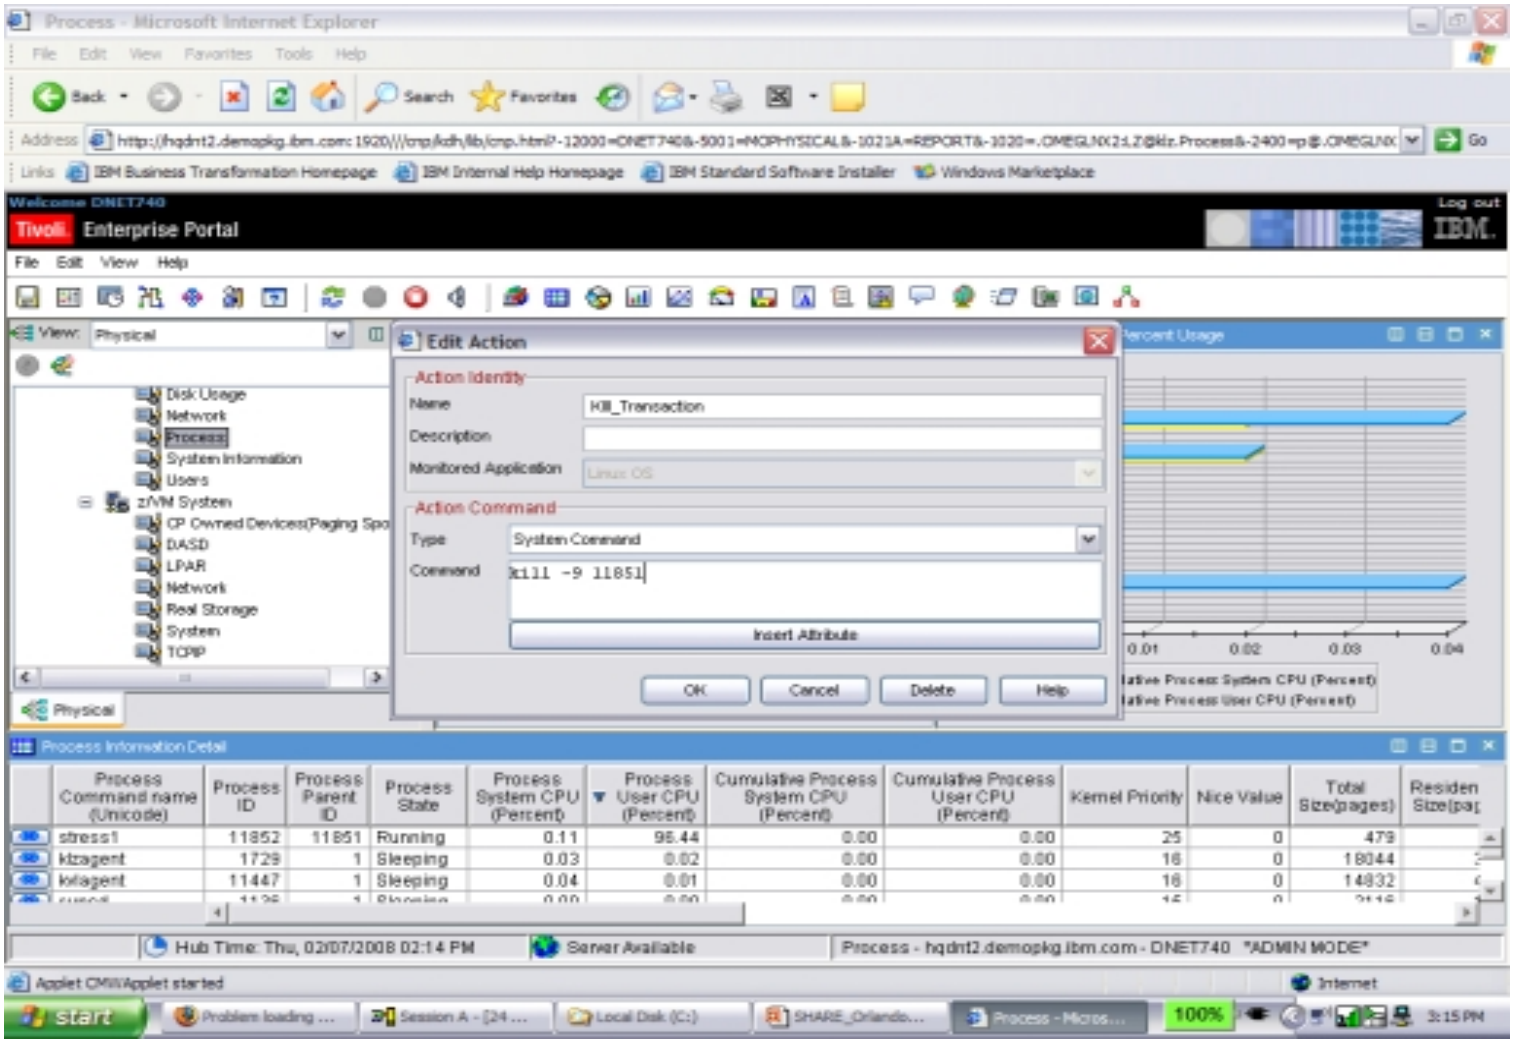

Operations Manager for z/VM

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

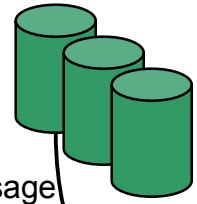

Monitor spool usage

Console monitoring

Service Virtual Machine being monitored

Operations Manager for z/VM

Take action

Service Virtual Machine being monitored

Console monitoring

Schedule tasks

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 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

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

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

OMEGAMON XE and Operations Manager for z/VM

Process Flow

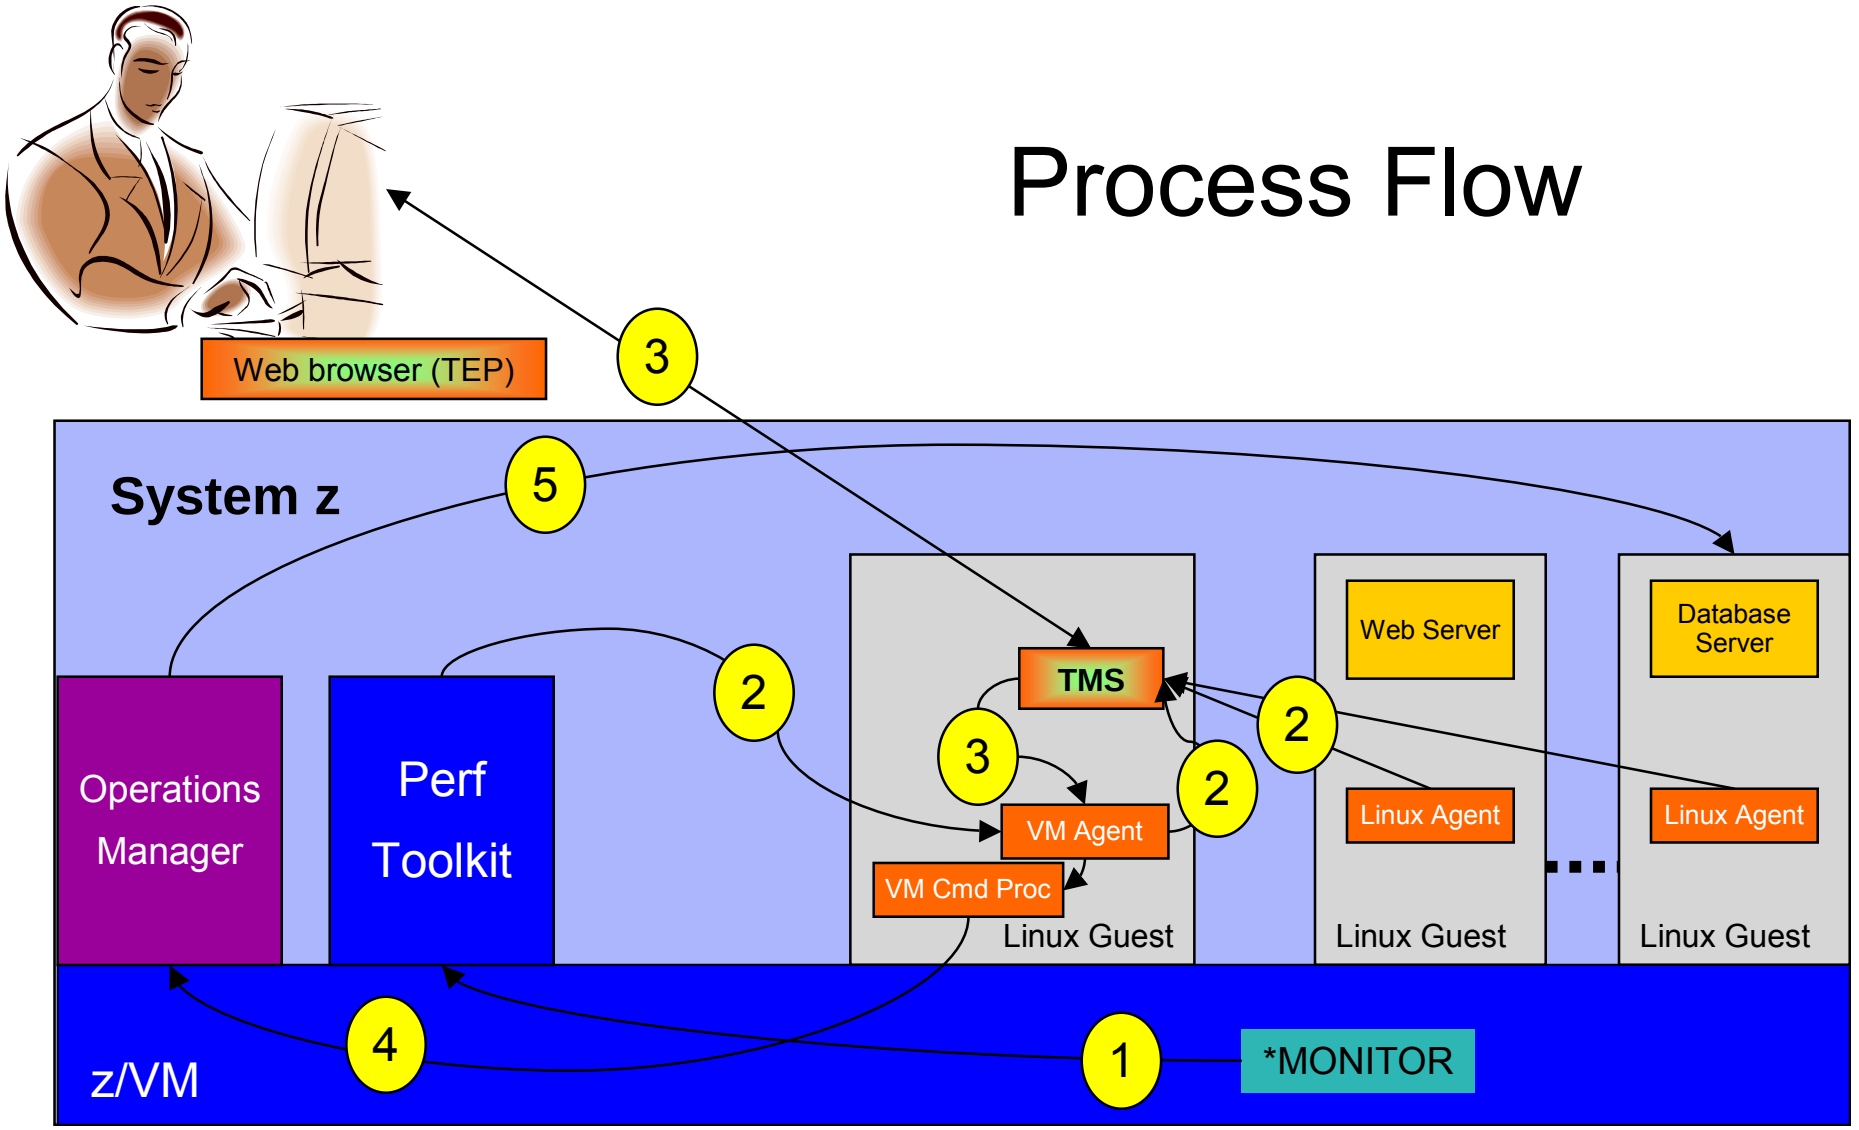

Adjusting resources for a Linux guest

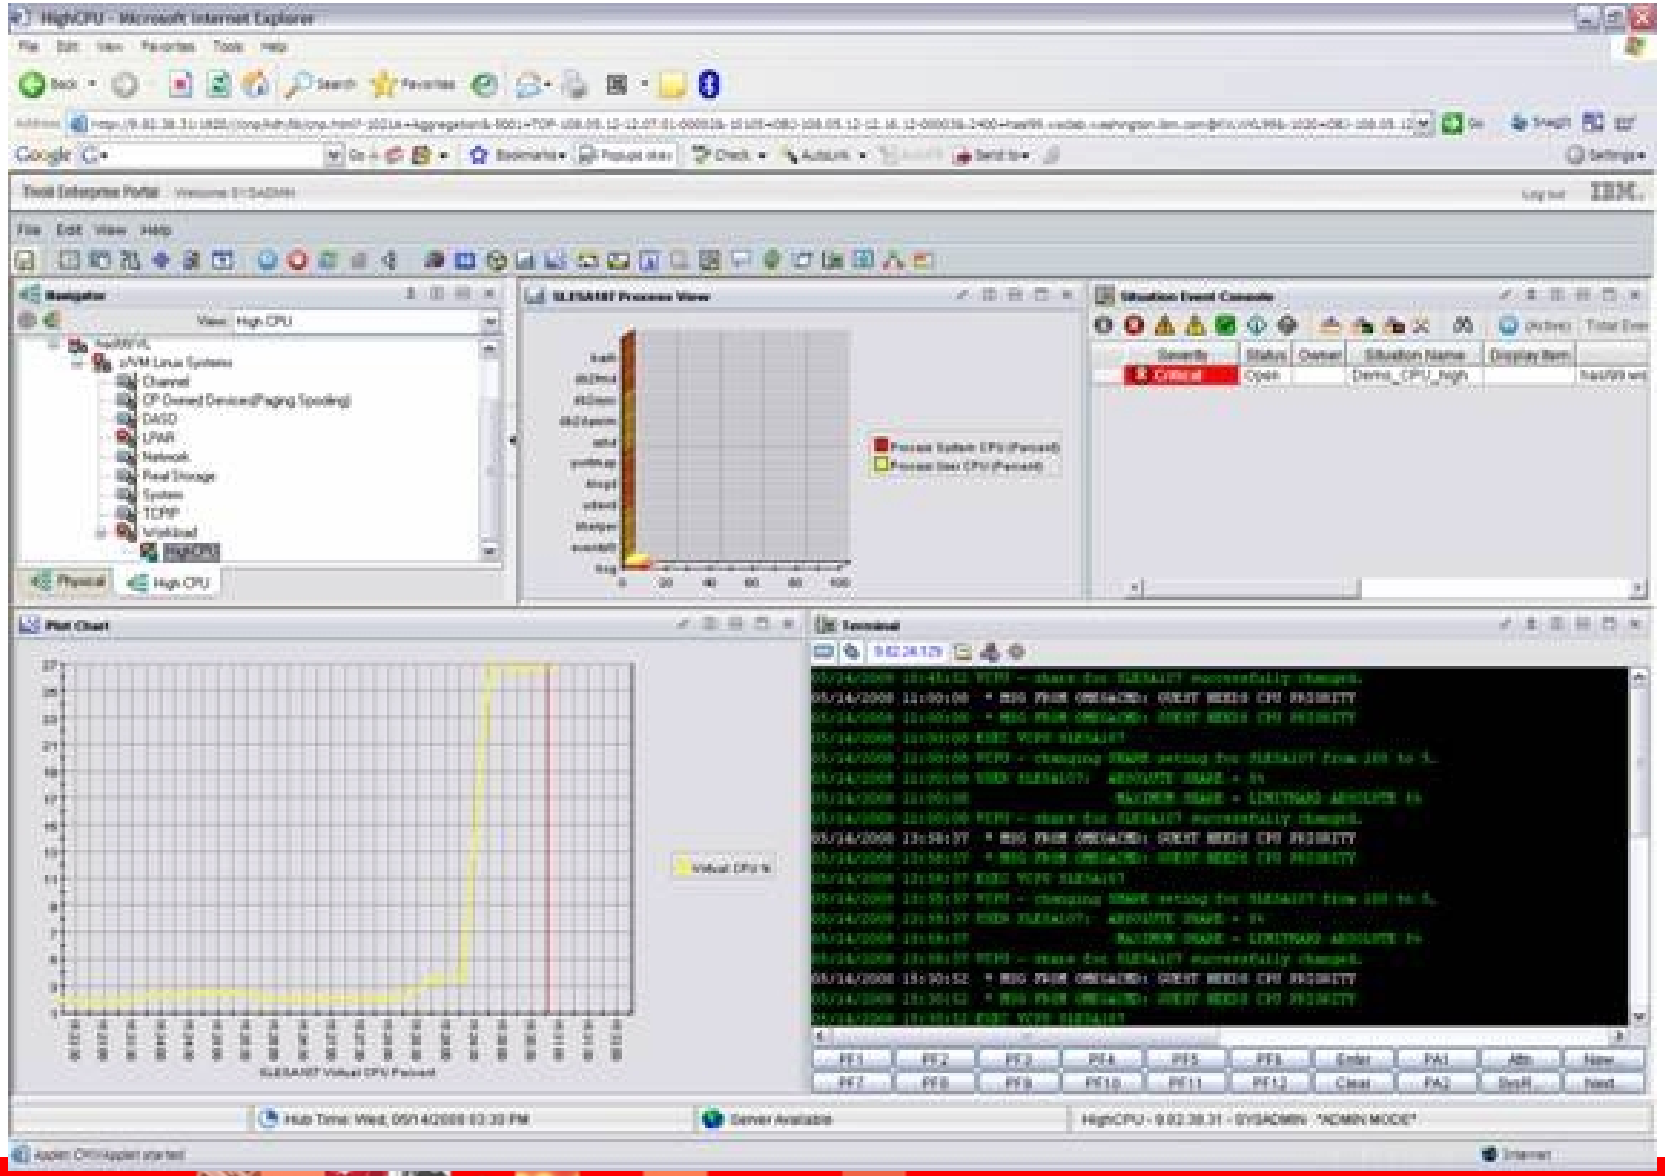

Operations Manager Configuration

- Define a rule to look for the message from OMEGAMON:

```
DEFRULE NAME(GUSTCPU),+
  MATCH(*GUEST NEEDS CPU PRIORITY*),+
  ACTION(GUESTCPU),+
  PARM(SLESA107)
```

- Define actions to highlight the message and call an EXEC to adjust CPU, passing the guest name

```
DEFACTN NAME(GUESTCPU),+
  INPUT(AHI),+
  NEXTACTN(GUSTCPUB)
*
DEFACTN NAME(GUSTCPUB),+
  COMMAND(EXEC VCPU &P),+
  ENV(LVM),+
  OUTPUT(LOG)
```

OMEGAMON Configuration

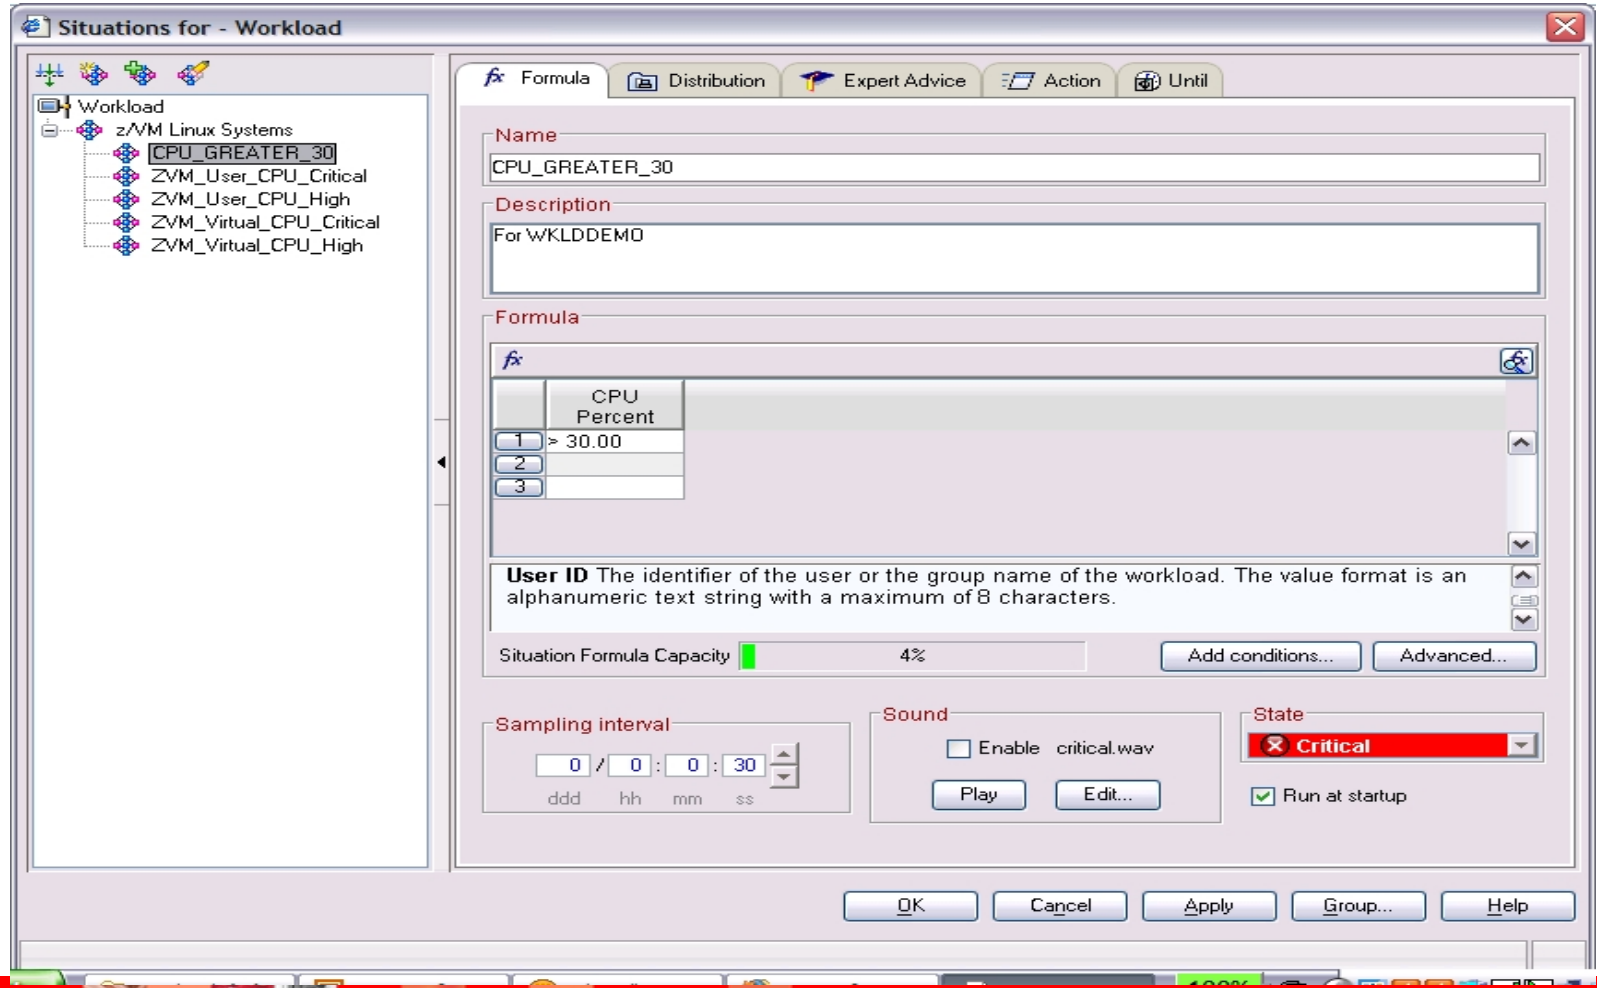

OMEGAMON Configuration

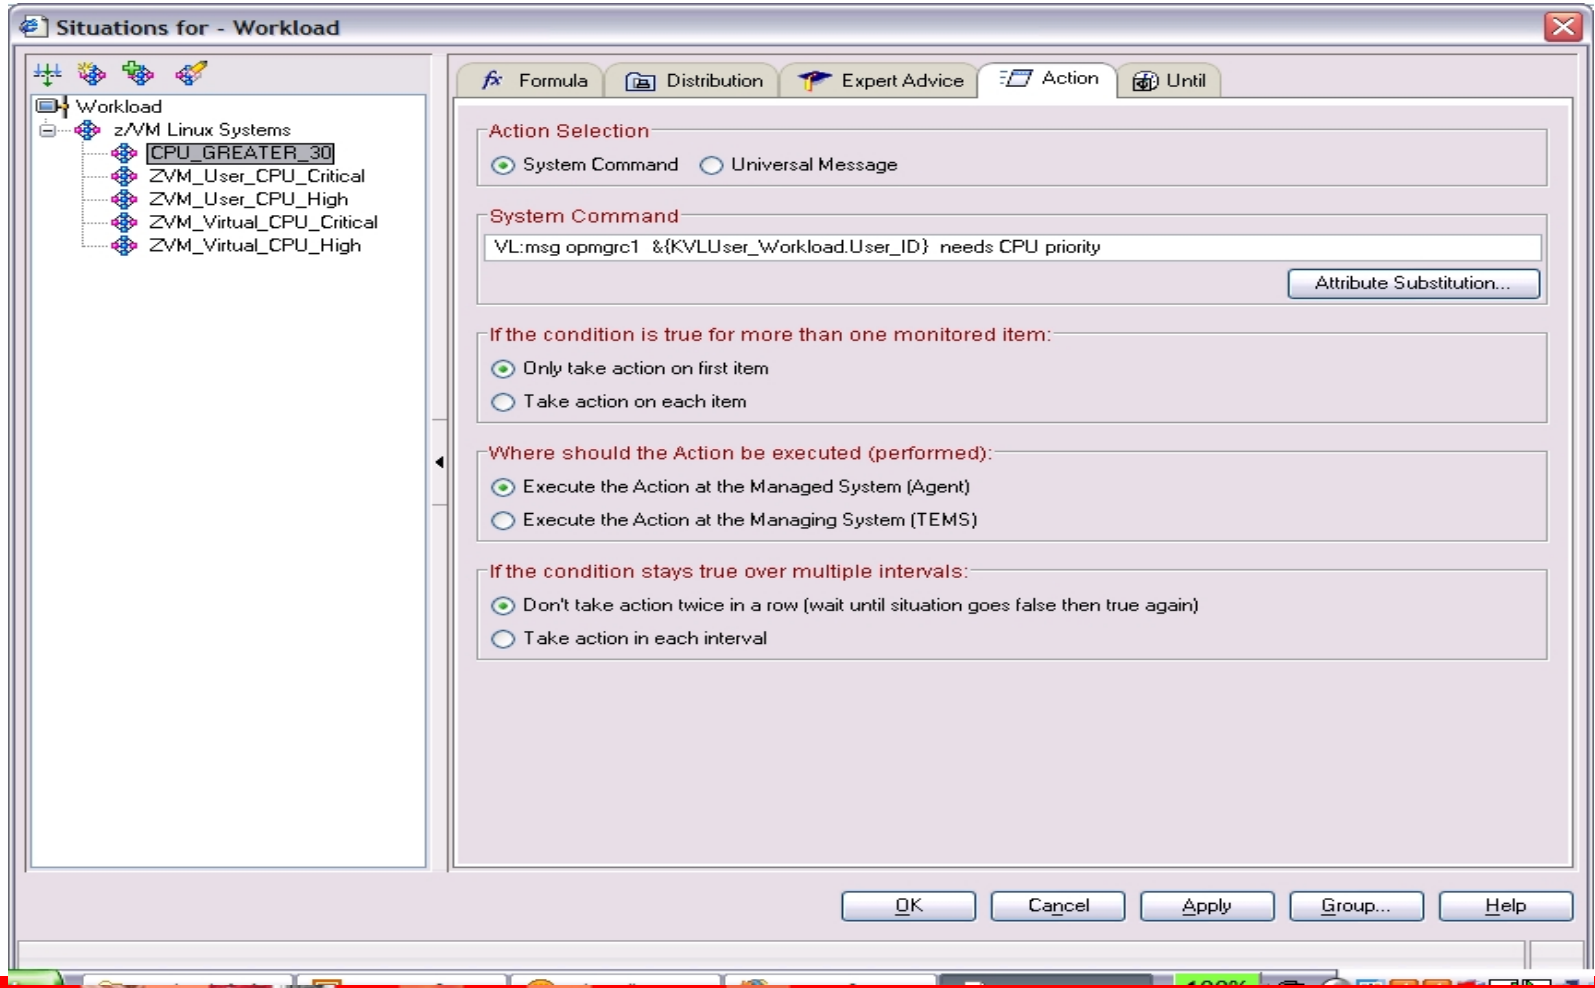

Additional Automation Scenarios

- Automating and scheduling the backup of a Linux guest
- Send a message or e-mail if spool usage is too high
  - Tools to find and view spool files based on size, owner, and/or date
- Automate the response to a disk full condition
- Take actions based on messages in Linux syslog data

Session 9164: Tuesday at 9:30am

Handouts available online, with screenshots

Related Sessions

- 9102 Introduction to Virtualization: z/VM Basic Concepts and Terms
- 9115 VM Performance Introduction
- 9106 VM Performance Update
- 9166 z/VM Performance Case Studies
- 9122 z/VM Tuning Revisited with Specialty Engines for z/OS
- 9164 Automation and Backup Scenarios for z/VM and Linux on System z
- 5201-5202 Explore the TEP Hands on Lab (Tuesday)

BACKUP

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

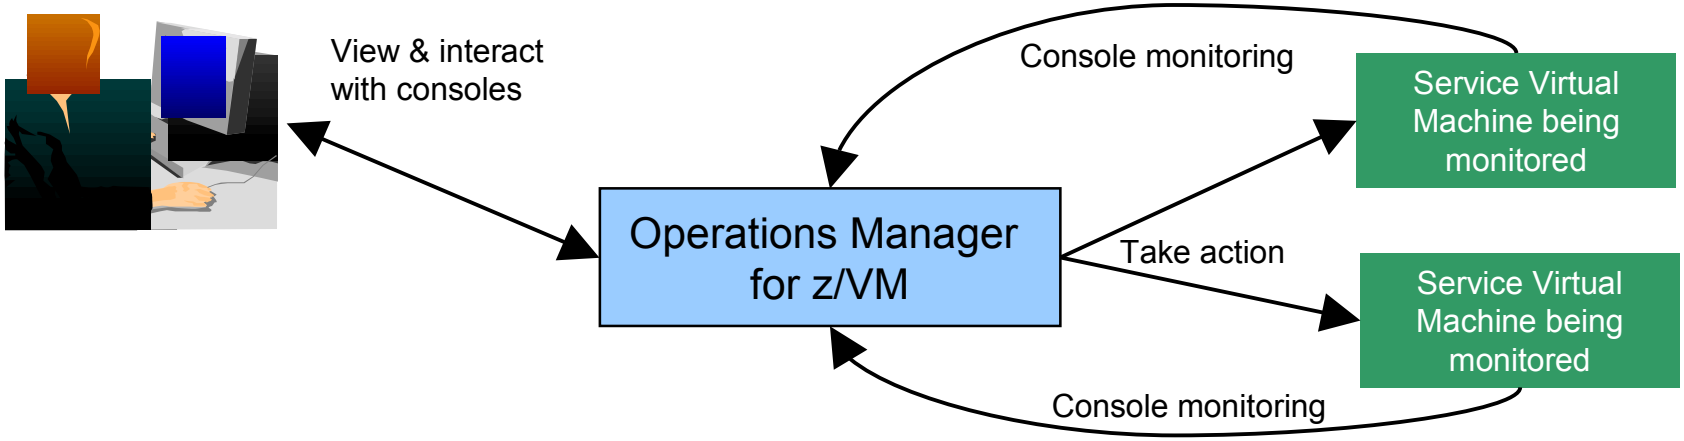

Automation
Routine activities done more effectively with minimal operations staff

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

Operations Manager for z/VM and OMEGAMON XE on z/VM and Linux

- OMEGAMON XE on z/VM and Linux
  - Performance monitoring for z/VM and Linux guests
  - Part of Tivoli Management Services (TMS) infrastructure
- Operations Manager for z/VM
  - Monitor consoles of z/VM service machines and guest user IDs
  - Take actions based on console messages
    - Respond to “take action” requests from OMEGAMON
  - Schedule routine tasks

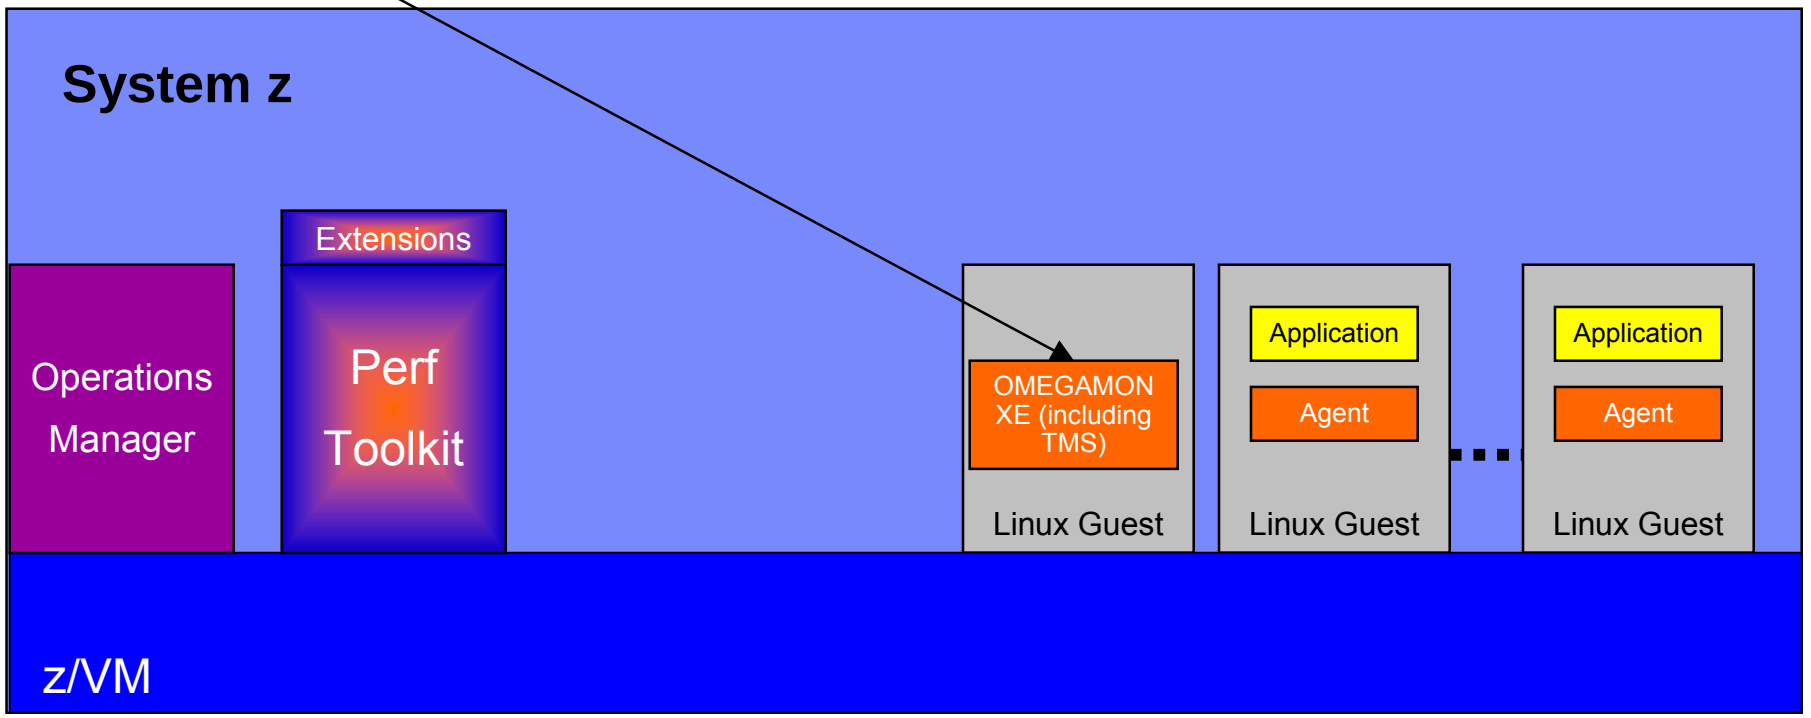

Tape Manager for z/VM

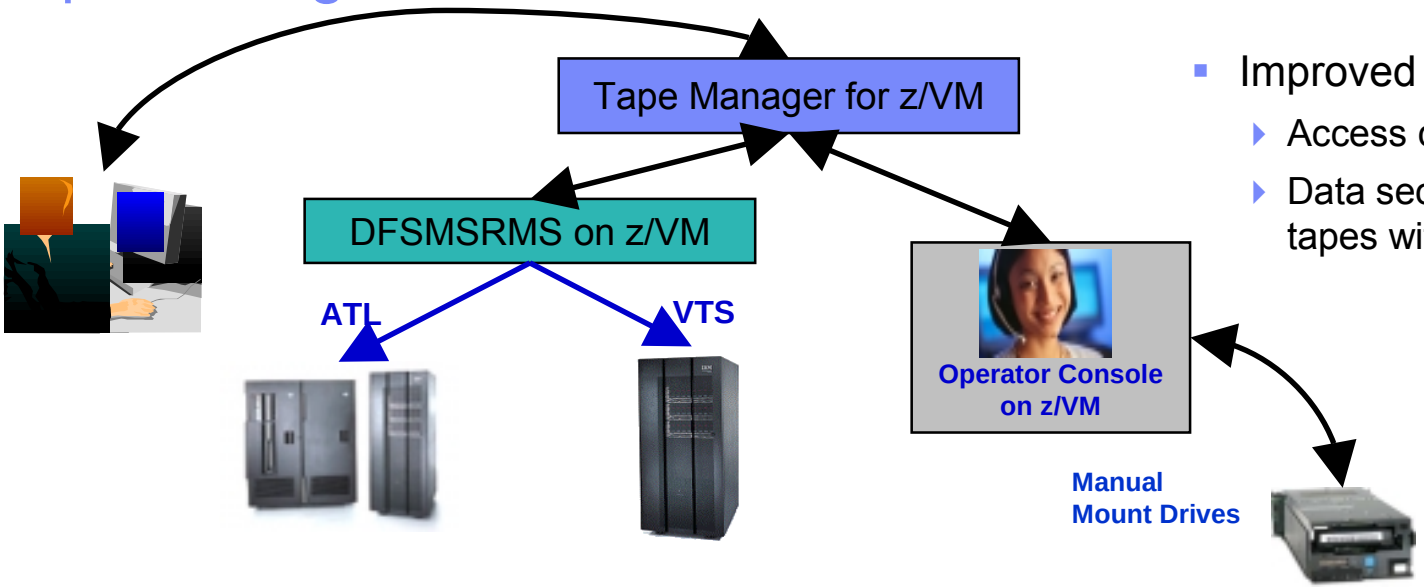

- Improved tape management
  - Access control
  - Data security erase option for tapes with sensitive data

- Effective management of tapes in ATL or VTS
  - Granular access control
  - Expiration processing
  - Notification for low threshold for tape resources
- Improved accuracy of manual tape processing
  - Automated interface to Operator for manual mounts
  - Verification of mount before given to requesting user
    - Internal tape label (verified again when user detaches the tape)
    - Read/Write verification

- Share tape devices with z/OS or other systems
- Integrated management of z/OS and z/VM tapes
  - Optionally use DFSMSrmm on z/OS as the tape catalog for z/VM and z/OS tapes
  - Tapes, access control, and retention managed by the existing RMM catalog
  - Accessible via Tape Manager on z/VM

Backup and Restore Manager for z/VM

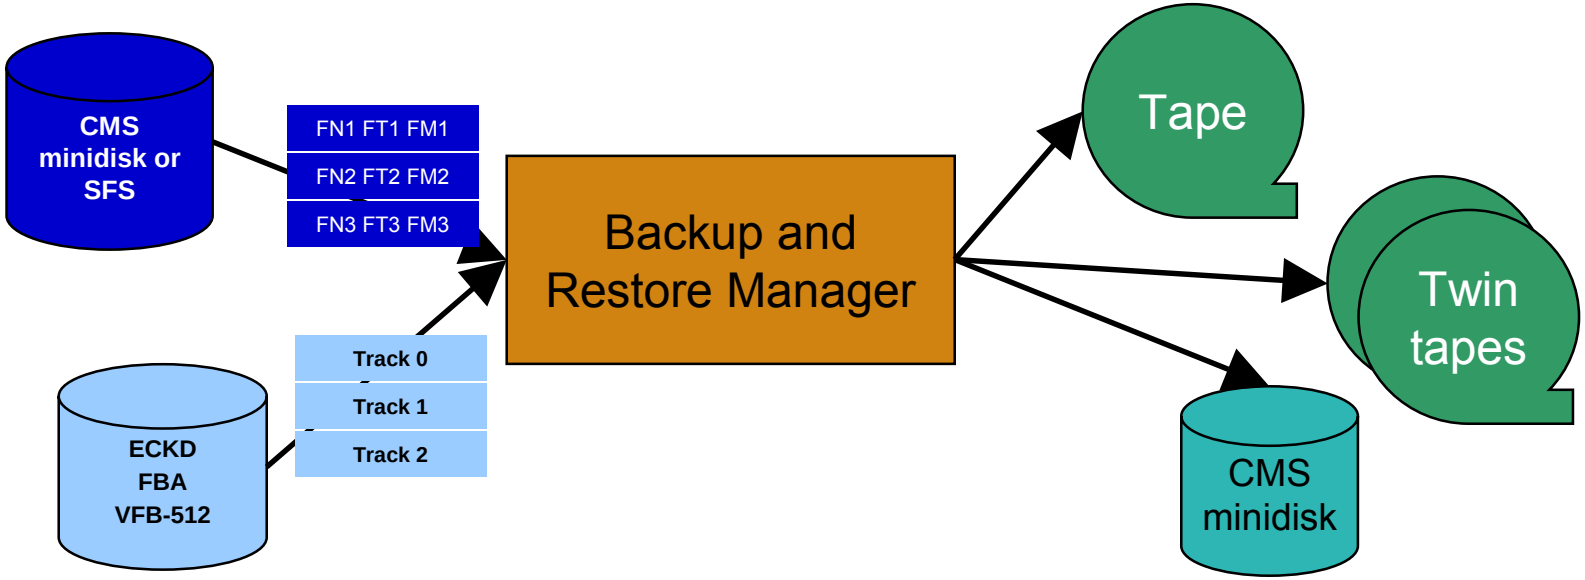

- **System backups available for Disaster Recovery**
- **Guest backups available for restoring to a previous state/level**
- **Backups of user data available for**
  - Restoring to a previous state/level
  - Replacing files accidentally erased
- **Flexible selection of data to back up**
  - Include/exclude minidisks, directories, DASD extents
  - Review of backup job before submission
- **Reduced backup window with concurrent processing**
  - Multiple worker service machines sharing the job
- **Users restore their own data**
  - No administrator interaction required
- **Management of backup data**
  - Retention set as part of the backup job
  - Automatic aging and pruning of the backup catalog
  - View/query the list of expired backups
- **Integration with Tape Manager for z/VM**

Backup and Restore Manager and Linux Guests

Using Backup and Restore Manager with Tivoli Store Manager

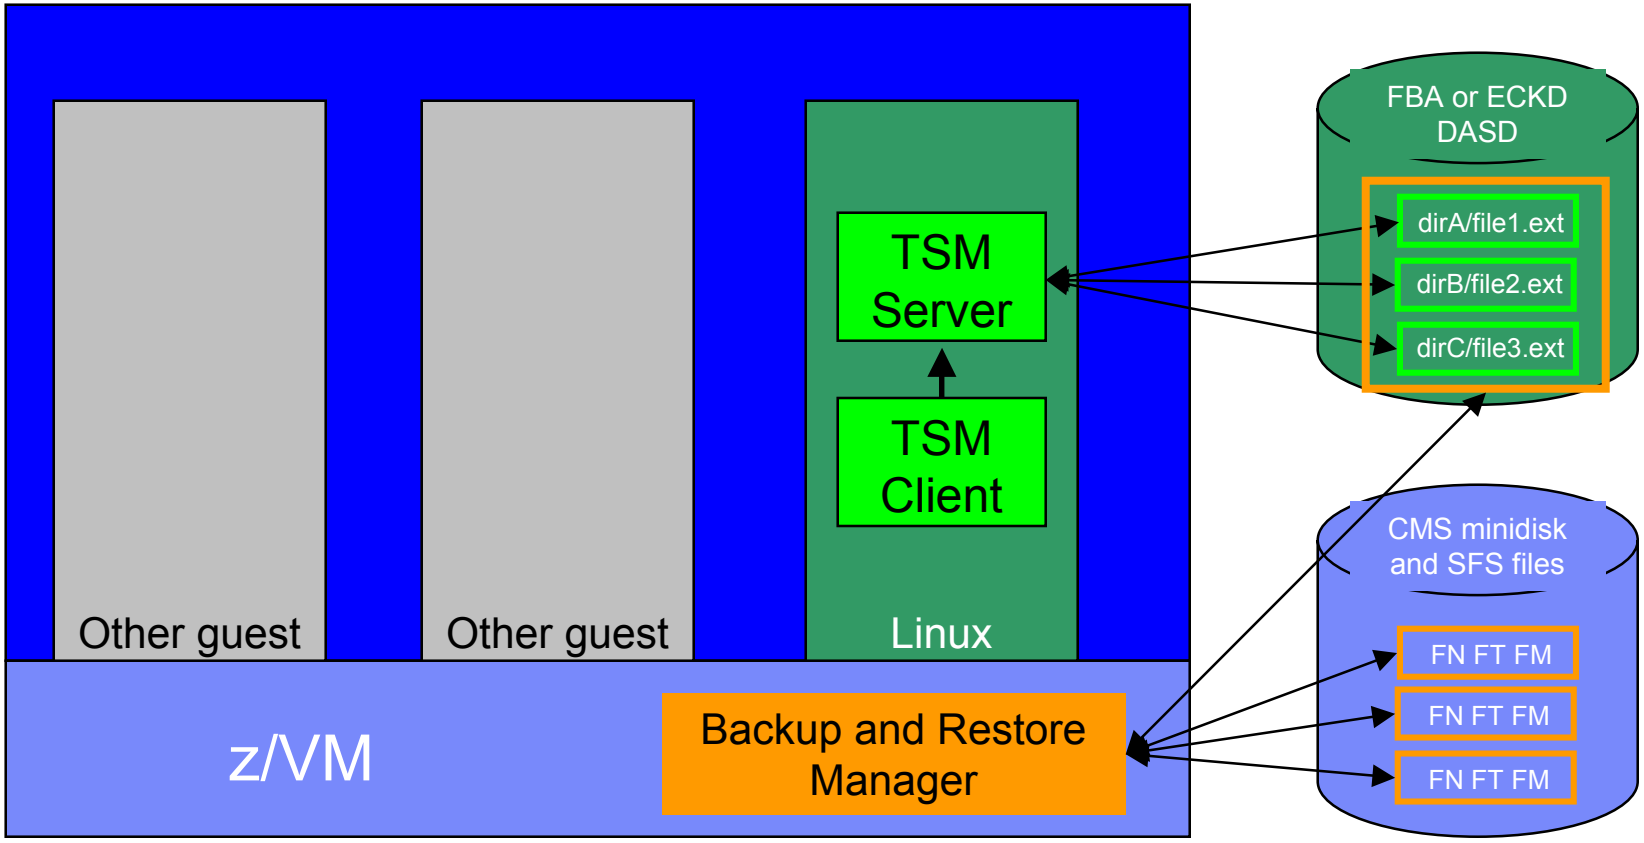

Archive Manager for z/VM

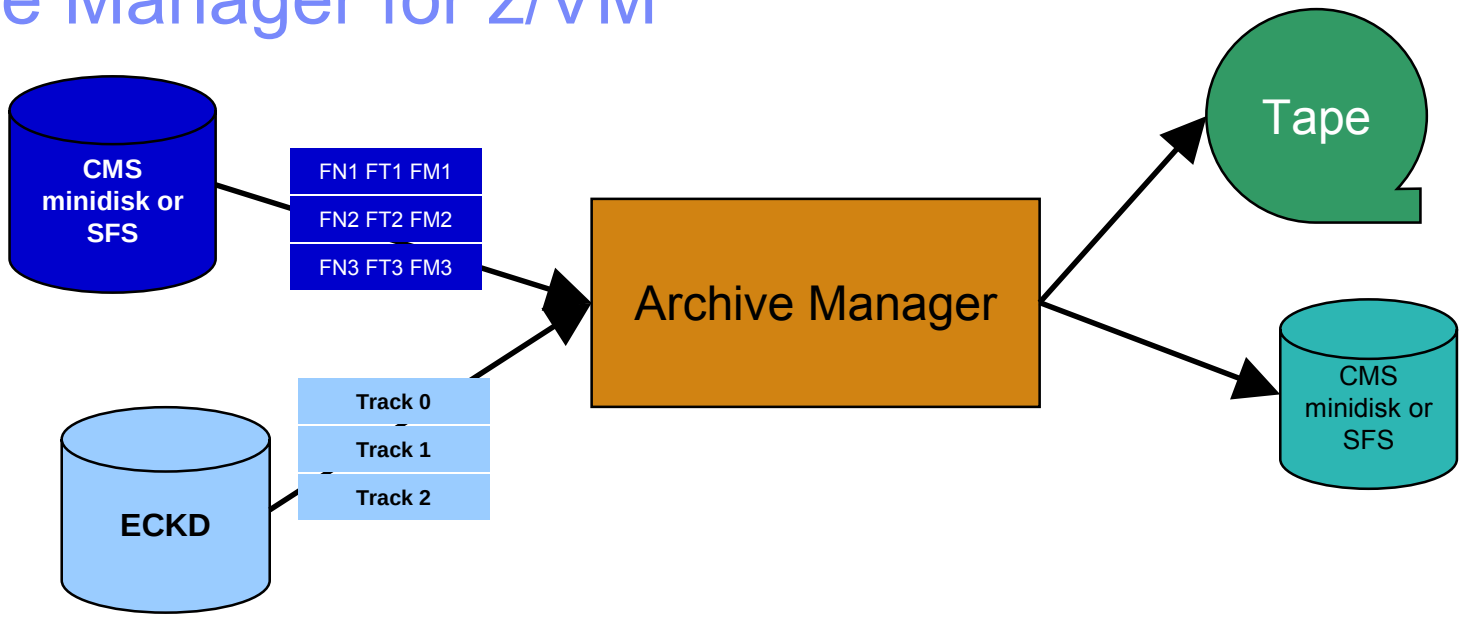

- Improve end user satisfaction and productivity
  - Users manage their own disk space
  - Move infrequently used files to tape or other disk
  - Archive and recall functions are controlled by the user
    - No administrator intervention required
  - Archived data staged to DASD, then tape if applicable
    - Users don't wait for a tape mount for archive request to complete
- Reduce DASD space requirements
  - Archive older files to less expensive storage media
  - Continue to provide users access to the archived data/files
- Control location, retention, and access to archived data
- Integration with Tape Manager for z/VM

For More Information

- Product Web sites
  - http://www.ibm.com/software/stormgmt/zvm/
    - Publications
    - Pre-requisites
    - Announcements
    - Support
- e-mail: Tracy Dean, tld1@us.ibm.com

The following pages contain screenshots and a list of all the data available via the OMEAGMON on z/VM and Linux and the Linux OS agents

z/VM Linux Default Workspace

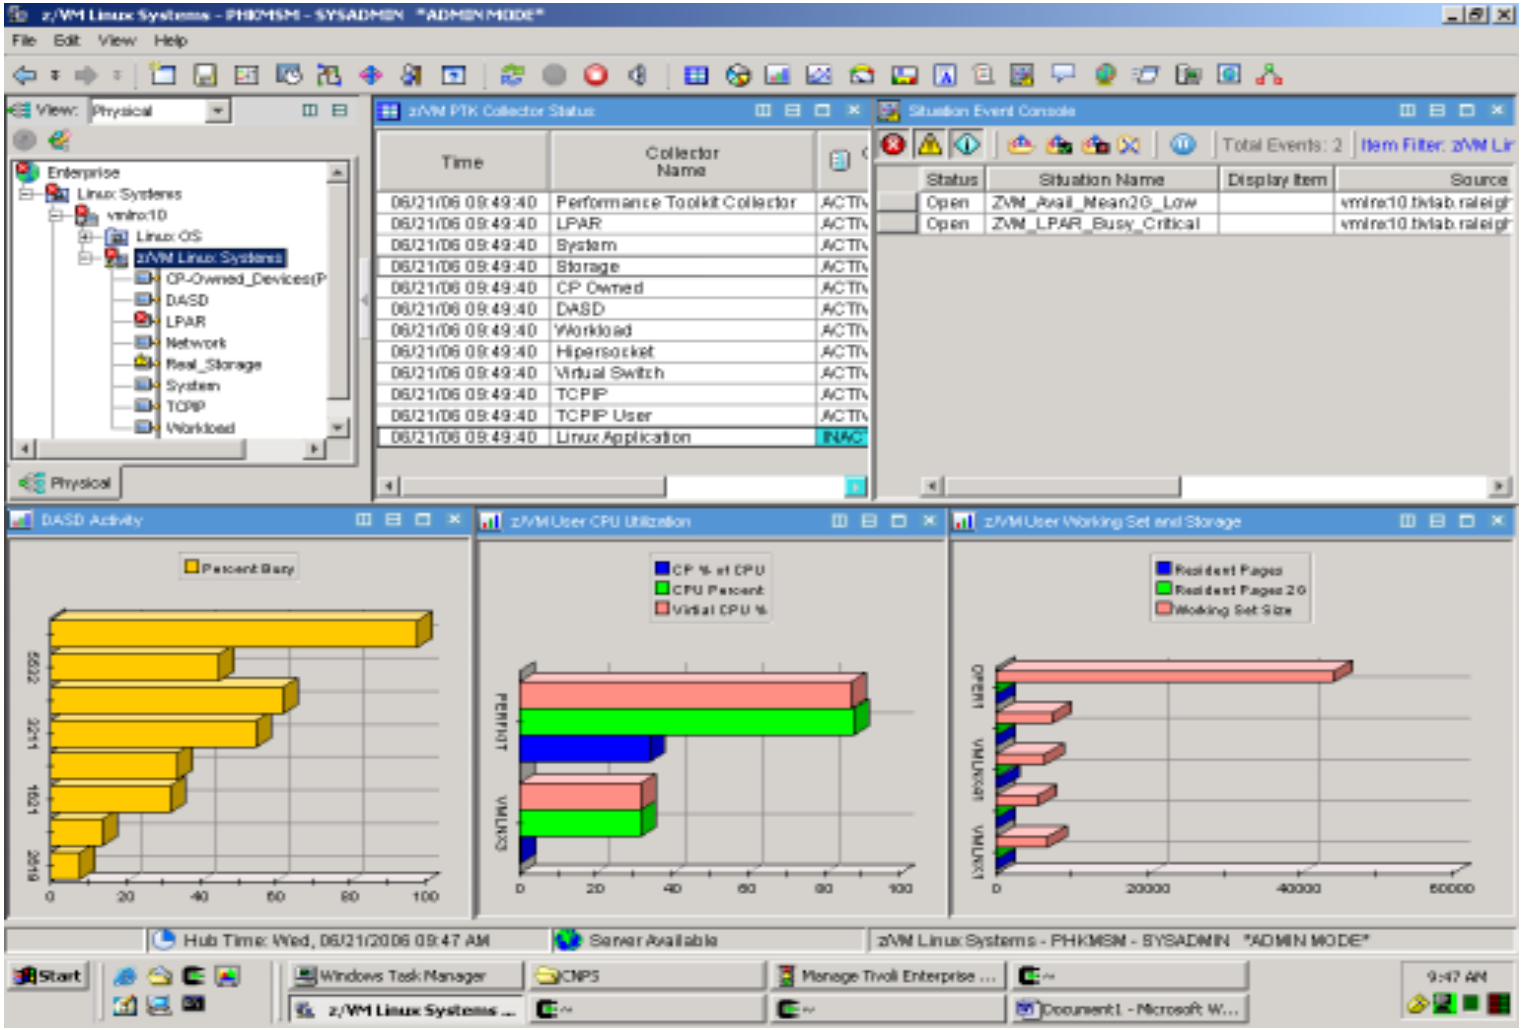

PAGING and SPOOLING Utilization

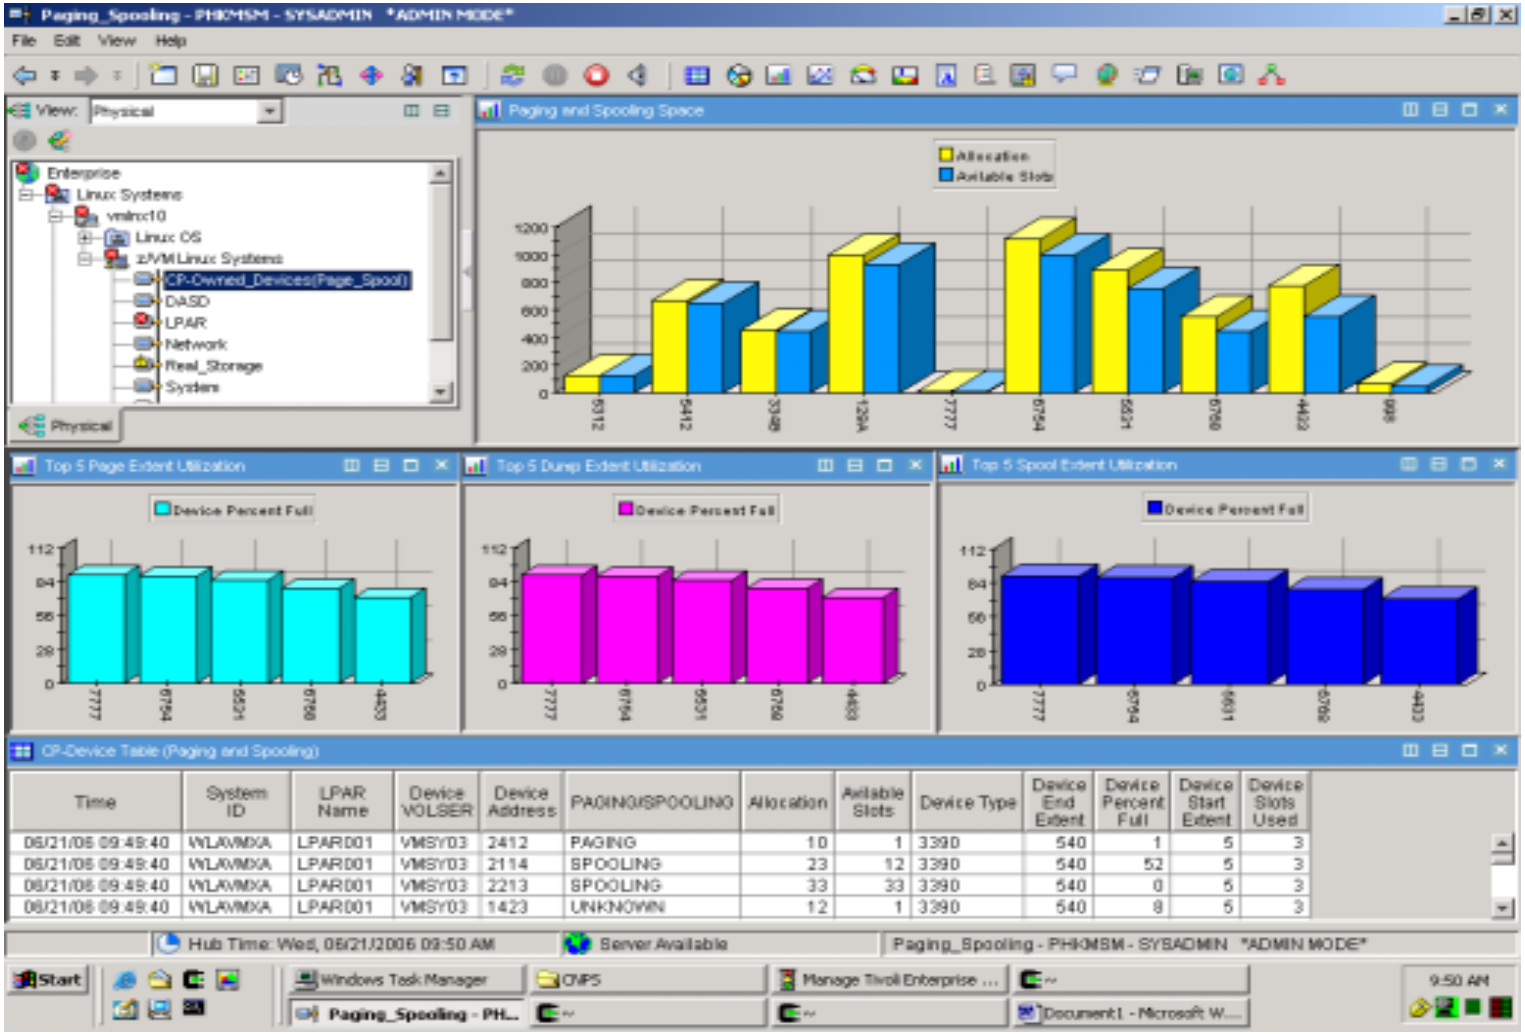

PAGING and SPOOLING Utilization

- **Paging_Spooling Workspace**
  - This workspace displays data on the paging and spooling devices for the whole z/VM system.
- **CP-Device Table**
  - All attributes are collected for the current reporting interval
- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy).
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Real address of CP-owned DASD.
  - Number of cylinders or blocks allocated for the CP-owned extent.
  - Number of slots available on the CP-Owned device at the time the sample was taken.
  - Type of DASD (for example: 3370 or 3380).
  - Primary Purpose of this device.
  - End extent allocated on CP-owned device. (cylinder or block number)
  - Percentage of space on the CP-Owned volume in use at the time the sample was taken.
  - Starting cylinder or block number for the CP-Owned device extent.
  - Number of used slots on the CP-Owned device at the time the sample was taken.
  - Volume serial number of CP-owned DASD.
- **Policies provide advanced automation processes**
- **Historical data option to show previous information**

DASD

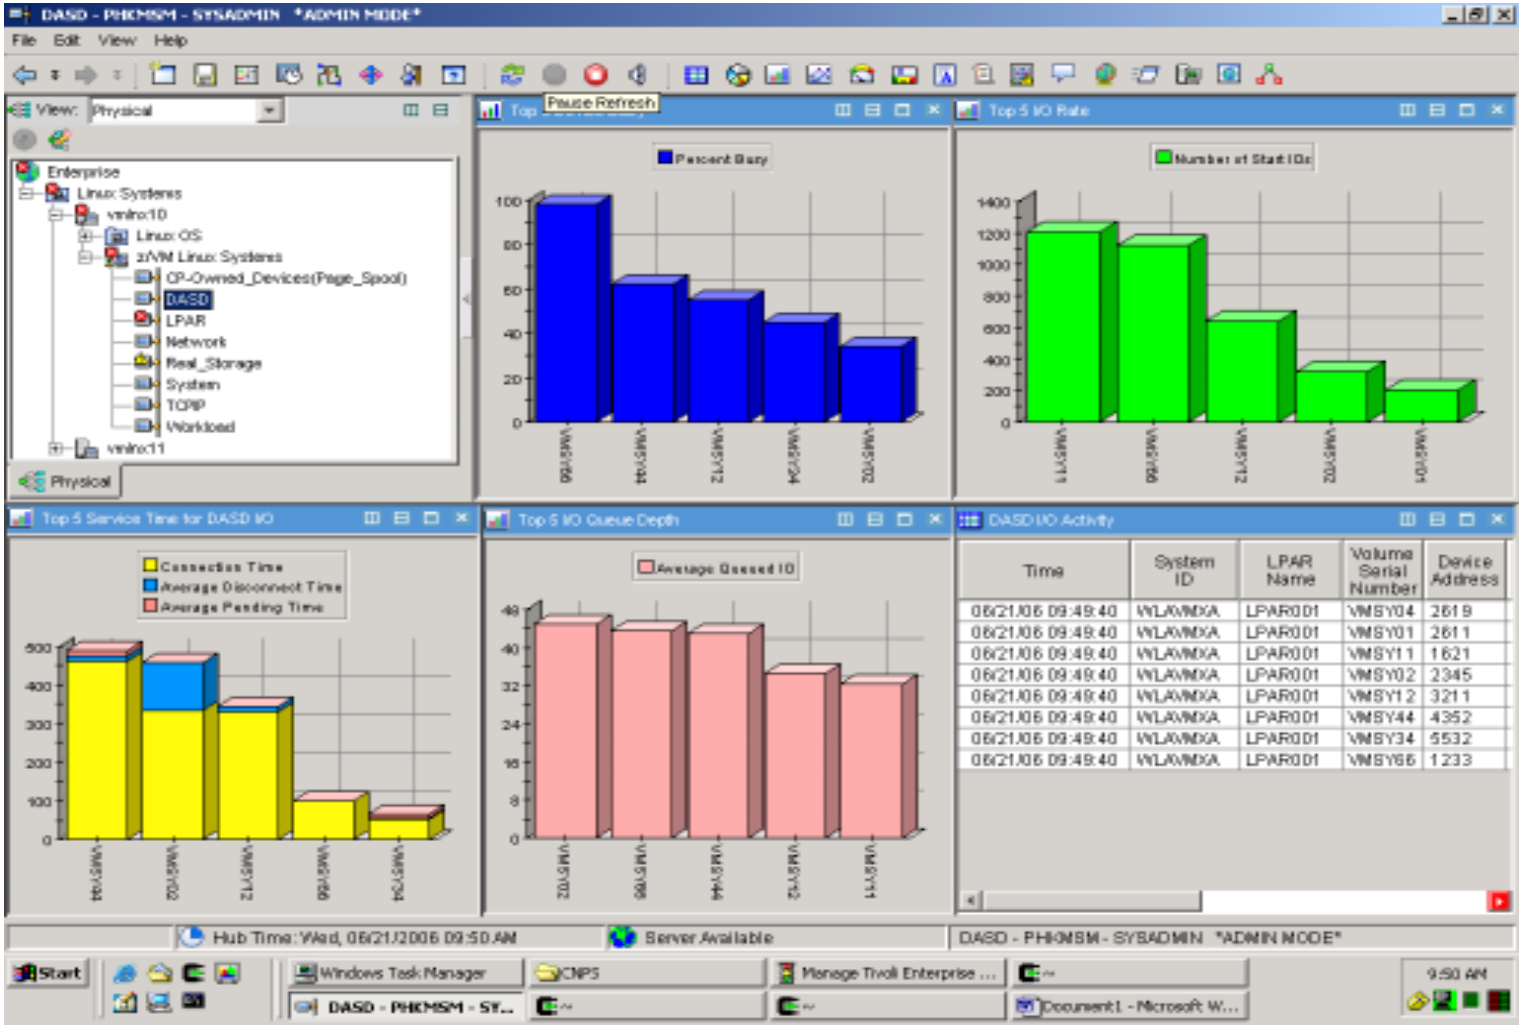

DASD

- **DASD Workspace**
  - The DASD workspace provides several views that show the busiest I/O devices on the overall z/VM system.
- **DASD I/O Activity Table (IO TABLE)**
  - All attributes are collected for the current reporting interval
- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy)
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Real address of device.
  - Type of device (for example: 3370 or 3380).
  - Average time this device was in CONNECT state during the interval.
  - Percentage of time the device was found busy.
  - Average number of I/Os queued on the device.
  - Number of I/Os started on this device.
  - I/Os per second to this device (I/O rate).
  - Average time this device was in DISCONNECT state during the reporting interval.
  - Average time this device was in PENDING state during the reporting interval.
  - Average service time for this device in milliseconds.
  - Volume serial number if DASD device.

LPAR Utilization

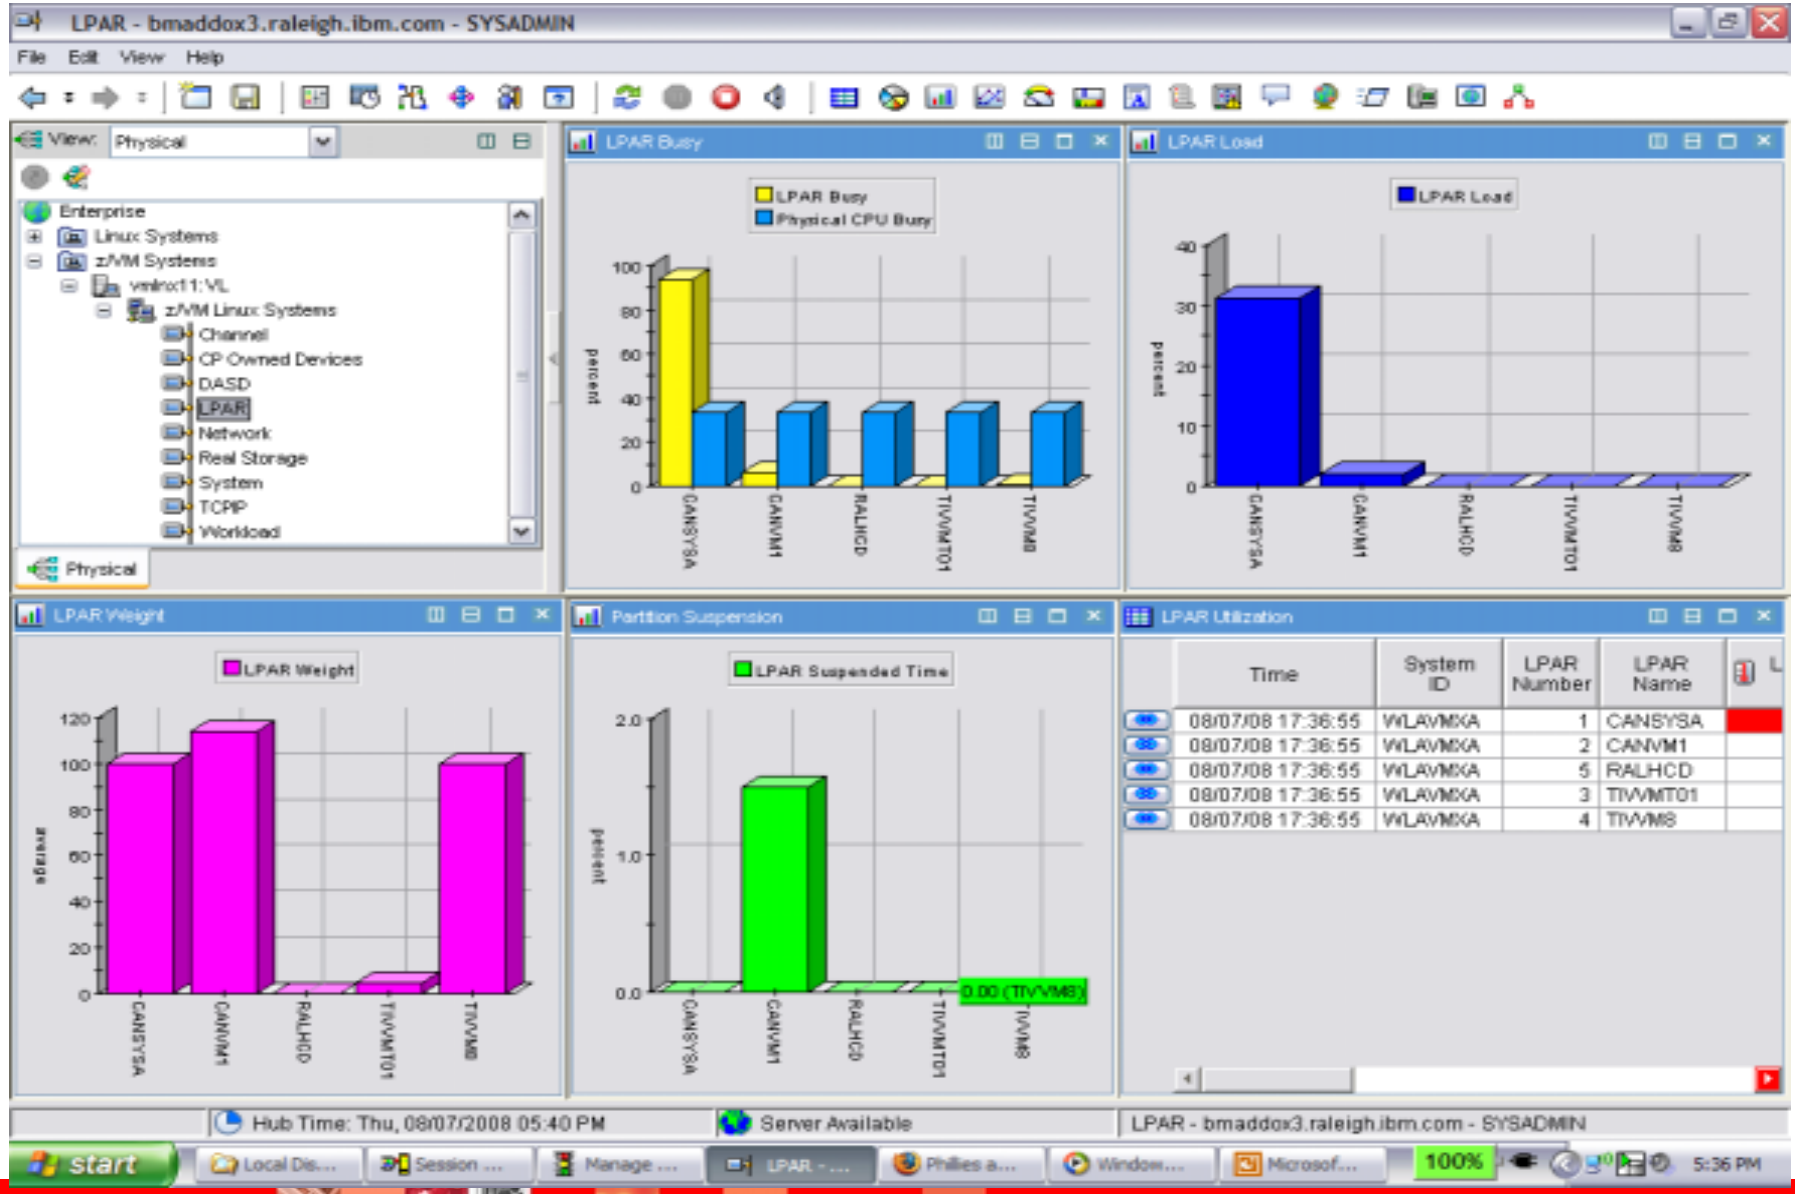

Processor Utilization

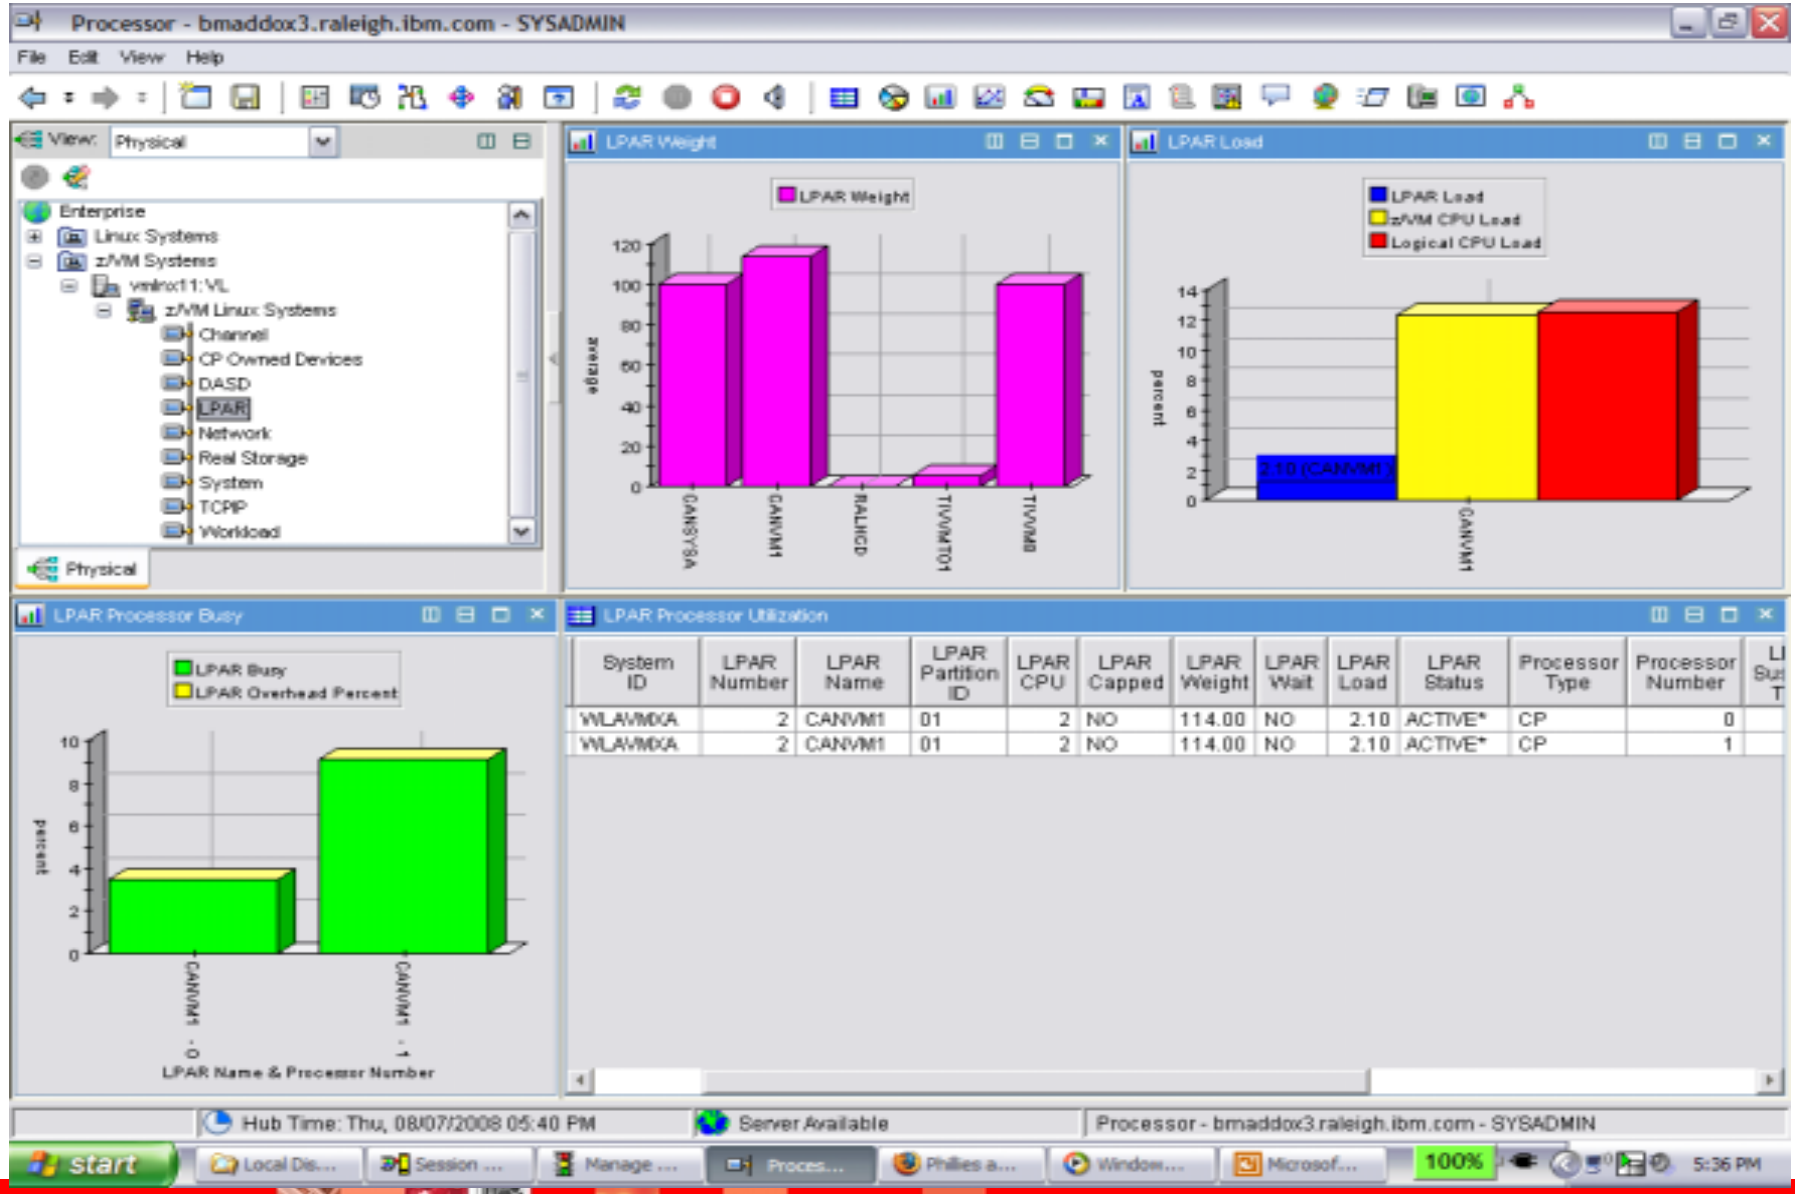

LPAR Utilization

- LPAR Utilization Workspace
  - The LPAR Utilization Workspace provides information about the overall utilization of the system complex. The LPAR Workspace is connected to the LPAR entry on the Navigator.

- **LPAR Utilization table (Data is taken from the IRA LPAR Table)**
  - All attributes are collected for the current reporting interval

- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy).
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Assigned name of the logical partition.
  - Utilization of the system based on the number of logical processors available.
  - Computed as:(Dispatch time / Elapsed time) * Number of LPsTotal amount of time that all of the logical processors for this LPAR were busy during the reporting interval.
  - Number of logical processors assigned to this LPAR.
  - Average percentage of elapsed time that logical processors were 'suspended', i.e. could not give service to the guest system due to LPAR management time and contention for real processors, where the 'suspended' time is calculated as the difference between elapsed time and the sum of processor busy time and voluntary wait time for the same processor as seen by the VM system that is active in the partition.
  - Amount of logical CPU busy which was due to LP dispatching overhead.

LPAR Utilization (cont)

- Average percentage of elapsed time that the logical processors spent for LPAR management. This information is available only on systems with the LPAR management time facility.
- Status of the logical partition during the reporting interval. Can be ACTIVE or INACTIVE. The partition that was used to collect the LPAR data will have an asterisk (*) appended (for example, 'ACTIVE*').
- Status of the WAIT bit for the logical processors within this LPAR. If any LP has the WAIT bit ON, this field will contain 'YES'.
- Average weight of all logical processors defined for this LPAR. The weight values for dedicated processors will be 1000.
- Utilization of the system based on the number of physical processors available
- The CPU type of the logical processors defined for the partition. Possible values are:
  - CP
  - ICF
  - IFL
  - ZIIP
  - ZAAP
  - Special
  - Unknown.

NETWORK Utilization (Hipersocket)

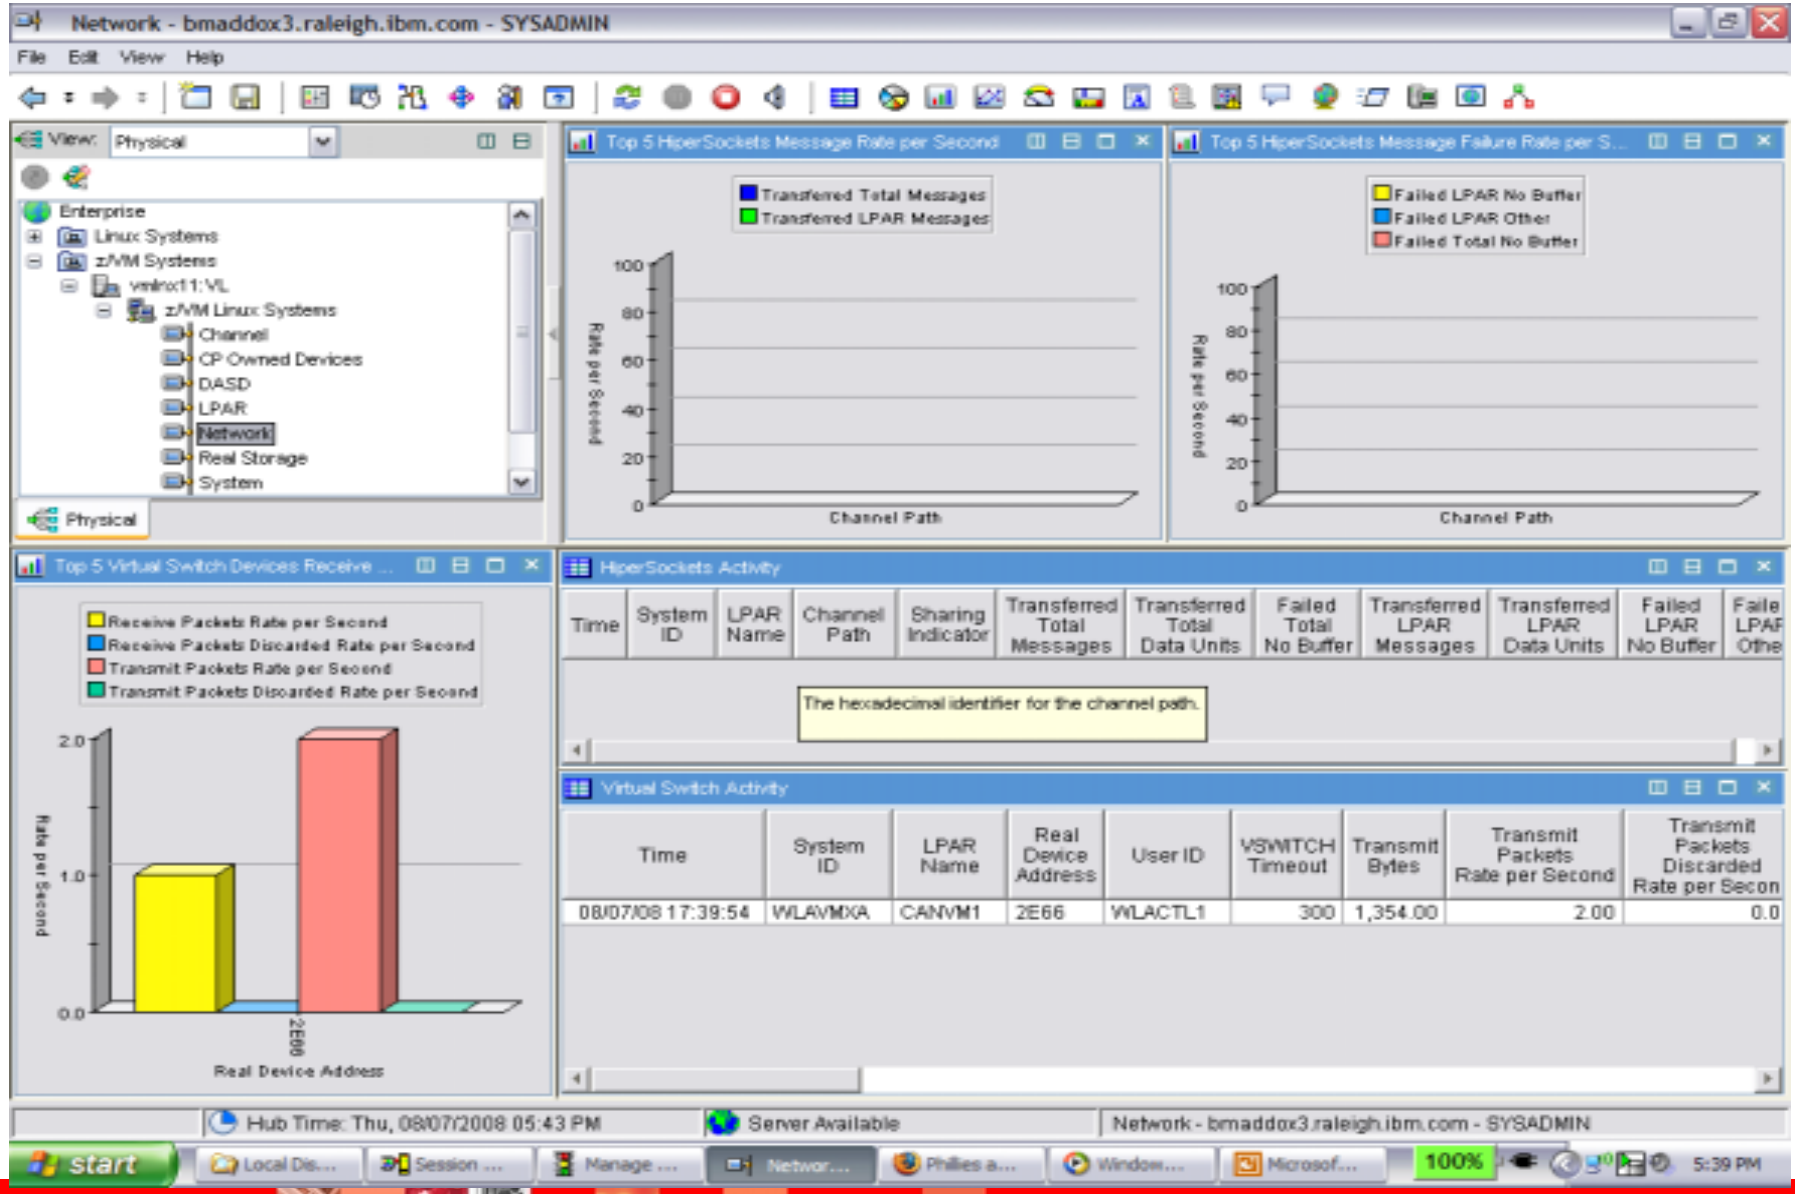

NETWORK Utilization (Hipersockets)

- **NETWORK Workspace**
  - This workspace displays data about the utilization of the hipersocket and virtual switch devices on the z/VM system
- **Hipersocket Utilization (HIPERSOCKET TABLE)**
  - All attributes are collected for the current reporting interval
- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy).
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Hex channel path identifier.
  - Sharing indicator for the channel YES - shared with other LPARs, NO - Dedicated channel.
  - Number of messages sent per second for the whole system.
  - Number of data units sent per second for the whole system.
  - Number of sends per second that failed due to no receiver buffer for the whole system.
  - Number of messages sent per second for this partition.
  - Number of data units sent per second for this partition.
  - Number of sends per second that failed due to no receiver buffer for this partition.
  - Number of sends per second that failed due to other problems for this partition.

NETWORK Utilization (Virtual Switch)

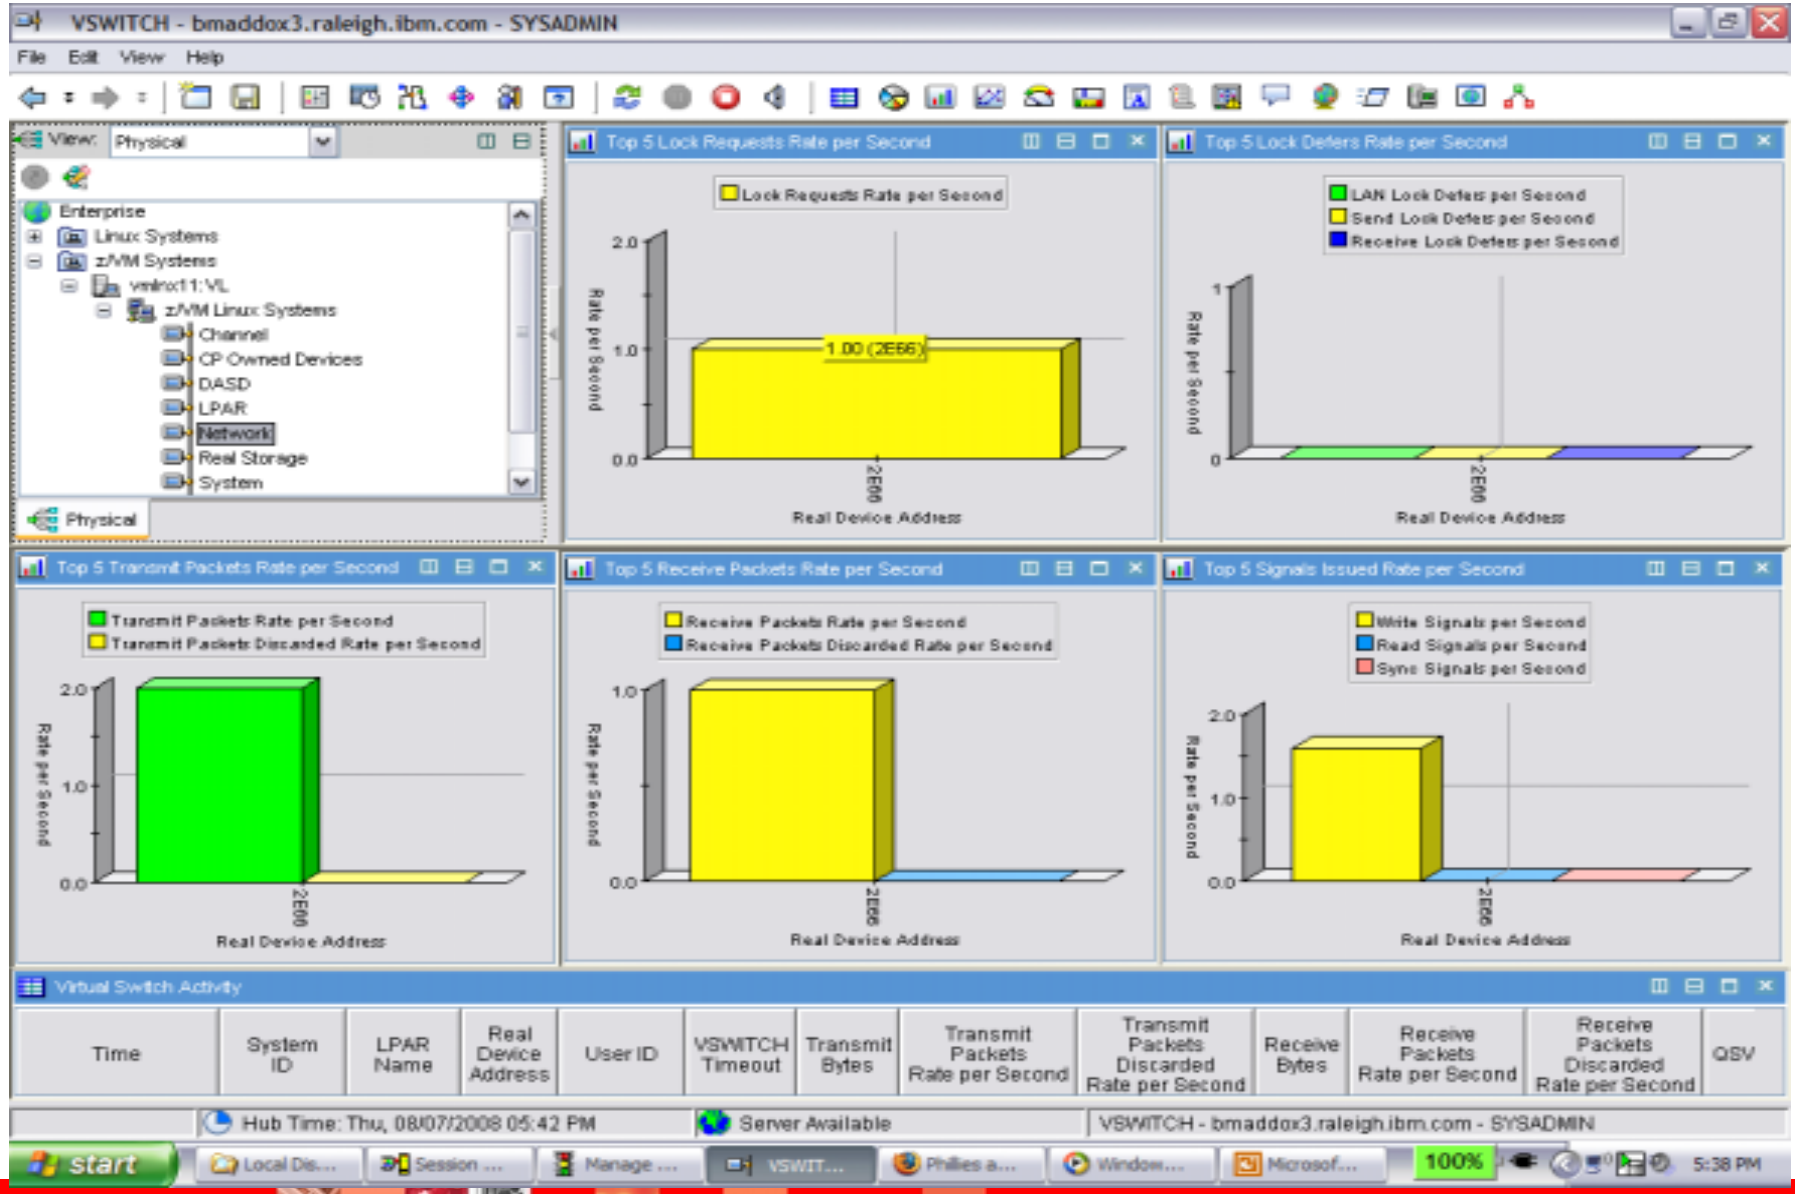

NETWORK Utilization (Virtual Switch)

- **NETWORK Workspace**
  - This workspace displays data about the utilization of the hipersocket and virtual switch devices on the z/VM system
- **Virtual Switch Utilization Table (VIRTSWITCH TABLE)**
  - All attributes are collected for the current reporting interval
- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy).
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Real Device Address.
  - User id of the virtual machine the device is currently attached to.
  - Timeout value for the virtual switch in seconds.
  - Number of bytes transmitted per second.
  - Number of packets transmitted per second.
  - Number of outbound packets discarded per second.
  - Number of bytes received per second.
  - Number of packets received per second.
  - Number of inbound packets discarded per second.
  - Queue storage value (Values 1 - 8).

Virtual Switch (cont)

- The name of the link aggregation port group in use for this virtual switch
- The number of input buffers processed per second
- The number of input queue overflows per second
- The number of interrupts, such as Peripheral Component Interconnect (PCI) interrupts, that resulted in input queue processing, per second.
- The number of interrupts, such as Peripheral Component Interconnect (PCI) interrupts, that resulted in input queue processing, per second
- The LAN Management IP address
- The number of times per second that the CP monitor waited for the network lock
- The load balancing interval
- The number of lock requests made for the network lock, per second
- The LAN Management Media Access Control (MAC) address
- The Open Systems Adapter (OSA) device microcode level
- The number of output buffers processed per second

Virtual Switch (cont)

- The number of output queue overflows per second
- The number of read signals issued per second
- The number of times per second that CP monitor waited for any lock when receiving data on this VSWITCH port
- The number of times per second that CP monitor waited for any lock when sending data from this VSWITCH port
- The session layer, either 2 or 3 of the Open Systems Interconnection (OSI) seven-layer model
- The number of sync signals issued per second
- The number of Link Aggregation Control Protocol (LACP) Packet Data Units (PDUs) received on this port
- The number of LACP PDUs sent on this port
- The number of marker PDUs received
- The number of marker PDUs sent to this port in response to receiving a marker PDU from the partner port
- The number of marker PDUs sent to this port
- The total number of times the virtual switch timed out while waiting for a marker response PDU for a marker request sent by CP monitor to a partner port
- The number of write signals issued per second

REAL STORAGE Utilization

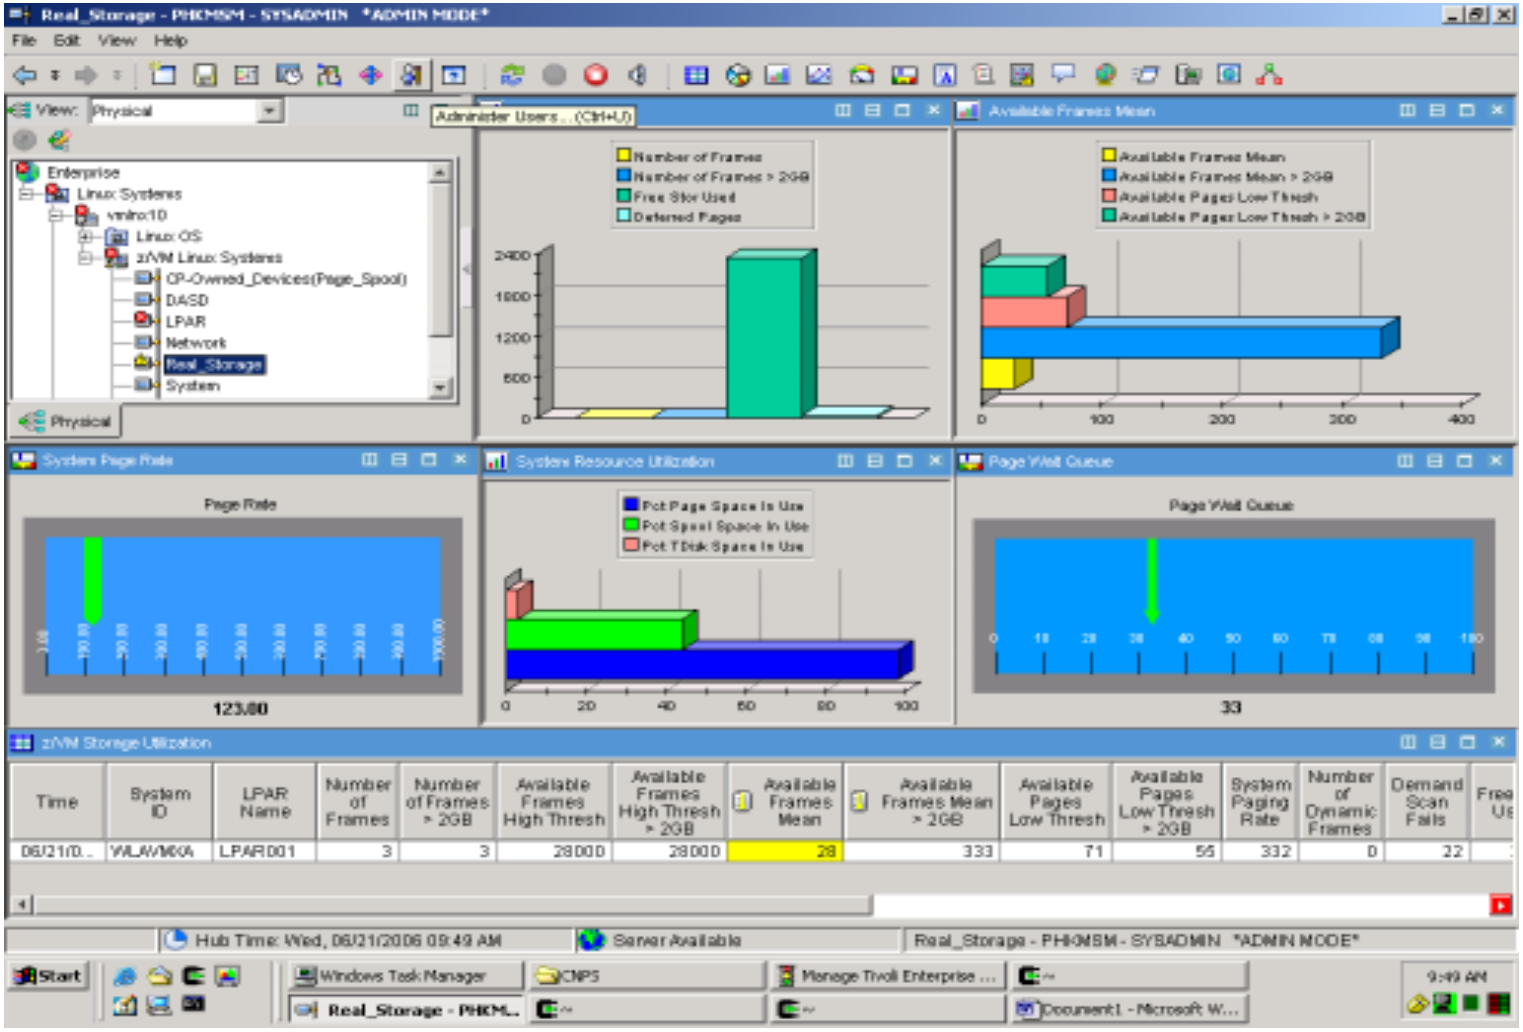

REAL STORAGE Utilization

- **REALSTORAGE Workspace**
  - The Real Storage Utilization workspace provides several views for the overall Storage and Paging activity for the z/VM system. Additionally, the Linux Group Paging Activity view displays z/VM paging activity specific to the Linux guests.

- **z/VM Storage Utilization (SYSTEM TABLE storage attributes only)**
  - This Table reflects only the portion of the System table that contains storage related attributes.
  - All attributes are collected for the current reporting interval

- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy).
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Number of frames currently on the available list.
  - High threshold for the available list replenishment subsystem.
  - Average number of page frames on the available list.
  - Low threshold for the available list replenishment subsystem.
  - System-wide I/O paging rate. Related statistics appear elsewhere by DASD volume.
  - Number of frames allocated to the dynamic paging area.
  - Number of times the demand scan was invoked and could not replenish the available list to its threshold.
  - Number of frames used by free-storage management.
  - Number of deferred pages waiting for a frame.
  - Total number of free-storage requests during the reporting interval.
  - Total number of free-storage releases during the reporting interval.
  - Number of pages per second being read in by the system.
  - System resource weight for paging. Used by the scheduler to decide how much of a bottleneck the paging resources are.
  - Average system-wide percent of paging space in use. Related statistics appear elsewhere by DASD volume.
  - Percent of SPOOL space in use for the entire system.
  - Percent of temporary disk space in use for the entire system.
  - Average number of page faults per second for single-page reads during the reporting interval.
  - Percentage of real storage available to the Dynamic Paging Area.
  - Average number of users in queue waiting to be dispatched.
  - Percent of all virtual machines in a page wait state.

REAL STORAGE Utilization (cont.)

- **REALSTORAGE Workspace (cont)**
- The Real Storage Utilization workspace provides several views for the overall Storage and Paging activity for the z/VM system. Additionally, the Linux Group Paging Activity view displays z/VM paging activity specific to the Linux guests.

- **Linux Storage Utilization (WORKLOAD TABLE storage attributes only)**
- This Table reflects only the portion of the Workload table that contains storage related attributes used by the Real Storage Workspace.
- All attributes are collected for the current reporting interval

- **Description**
- TOD clock at start of interval (Approximately 1 second accuracy).
- SYSID of z/VM System.
- Assigned logical partition number.
- Userid or group name.
- Rate of page-ins and page-outs for this workload (in pages/sec.).

- **Linkage and Secondary Workspaces:**
- A link exists from the Linux Storage Utilization Table. Each row in the table is keyed to a Linux Guest. Selecting the link for a specific row will take the user to the Linux Storage Utilization Workspace for the selected Linux Guest.
- There are no Secondary Workspaces for the System Workspace

SYSTEM Utilization

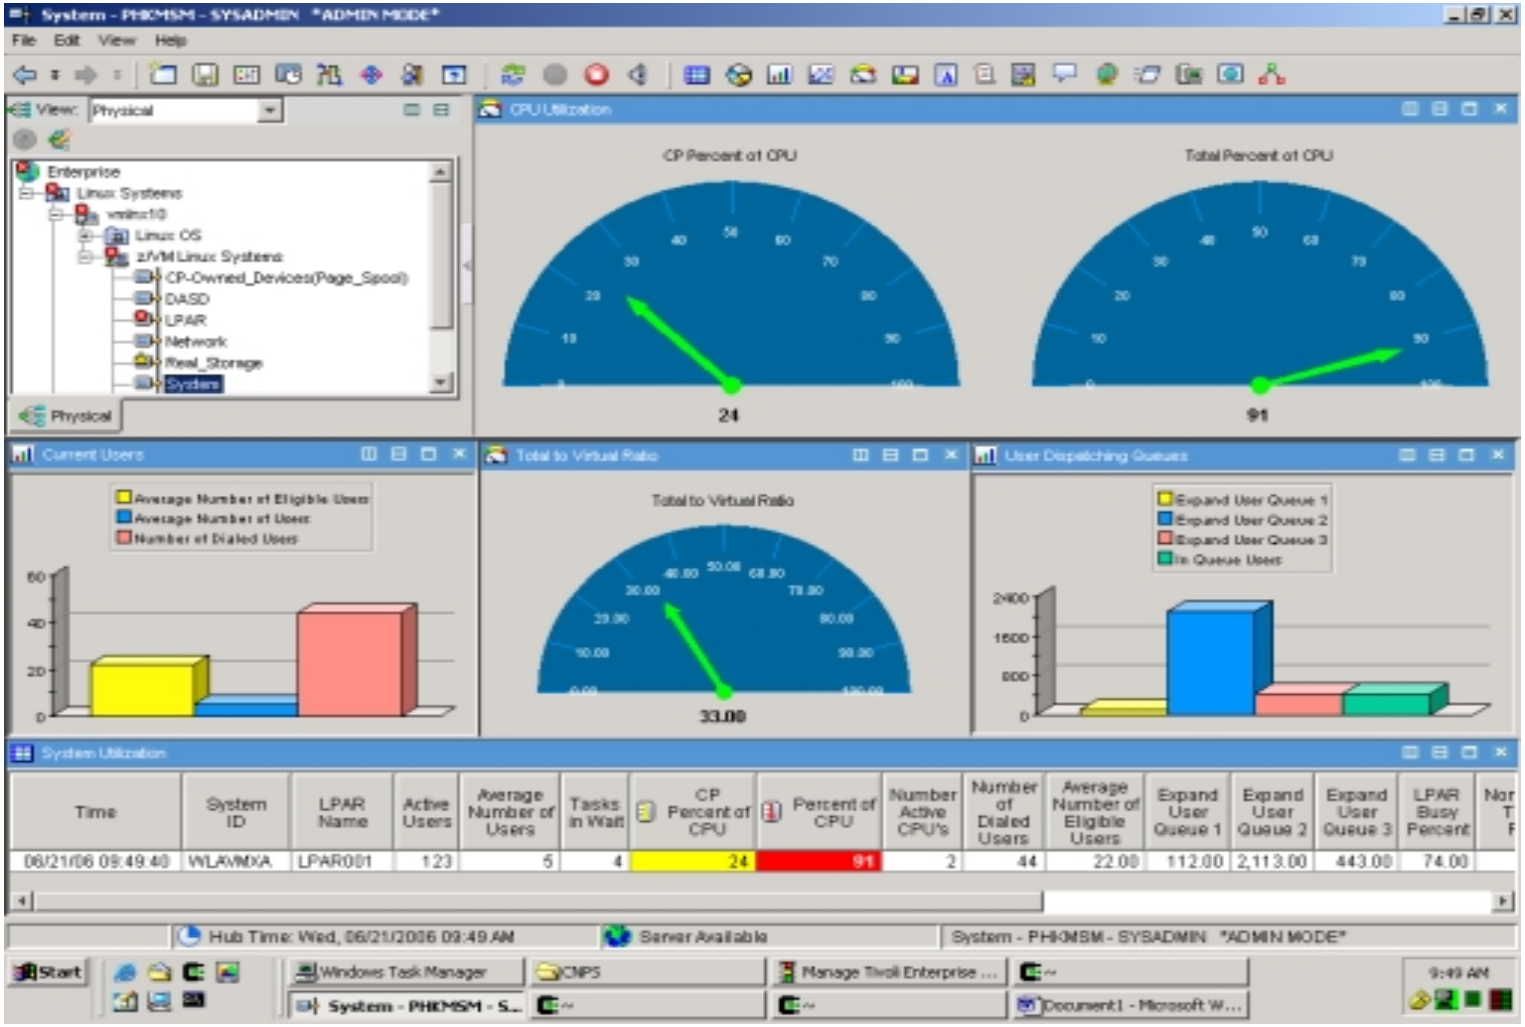

SYSTEM Utilization

- **The SYSTEM Workspace**
  - The System Utilization workspaces provide a view into the CPU utilization for the VM LPAR. The System Workspace is connected to the SYSTEM entry on the Navigator.

- All attributes are collected for the current reporting interval

- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy).
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Number of users who had any activity since the last sampling interval.
  - Average number of users logged on.
  - Number of tasks that cannot be executed because they are waiting for a frame.
  - Percentage of CPU utilized by CP.
  - Total CPU utilization (CP and virtual combined). If you are running multiple processors, this value is the sum of CPU utilization for all processors and can be greater than 100%.
  - Number of active processors.
  - Number of users who are dialed to VM.
  - Average number of users waiting in the eligible list.
  - Number of short running (interactive) users in the eligible list for the E1 queue.
  - Number of medium-running users in the eligible list for the E2 queue.
  - Number of long-running users in the eligible list for the E3 queue.
  - Total number of trivial transactions processed during the reporting interval.
  - Ratio of total CPU time to virtual CPU time.
  - Average number of users in queue waiting to be dispatched.
  - Percent of all virtual machines in an I/O wait state.

- **Linkage and Secondary Workspaces:**
  - A link is established on the table to the System_Terminal Workspace. This is a direct link to the workspace and does not require any DWL connections.
  - There are no Secondary Workspaces for the System Workspace

TCPIP Utilization - Server

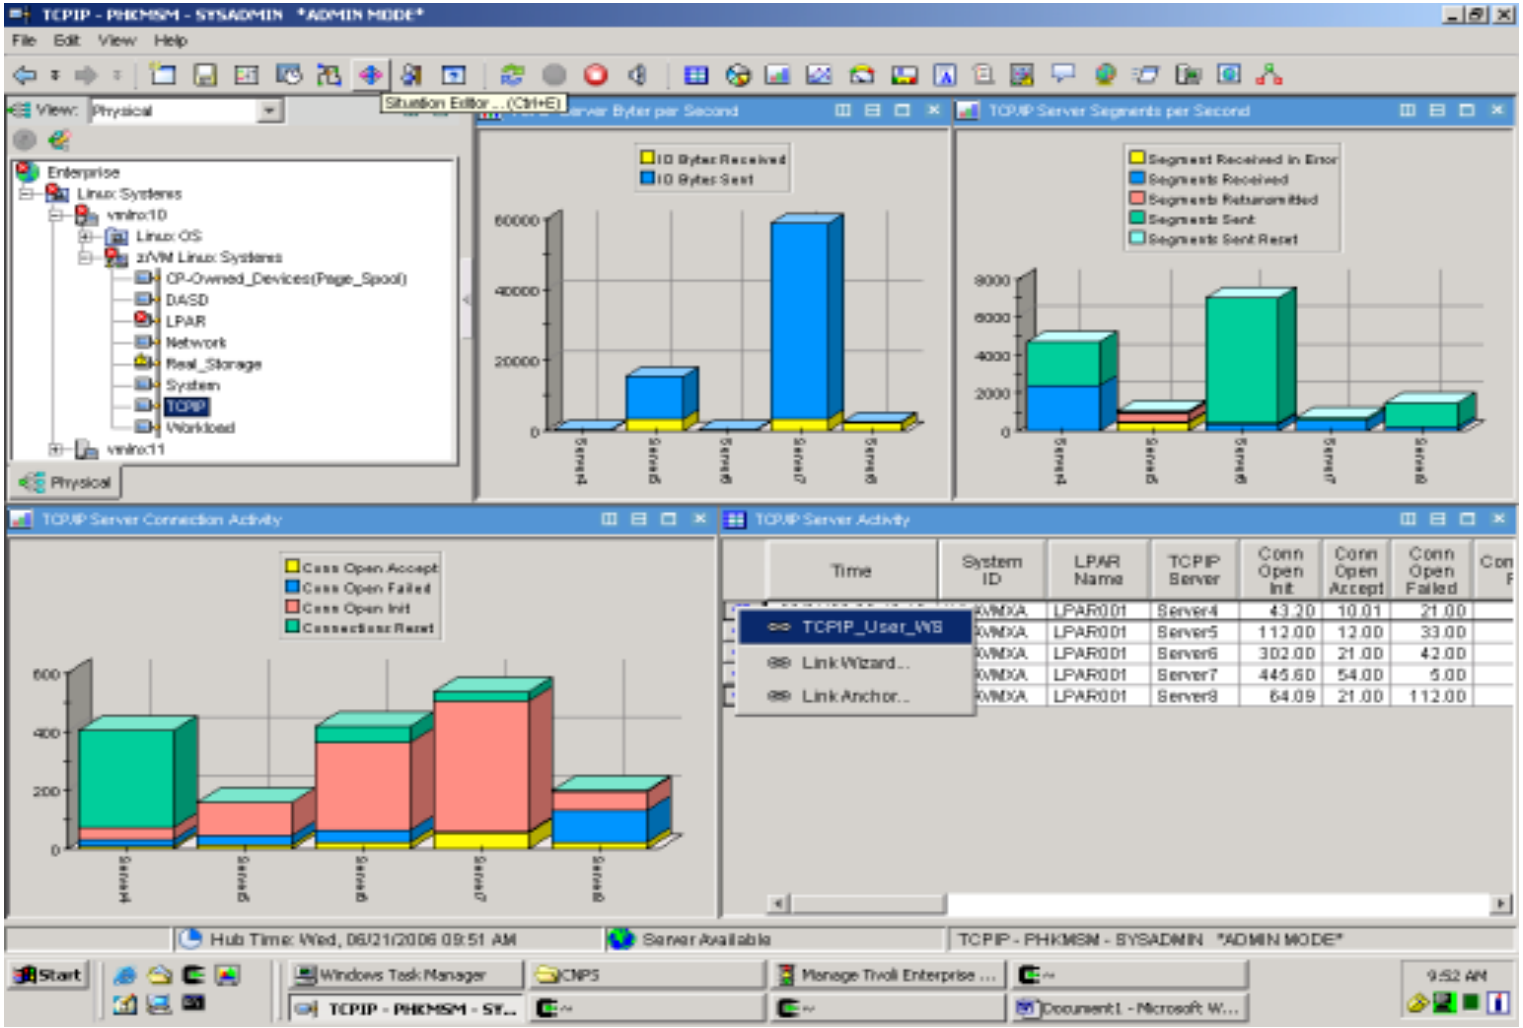

TCPIP Utilization - Server

- **TCPIP Workspace**
  - This Workspace displays data about the TCPIP Servers running on the z/VM system.

- **TCPIP Server Activity (TCPIP Table)**
  - All attributes are collected for the current reporting interval

- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy).
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Name of the TCP/IP Server.
  - Rate per second at which TCP connection open requests were initiated.
  - Rate per second at which TCP connection open requests were accepted.
  - Rate per second for TCP connection open failures.
  - Reset rate per second for TCP Connections.
  - Read requests per second.
  - Write requests per second.
  - Number of bytes received per second.
  - Number of bytes sent.
  - TCP segments received rate per second.
  - TCP segments transmit rate per second.
  - Rate at which TCP segments were retransmitted, per second.
  - Rate at which TCP segments were received that had errors, per second.
  - Rate at which TCP segments were transmitted that included a reset, in seconds.
  - ARP requests received rate per second.
  - Rate at which ARP replies were transmitted, per seconds.
  - Rate at which ARP requests were transmitted, per second.

TCPIP Utilization – Server (cont)

- **TCPIP Workspace (cont)**
  - This Workspace displays data about the TCPIP Servers running on the z/VM system.

- **TCPIP Server Activity (TCPIP Table) (cont)**
  - All attributes are collected for the current reporting interval

- **Description (cont)**
  - Activity control block pool level.
  - Client control block pool level.
  - Regular envelope pool level.
  - Large envelope pool level.
  - Raw IP Control Block Pool level.
  - Socket control block pool level.
  - BSD-type socket control block pool level.
  - TCP control block pool level.
  - UDP control block pool level.
  - Regular data buffer pool level.
  - Small data buffer pool level.
  - Tiny data buffer pool level.
  - Segment acknowledgement control block pool level.
  - Fixed page storage pool level.

- **Linkage and Secondary Workspaces:**
  - A link exists from the TCPIP Server Activity Table. Each row in the table is keyed to a TCP/IP Server. Selecting the link for a specific row will take the user to the TCPIP User Workspace for the selected Server.
  - There are no Secondary Workspaces for the System Workspace

TCPIP Utilization - Users

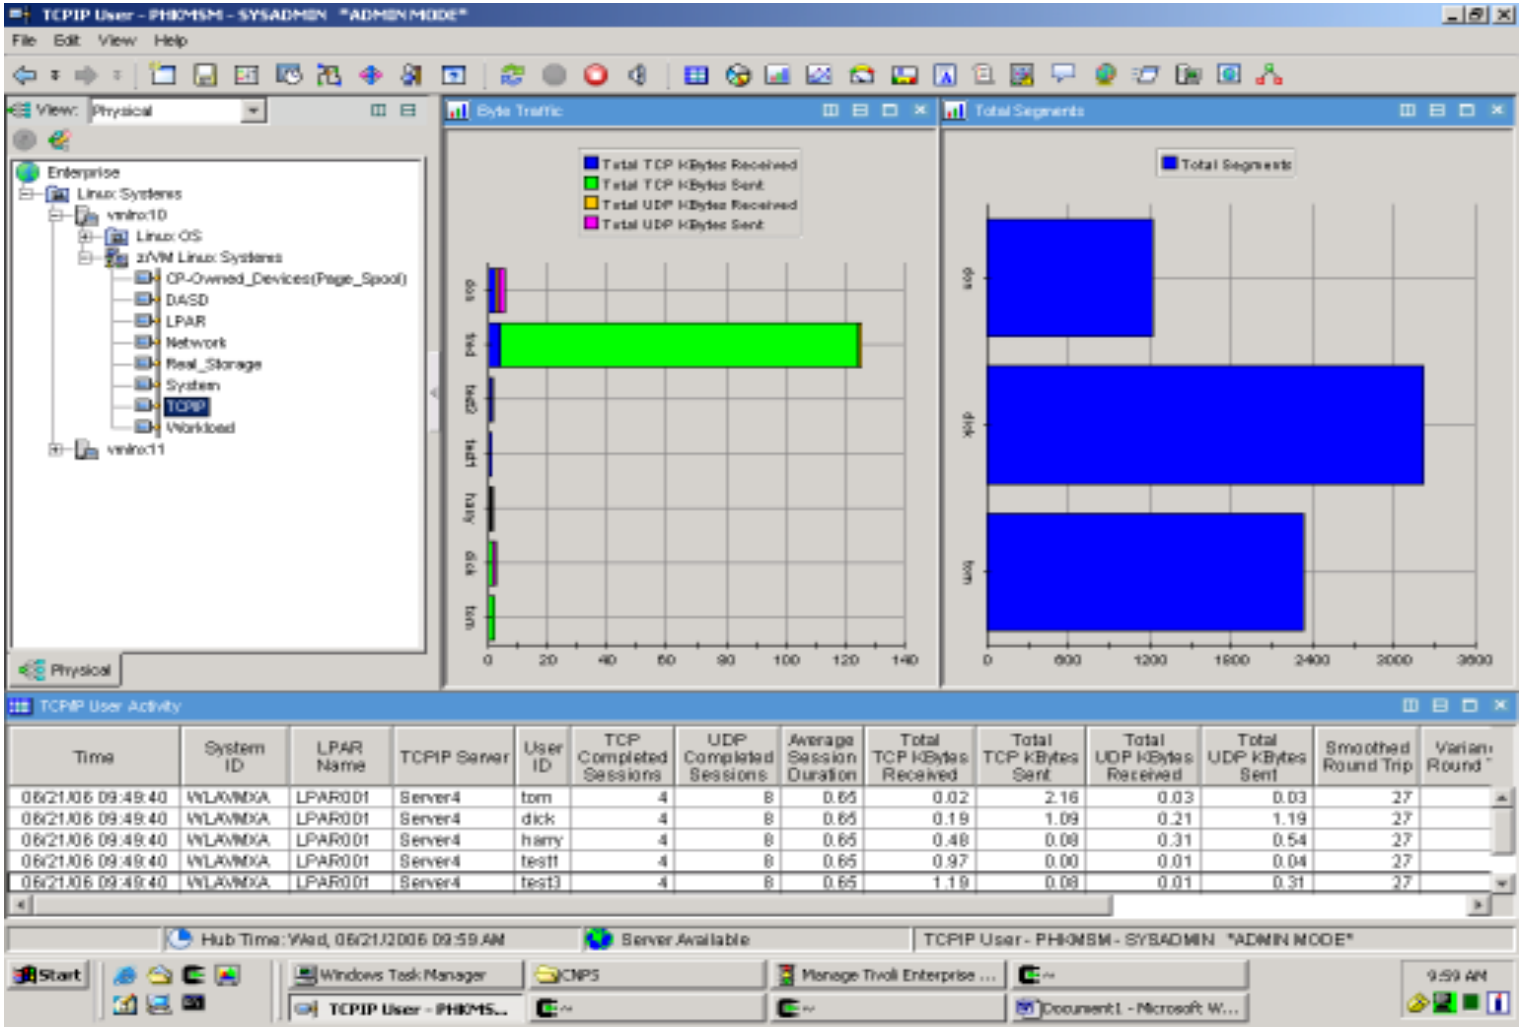

TCPIP Utilization - Users

- **TCPIP User Workspace**
  - This workspace displays data about the main users of the TCPIP function for the server selected on the previous workspace.

- **TCPIP User Activity (TCPIP USER)**
  - All attributes are collected for the current reporting interval

- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy).
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Name of the TCP/IP server that the user is connected to.
  - Userid of the virtual machine that handled the local side of the conn.
  - Number of completed TCP sessions.
  - Number of completed UDP sessions.
  - Average elapsed time from open to close for sessions.
  - Number of bytes received during TCP sessions.
  - Number of bytes sent during TCP sessions.
  - Number of bytes received during UDP sessions.
  - Number of bytes sent during UDP sessions.
  - Smoothed round trip time, in seconds.
  - Round trip variance time, in seconds.
  - Total number of segments.
  - Maximum number of unacknowledged segments.
  - Maximum input buffer queue size.
  - Maximum output buffer queue size.

- **Linkage and Secondary Workspaces:**
  - Link from the TCPIP User Activity Table to the Linux Network Workspace for the selected Linux Guest system. (Each row in the table will be for a specific guest system). This function requires DWL to be implemented.
  - There are no Secondary Workspaces for the System Workspace

System Terminal Workspace

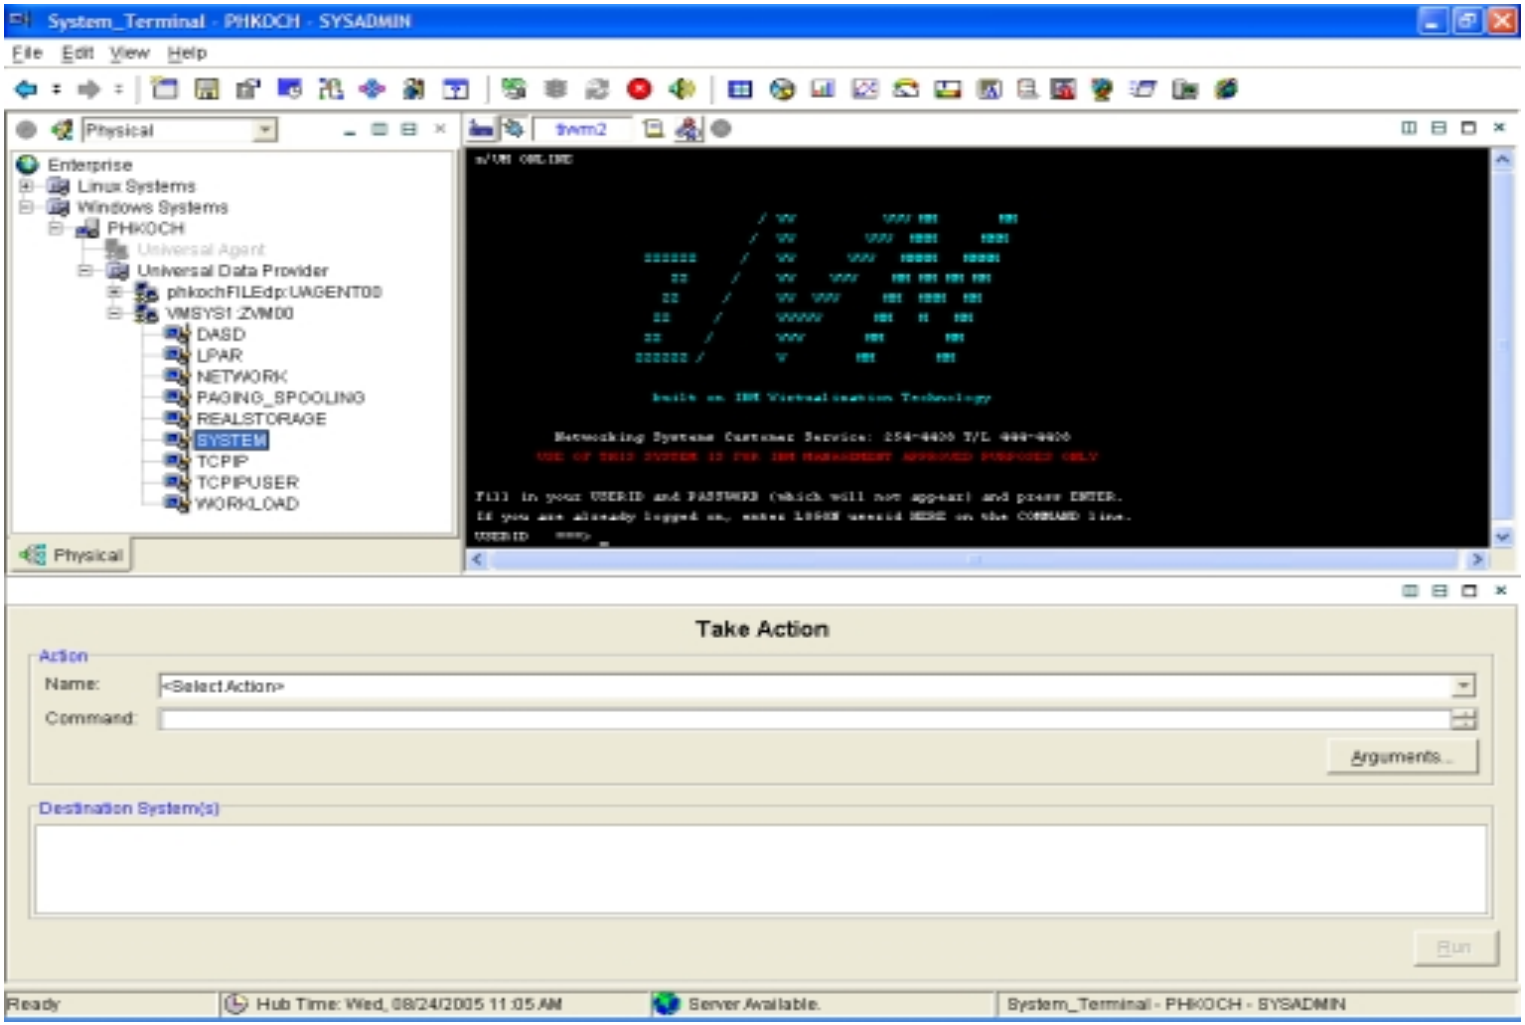

WORKLOAD (z/VM User ID) Activity

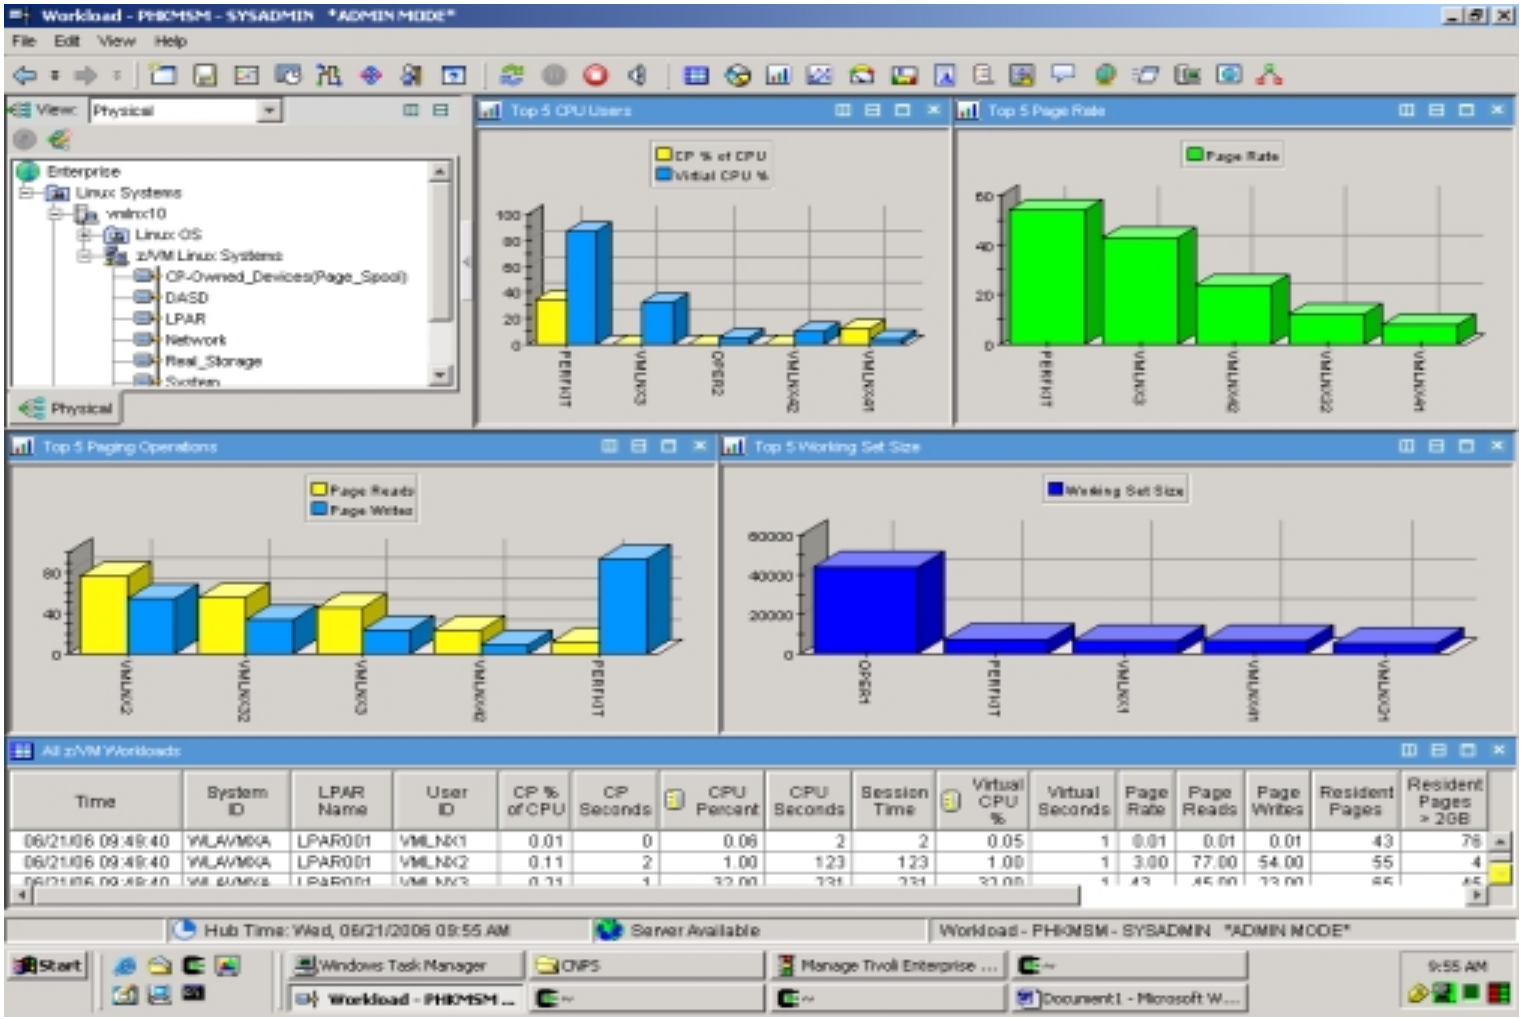

WORKLOAD (z/VM User ID) Activity

- **Workload Workspace**
  - This workspace displays the system usage (by userid/workload) for all users on the z/VM system.

- **All z/VM Workloads (Workload Table)**

- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy).
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Userid or group name.
  - Number of Virtual CPUs (VCPUs)
  - Percent of total CPU used by the system to manage this workload.
  - Percent of total CPU used by the system to manage this workload scaled by number of VCPUs
  - Total CP seconds used by this workload (to nearest second).
  - Percent of total CPU used by the system to manage this workload.
  - Percent of total CPU used by the system to manage this workload. scaled by number of VCPUs
  - Total CPU seconds used by this workload (to nearest second).
  - Total time this workload was logged on (to nearest second), or aggregation of group.
  - Percent of virtual CPU utilization for the workload specified.
  - Percent of virtual CPU utilization for the workload specified scaled by number of VCPUs
  - Total virtual CPU seconds used by this workload (to nearest second).

WORKLOAD (z/VM User ID) Activity (cont)

- **Description**
  - The rate of page-ins and page-outs for this workload (in pages/sec.).
  - Number of page reads over the specified period of time.
  - Number of page writes over the specified period of time.
  - The current number of pages physically in main storage for this workload.
  - Average storage size for this workload.
  - The number of megabytes of expanded storage attached to this workload.
  - The number of expanded pages moved for this workload.
  - The number of expanded storage blocks allocated to this workload by CP for paging.
  - A user's projected working set size. This value is calculated each time a user drops from queue, and is based on the number of pages referenced during the last stay in queue.
  - Name of the group that this workload belongs to [Primarily used to determine which VMs are Linux guest hosts.]

- **Linkage and Secondary Workspaces:**
  - No Links from this Workspace
  - Secondary Workspaces from the WORKLOAD Navigator is the Linux Workload Workspace
    - Right Mouse click on WORKLOAD Navigator to go to Linux Workload workspace.

Linux Workload Workspace

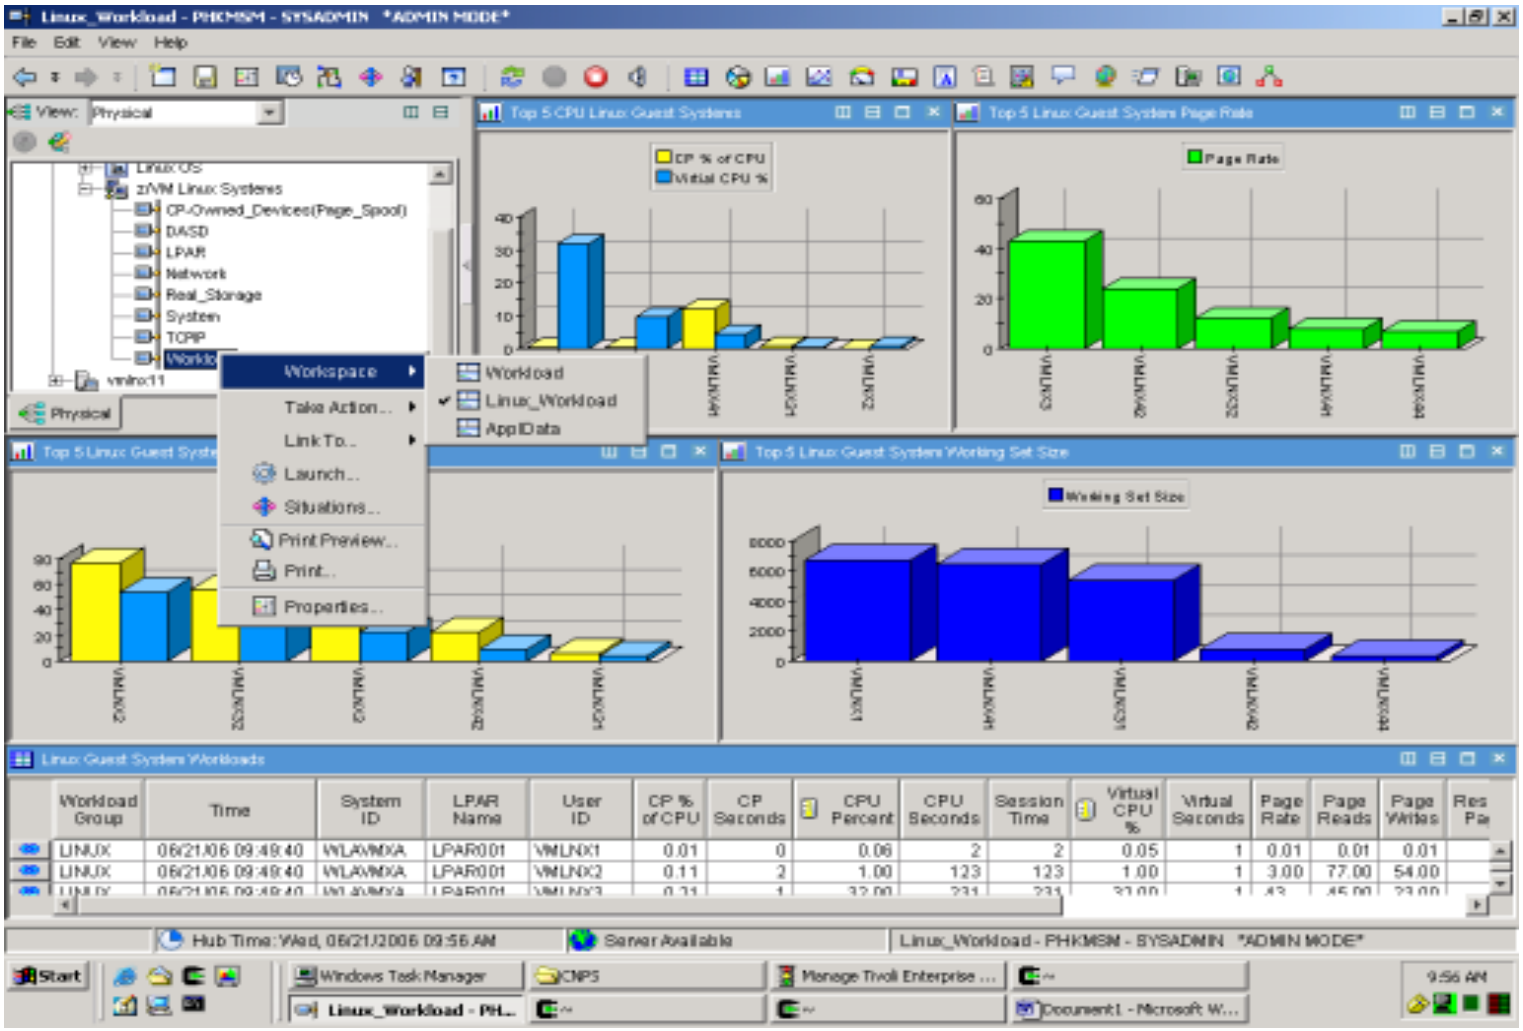

Linux Workload Workspace

- **Linux Workload Workspace**
  - This workspace displays the same information as the Workload Workspace, but only for the guest systems which have the GROUP set to Linux.

- **Linux Workloads (Workload Table)**

- **Description**
  - TOD clock at start of interval (Approximately 1 second accuracy).
  - SYSID of z/VM System.
  - Assigned logical partition number.
  - Userid or group name.
  - Percent of total CPU used by the system to manage this workload.
  - Total CP seconds used by this workload (to nearest second).
  - Percent of total CPU used by the system to manage this workload.
  - Total CPU seconds used by this workload (to nearest second).
  - Total time this workload was logged on (to nearest second), or aggregation of group.
  - Percent of virtual CPU utilization for the workload specified.
  - Total virtual CPU seconds used by this workload (to nearest second).
  - The rate of page-ins and page-outs for this workload (in pages/sec.).
  - Number of page reads over the specified period of time.
  - Number of page writes over the specified period of time.
  - The current number of pages physically in main storage for this workload.
  - Average storage size for this workload.
  - The number of megabytes of expanded storage attached to this workload.
  - The number of expanded pages moved for this workload.
  - The number of expanded storage blocks allocated to this workload by CP for paging.
  - A user's projected working set size. This value is calculated each time a user drops from queue, and is based on the number of pages referenced during the last stay in queue.
  - Name of the group that this workload belongs to [Primarily used to determine which VMs are Linux guest hosts.]

- **Linkage and Secondary Workspaces:**
  - Link from the Linux Workloads table to the OMEGAMON XE for Linux System Information Workspace for the selected Linux Guest system (by row).
  - There are no Secondary Workspaces for the Workloads Workspace

ApplData Workspace

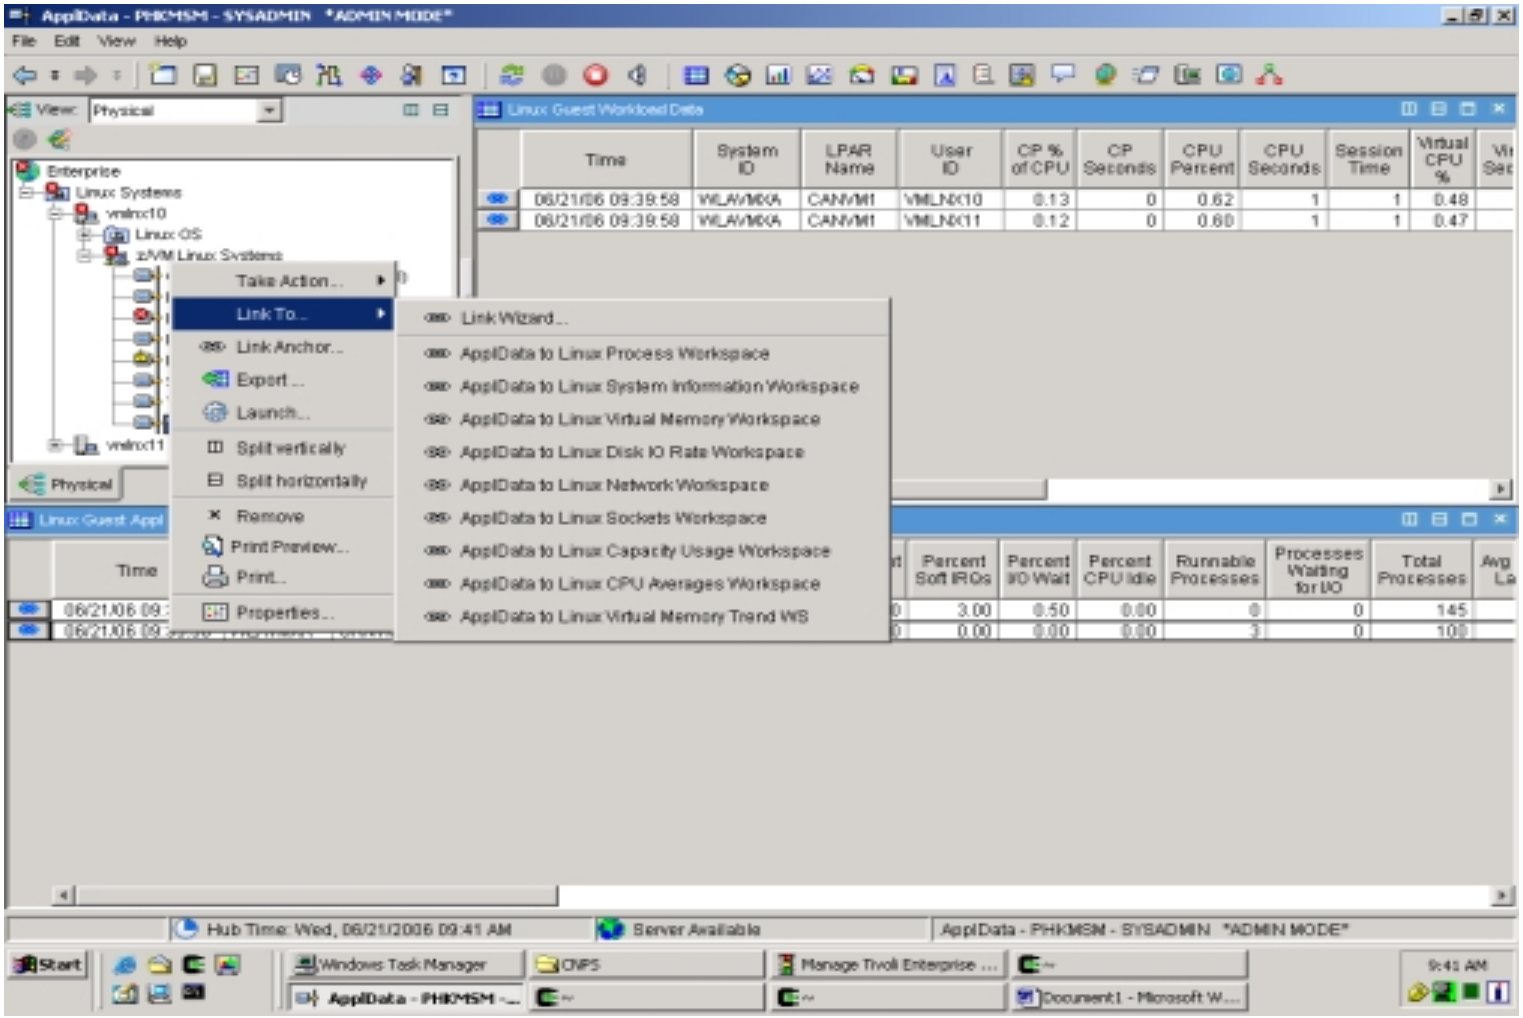

ApplData Workspace

- **ApplData Workspace**
  - This workspace displays information about Linux(R) workload activity and storage utilization
- **Description**
  - The name that uniquely identifies the active z/VM system.
  - The name assigned to the logical partition.
  - The user identification or group name of the Linux guest.
  - The number of virtual central processing units (CPUs) defined for the Linux guest system.
  - The percent of total CPU used by this Linux virtual machine.
  - The percent of CPU used by this Linux virtual machine, running in user mode.
  - The percent of CPU used by this Linux virtual machine, running in kernel mode.
  - The percent of CPU used by the Linux virtual machine, running in 'nice' mode (with modified priority).
  - The percent of interrupts (IRQs).
  - The percent of soft interrupts (IRQs).
  - The percent of time spent by the virtual machine in an I/O wait state
  - The percent of time spent by the virtual machine in a CPU idle state.
  - The number of runnable processes at sampling time.
  - The number of processes waiting for I/O.
  - The total number of processes at sampling time.
  - The average number of processes found running during the last minute.
  - The average number of processes found running during the last five minutes.
  - The average number of processes found running during the last fifteen minutes.
  - The total size of the main memory, in megabytes
  - The percent of main memory used.
  - The percent of main memory used.
  - The total size of the high memory, in megabytes.
  - The percent of high memory used.

ApplData Workspace (cont)

- **Description**

The size of the memory that is usable by more than one process, in megabytes.

- The size of the memory that is reserved for buffers and for free cache, in megabytes.
- The size of the memory that is used for buffers, in megabytes.
- The total amount of swap space, both used and available, in megabytes.
- The percent of swap space used.
- The number of pages swapped in, at the rate of 4-kilobyte pages per second.
- The number of pages swapped out, at the rate of 4-kilobyte pages per second.
- The rate of page allocations (the number of pages obtained from the available list), in 4 kilobyte pages per second.
- The rate per second of major page faults for the process.
- The rate per second of minor page faults for the process.
- The block I/O data read rate, in kilobytes per second.
- The block I/O data write rate, in kilobytes per second.
- The number of networking interfaces defined.
- The rate per second of packets received.
- The rate per second of packets transmitted.
- The rate per second of bytes received.
- The rate per second of bytes received.
- The number of bad packets received, per second.
- The rate per second of packet transmit problems.
- The rate per second, of no space found in Linux buffers.
- The rate per second, of no space available in Linux.
- The rate per second of collisions while transmitting.

ApplData Workspace (cont)

Linkage:

- Link from ApplData to Linux Process
- Link from ApplData to Linux System Information Workspace
- Link from ApplData to Linux Virtual Memory Workspace
- Link from ApplData to Linux Disk IO Rate Workspace
- Link from ApplData to Linux Network Workspace
- Link from ApplData to Linux Sockets Workspace
- Link from ApplData to Linux Capacity Usage Workspace
- Link from ApplData to Linux CPU Averages Workspace
- Link from ApplData to Linux Virtual Memory Trend Workspace

Channel Workspace

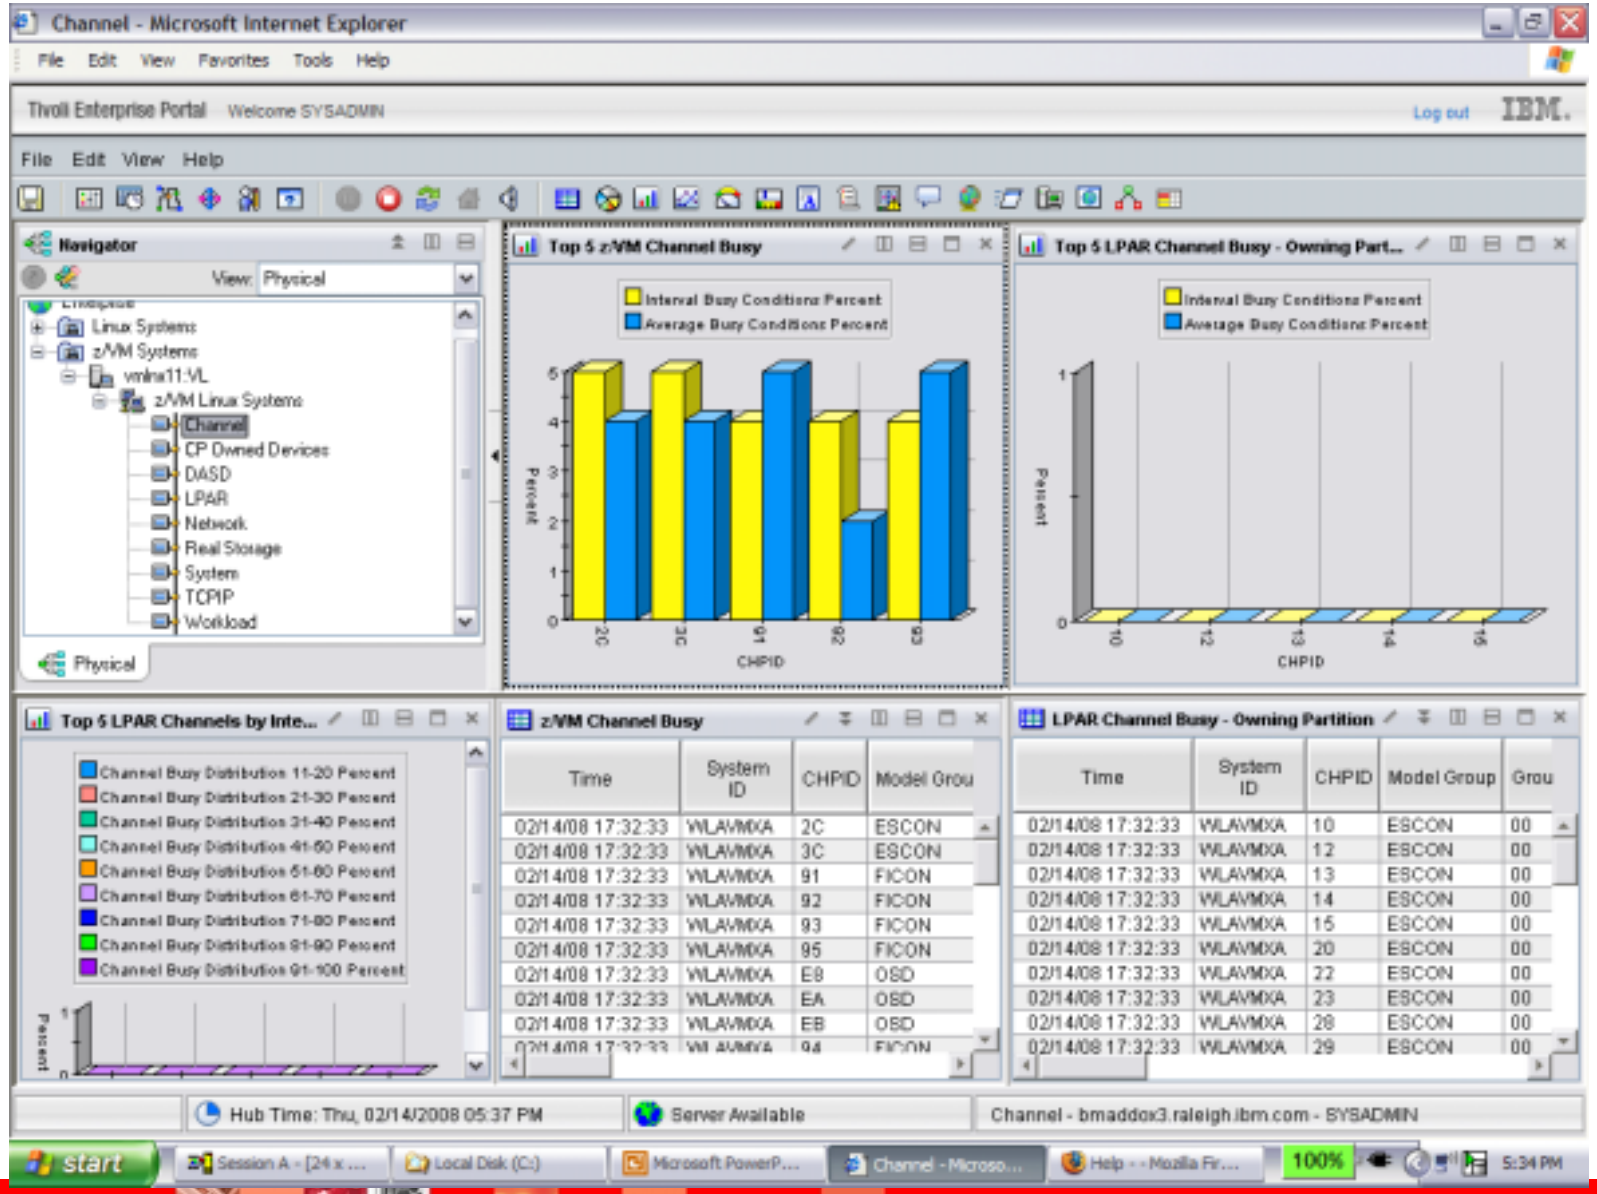

Channel Workspace

- **Channel Workspace**
  - This workspace displays channel load data for all active channel paths
- **Description**
  - The name assigned to the logical partition
  - Description of the channel model group
  - The average percentage of 'busy' conditions found
  - CHPID value assigned to each installed channel path of the system
  - The channel load distribution for each of the active channels
  - The hexadecimal channel model group qualifier
  - The percentage of 'busy' conditions found for the last interval

FICON Channel Workspace

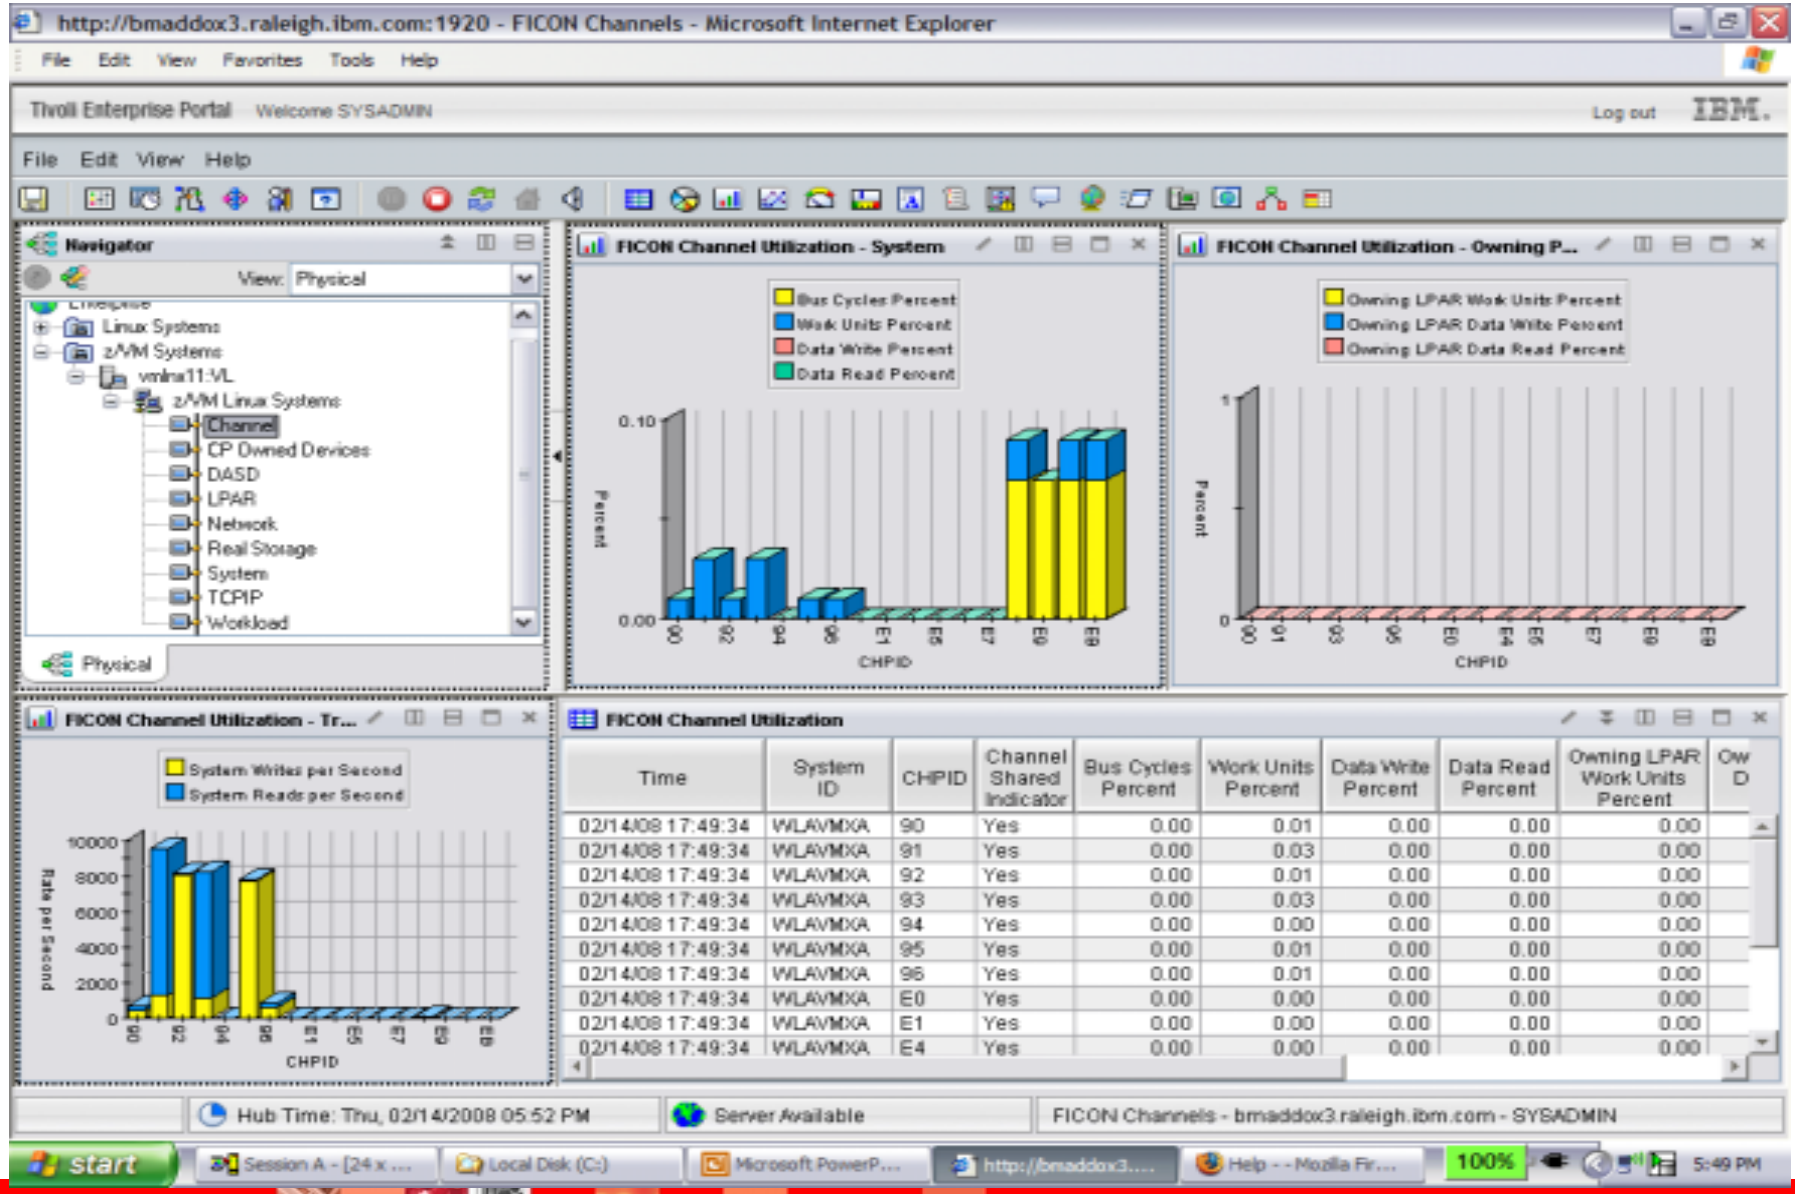

FICON Channel Workspace

- **FICON Channel Workspace**
  - This workspace displays FICON channel load data.
- **Description**
  - The name assigned to the logical partition
  - The bus cycles utilization for the whole system
  - CHPID value assigned to each installed channel path of the system
  - The channel shared indicator
  - The data units read utilization for the whole system
  - The data units write utilization for the whole system
  - The work units utilization for the entire system
  - The data units read utilization for the owning logical partition
  - The data units write utilization for the owning logical partition
  - The work units utilization for the owning logical partition
  - The total number of bytes read per second for the entire system
  - The total number of bytes written per second for the entire system

Minidisk Cache Workspace

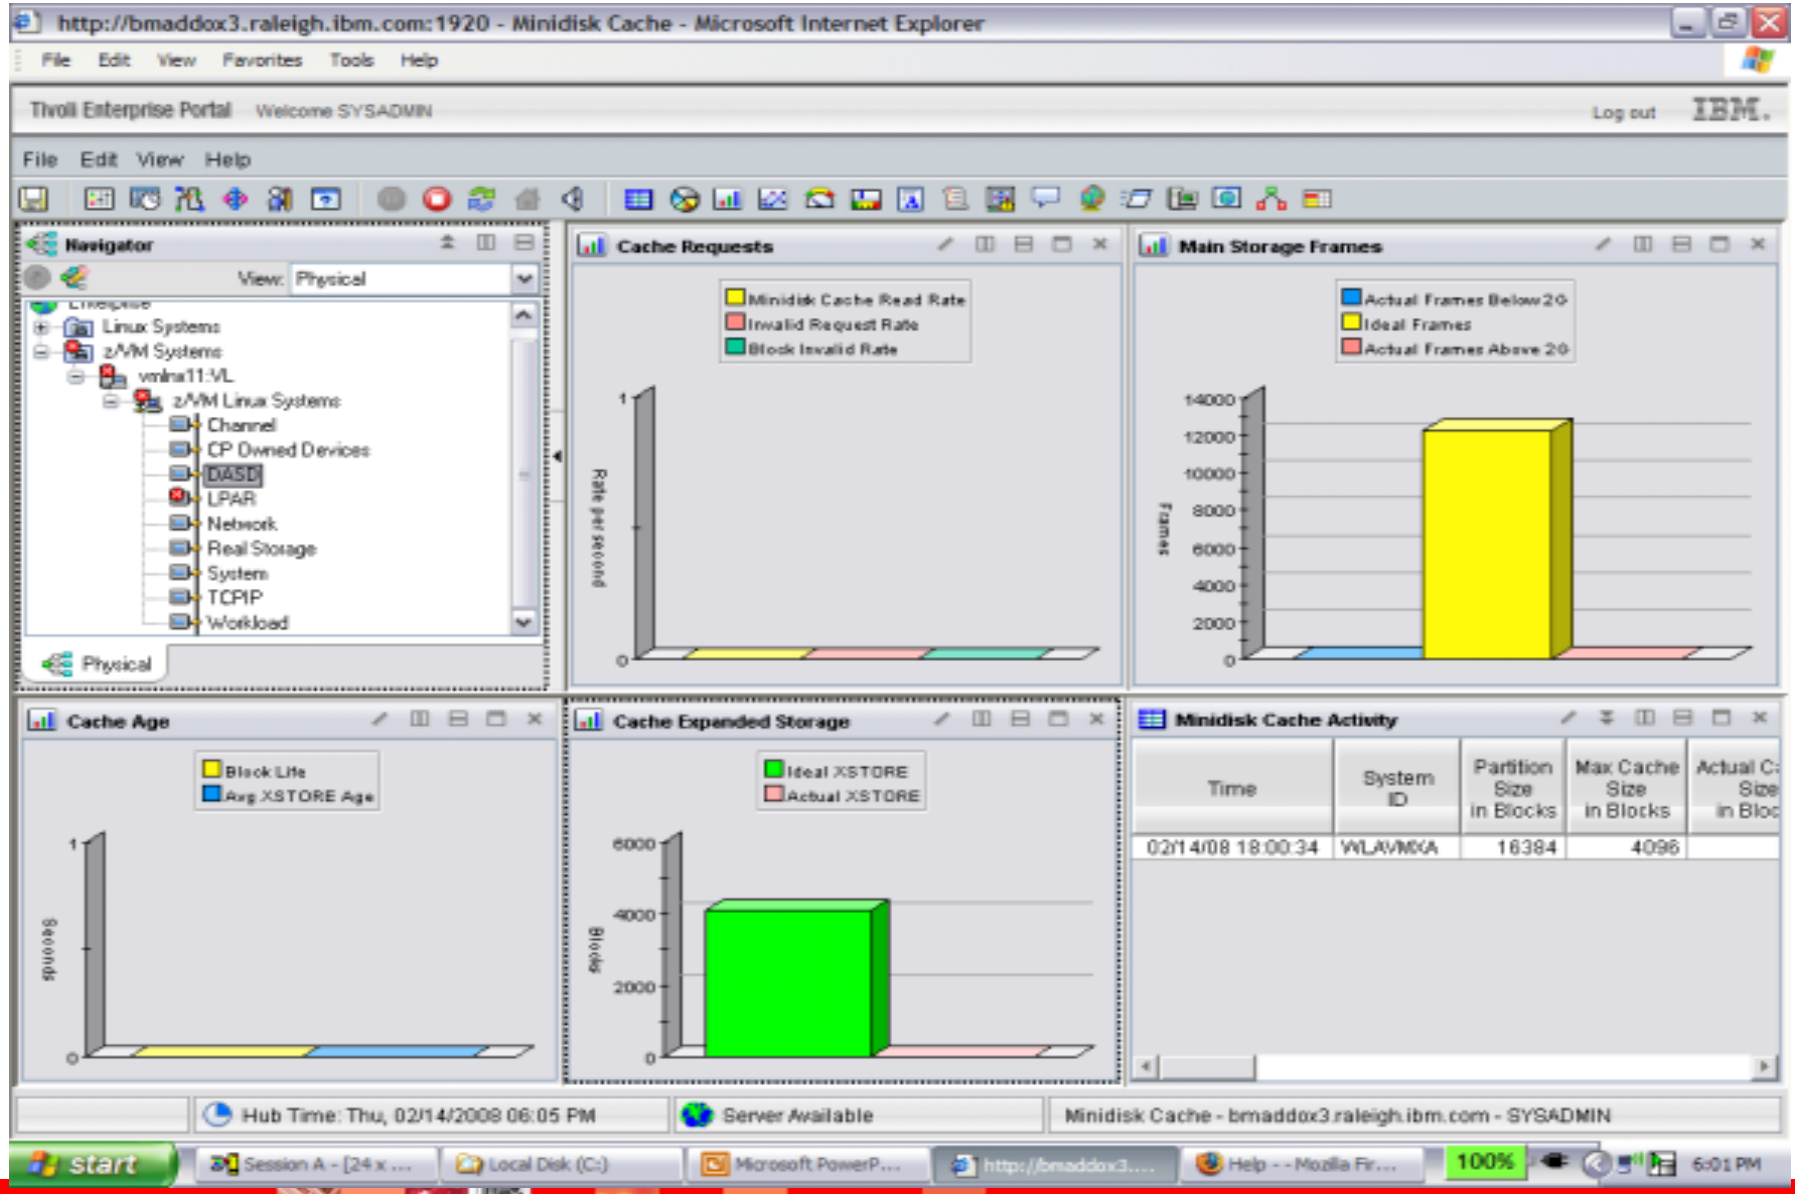

Minidisk Cache Workspace

- **Minidisk Cache Workspace**
  - This workspace displays minidisk cache storage data.
- **Description**
  - The number of expanded storage blocks used for minidisk caching
  - The actual number of main storage page frames used for the minidisk cache below the 2 GB line
  - The actual number of main storage page frames used for the minidisk cache above the 2 GB line
  - The actual number of expanded storage blocks used for the minidisk cache
  - The average age of paging XSTORE blocks
  - The bias for minidisk cache usage of real storage
  - The number of blocks per second that were invalidated following an invalidation request
  - The estimated average age in seconds of a minidisk cache block
  - The done rate for successfully translated channel command words
  - The number of CCWs that were found to be not eligible for translation
  - The rate per second that blocks could not be moved into the minidisk cache because their user's fair share limit was exceeded
  - The fair share limit for the minidisk cache
  - The percentage of requests that were full hits
  - The ideal number of main storage page frames in the minidisk cache

Minidisk Cache Workspace (cont)

- The ideal number of expanded storage blocks in the minidisk cache
- The insertions per second into the in-transit waiting queue
- The number of requests per second to invalidate minidisk cache blocks, due to an I/O to a virtual device via a non-cacheable I/O interface
- The MDC bias for main storage
- The bias for minidisk cache use of expanded storage
- The maximum number of expanded storage blocks that can be used for minidisk caching
- The maximum number of main storage page frames to be used for the minidisk cache
- The maximum number of expanded storage blocks to be used for the minidisk cache
- The number of read requests to the minidisk cache, per second, where all the requested blocks were found in the cache
- The minimum number of main storage page frames to be used for the minidisk cache
- The minimum number of expanded storage blocks to be used for the minidisk cache
- The aborted translation attempts per second for network devices

CCW Translation Workspace

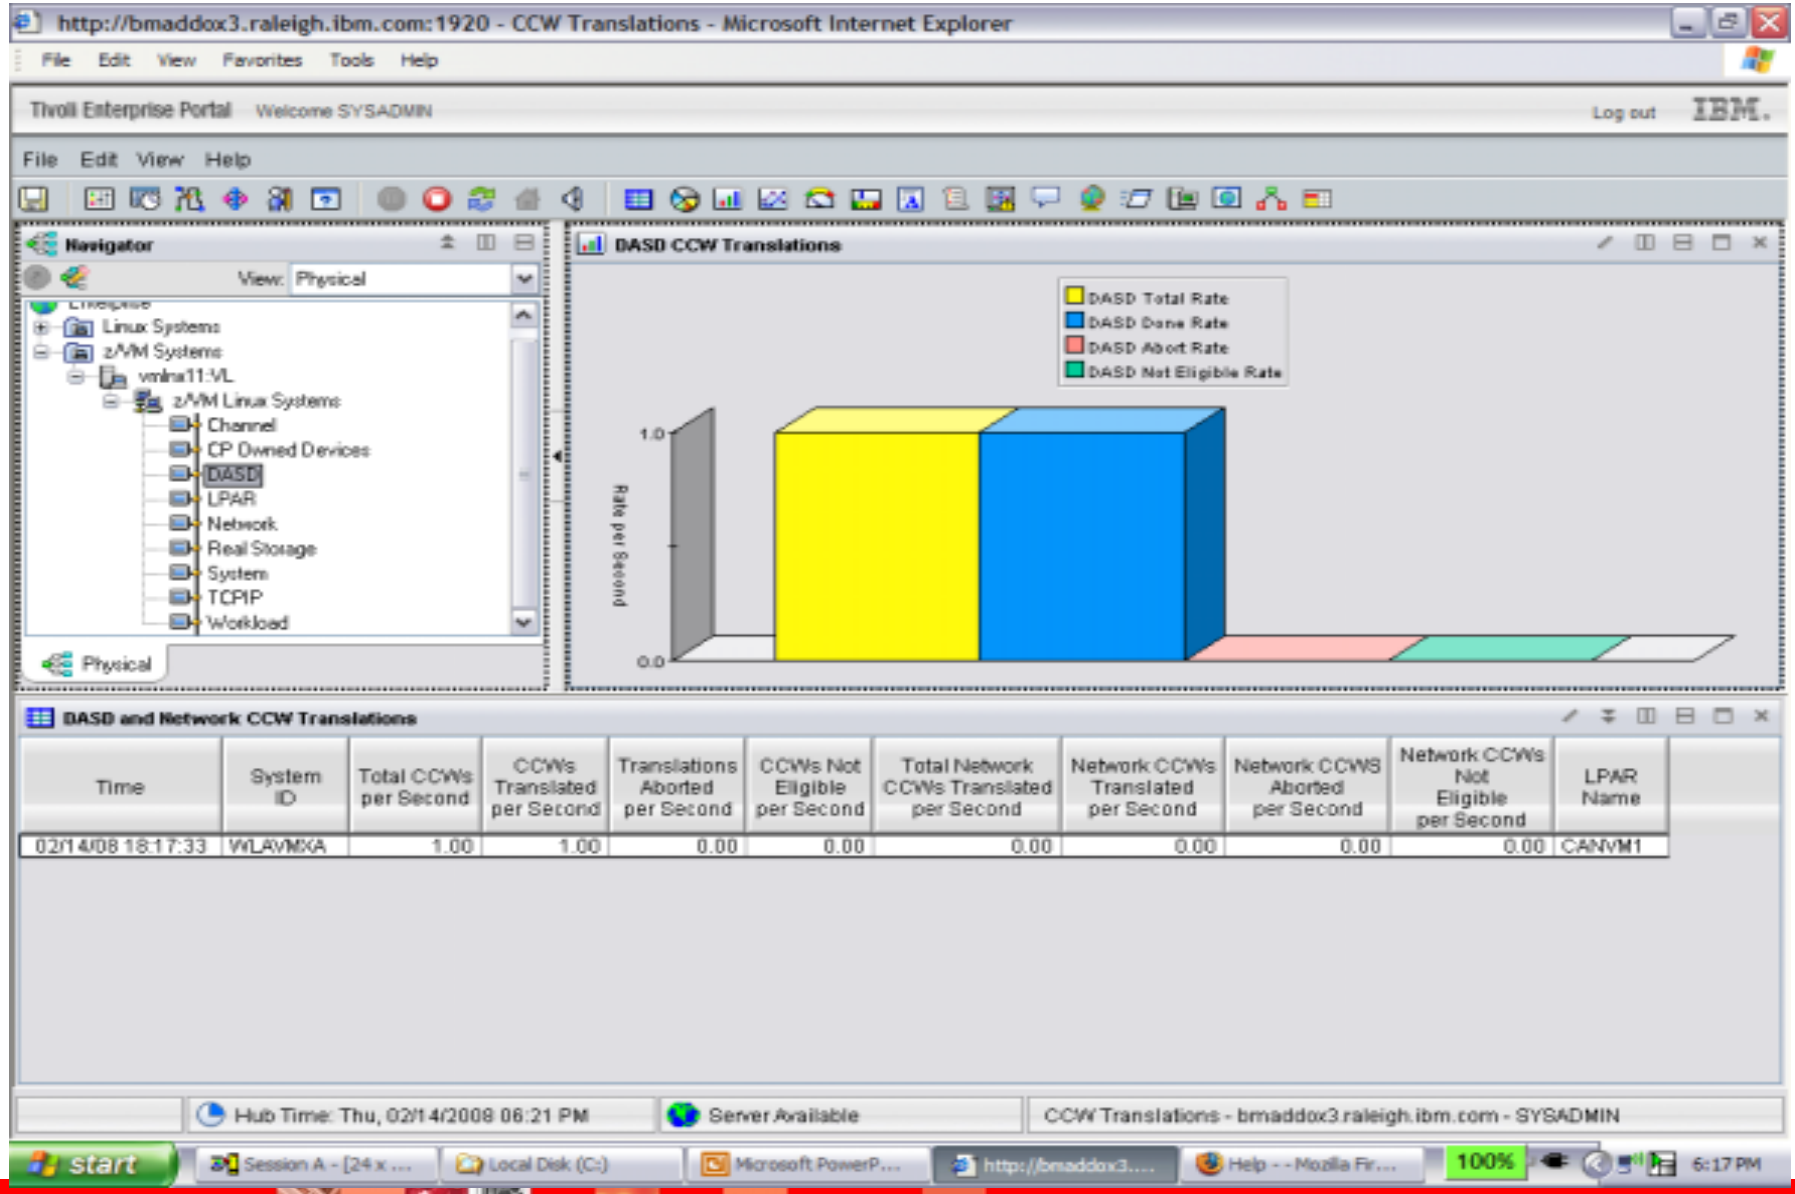

CCW Translation Workspace

- The done rate for networks
- The number of CCWs that were found to be not eligible for translation
- he number of main storage pages deleted from cache per second
- The rate at which minidisk cache blocks are moved, per second, from expanded storage to central storage
- The rate at which minidisk cache blocks are moved per second from central storage to expanded storage
- The size of the CP partition in expanded storage
- The rate at which the steal function was invoked to steal expanded storage pages from the minidisk cache
- The total number of CCWs handled per second for DASD
- The total number of CCWs handled per second for networks
- The aborted translation attempts per second for DASD
- The number of XSTORE pages deleted from cache
- The number of times the steal function was invoked to steal expanded storage pages from cache

DASD Cache Workspace

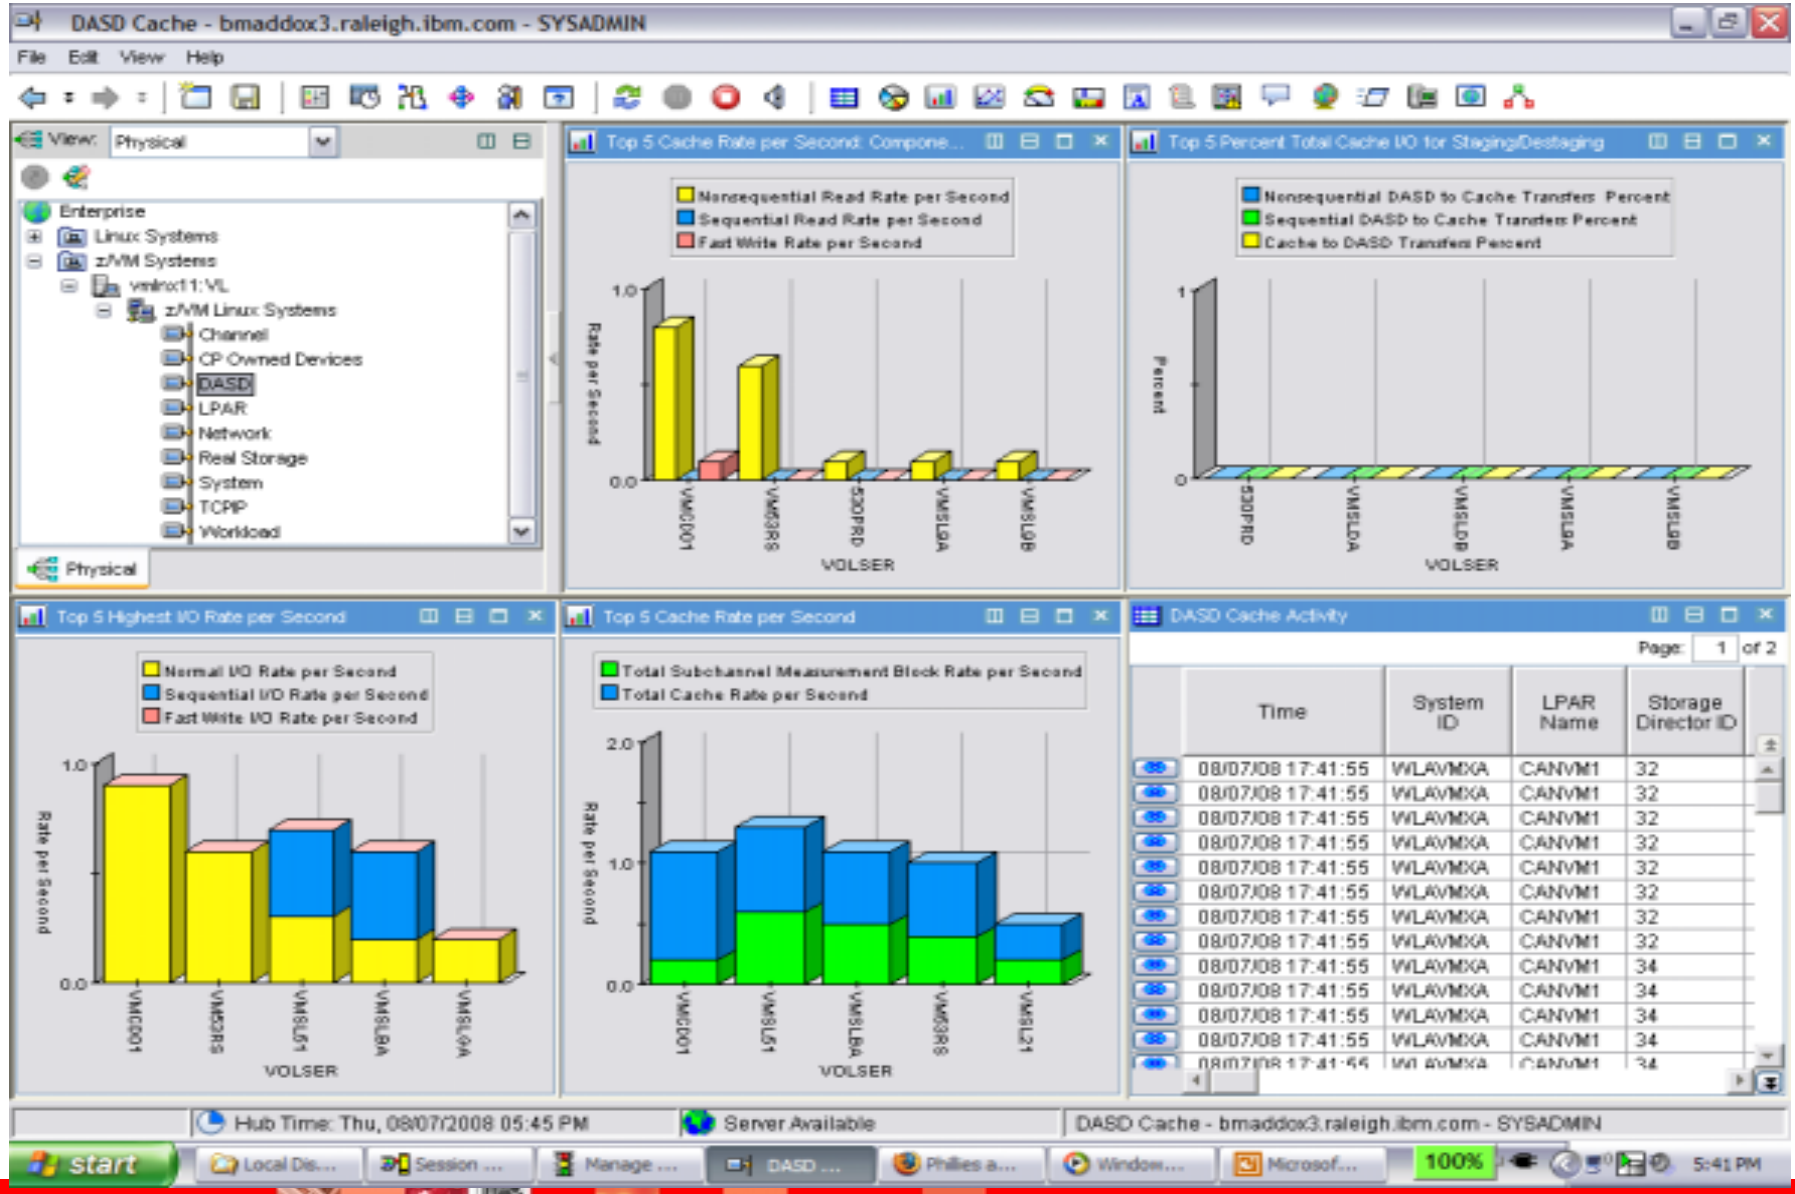

DASD Cache Workspace

- The base address for the Parallel Access Volume (PAV).
- The percentage of I/O operations that bypassed caching voluntarily
- The number of bypass cache requests per second
- The percentage of cache fast write operations, based on total write activity
- The percentage of cache fast write hits, based on the sum of all cache fast-write operations.
- The percentage of cache to DASD transfer operations
- The percentage of cachable read operations, based on total I/O activity
- The percentage of read hits, based on the sum of all cachable read operations
- The percentage of total hits (read + DASD fast write + cache fast write), based on the sum of all cachable read and write operations
- The percentage of write hits (DASD fast write + cache fast write), based on the sum of all DASD and cache fast write operations
- The caching status. Possible options are:
  - ACTIVATED
  - DEACTIVATED
  - DEACTIVE PENDING
  - UNKNOWN

DASD Cache Workspace (cont)

- The storage controller subsystem identifier.
- The percentage of DASD fast write operations that were forced to bypass the cache and access DASD directly due to nonvolatile storage constraints
- The percentage of DASD fast write hits, based on the sum of all DASD fast-write operations
- The percentage of DASD fast write operations, based on total write activity.
- The status of DASD fast write Possible options are:
  - ACTIVATED
  - DEACTIVATED
  - DEACTIVE PENDING
  - UNKNOWN
- The real address of the storage device managed by the z/VM Control Program.
- The type and model of the device managed by the z/VM Control Program
- The dual copy indicator. Possible options are:
  - DUPLEX PAIR AVAILABLE
  - DUPLEX PAIR PENDING
  - FAILED DUPLEX
  - SUSPENDED DUPLEX
  - DUPLEX NOT ACTIVE
  - UNKNOWN

DASD Cache Workspace (cont)

- The number of fast write I/O requests per second over this storage director
- The total fast-write rate per second (cache fast write + DASD fast write).
- The percentage of fast-write read requests
- The percentage of fast-write read requests that did not need DASD access
- The sum of normal, sequential, and fast write (for IBM DASD subsystems only) I/O requests, per second (read + write
- The number of inhibit cache load requests per second.
- The name assigned to the logical partition
- The percentage of nonsequential DASD to cache transfer operations
- The nonsequential read rate (read normal + read cache fast write), per second
- The number of normal I/O requests, per second, over this storage director
- The percentage of normal read requests
- The normal read percentage. That is, the percentage of read requests that did not need DASD access
- The overall percentage of read requests (normal, sequential, and fast write) to the total I/O activity of the device

DASD Cache Workspace (cont)

- The overall percentage of read hits. That is, the percentage of read requests that did not need DASD access (normal, sequential, and fast write)
- The percentage of sequential DASD to cache transfer operations.
- The number of sequential I/O requests, per second, over this storage director
- The percentage of sequential read requests. That is, the percentage of read requests that did not require DASD access
- The percentage of sequential read hits
- The sequential read rate (read sequential) per second.
- 3880-13/23 storage director ID. For IBM DASD subsystems, the last two hexadecimal digits of the service set identifier (SSID) will be inserted
- The number of timeouts that occurred while waiting for data from the control unit
- The total I/O rate, per second, for the disk as it is recorded by the cache control unit. That is, where multiple systems are connected to one control unit, the total I/O activity from all systems is shown
- The I/O rate per second, as indicated by subchannel measurement block data for the system that does the monitoring
- The overall percentage of write hits to write requests. On 3880 control units, this refers to ″write normal″ hits

Control Unit Cache Workspace

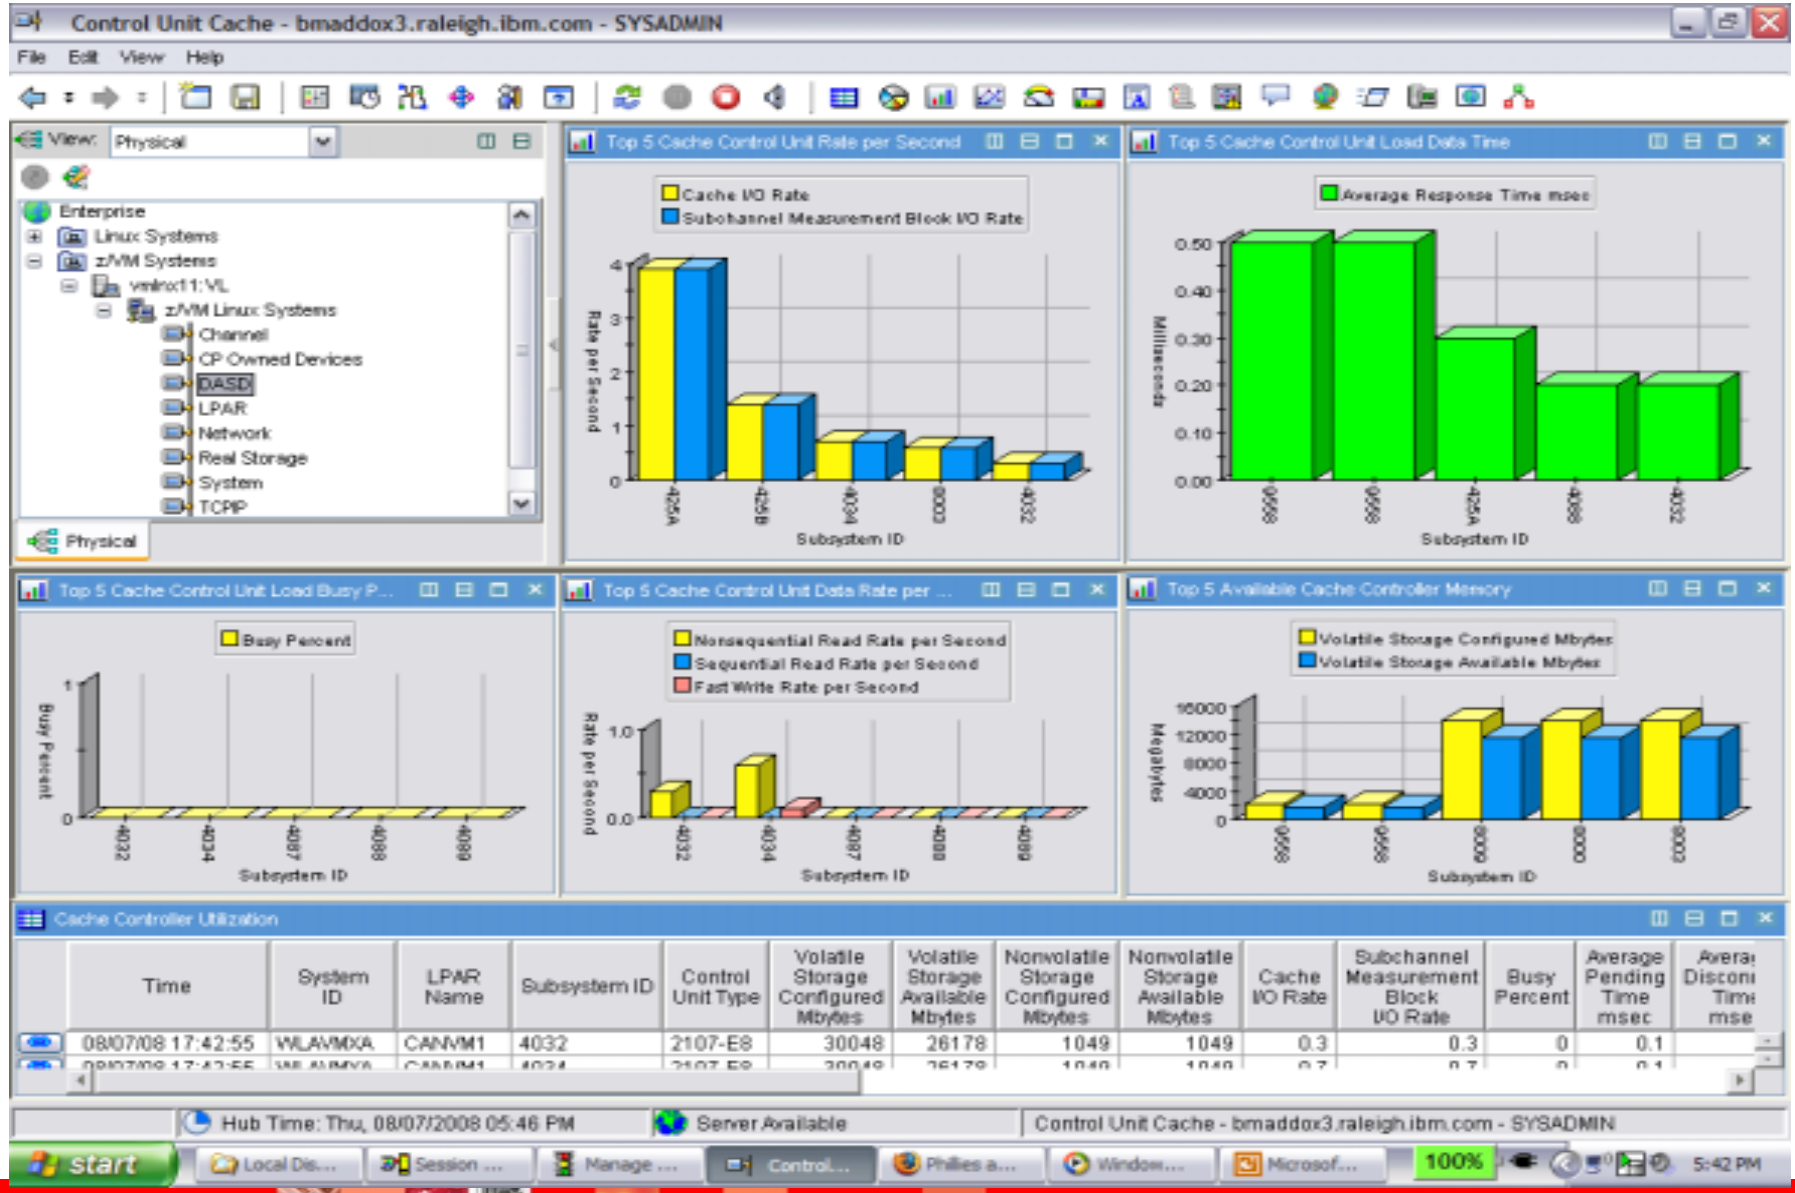

Control Unit Cache Workspace

- The average connected time, in milliseconds.
- The average disconnected time, in milliseconds.
- The average function pending time, in milliseconds
- The average percentage of read hits
- The average response time, in milliseconds, for the device.
- The average service time, in milliseconds
- The average percentage of write hits, for DASD and cache fast write operations
- The average busy percentage for all connected disks
- The total I/O activity as it is recorded by the cache control unit
- The average percentage of cache fast-write operations, based on the sum of all write operations
- The average percentage of cache fast write hits
- The type and model of the control unit

Control Unit Cache Workspace (cont)

- The average percentage of DASD fast write hits
- The average percentage of DASD fast-write operations, based on the sum of all write operations
- The total fast-write rate per second
- The total nonsequential read rate per second.
- The amount of nonvolatile storage that is available
- The size of configured nonvolatile storage
- The average percentage on read operations, based on the sum of read and write
- The total sequential read rate per second.
- The total I/O activity as determined from count fields in the subchannel measurement blocks of the system that does the monitoring
- The subsystem identifier for the control unit
- The amount of cache storage that is available
- The size of configured cache storage

Spin Locks Workspace

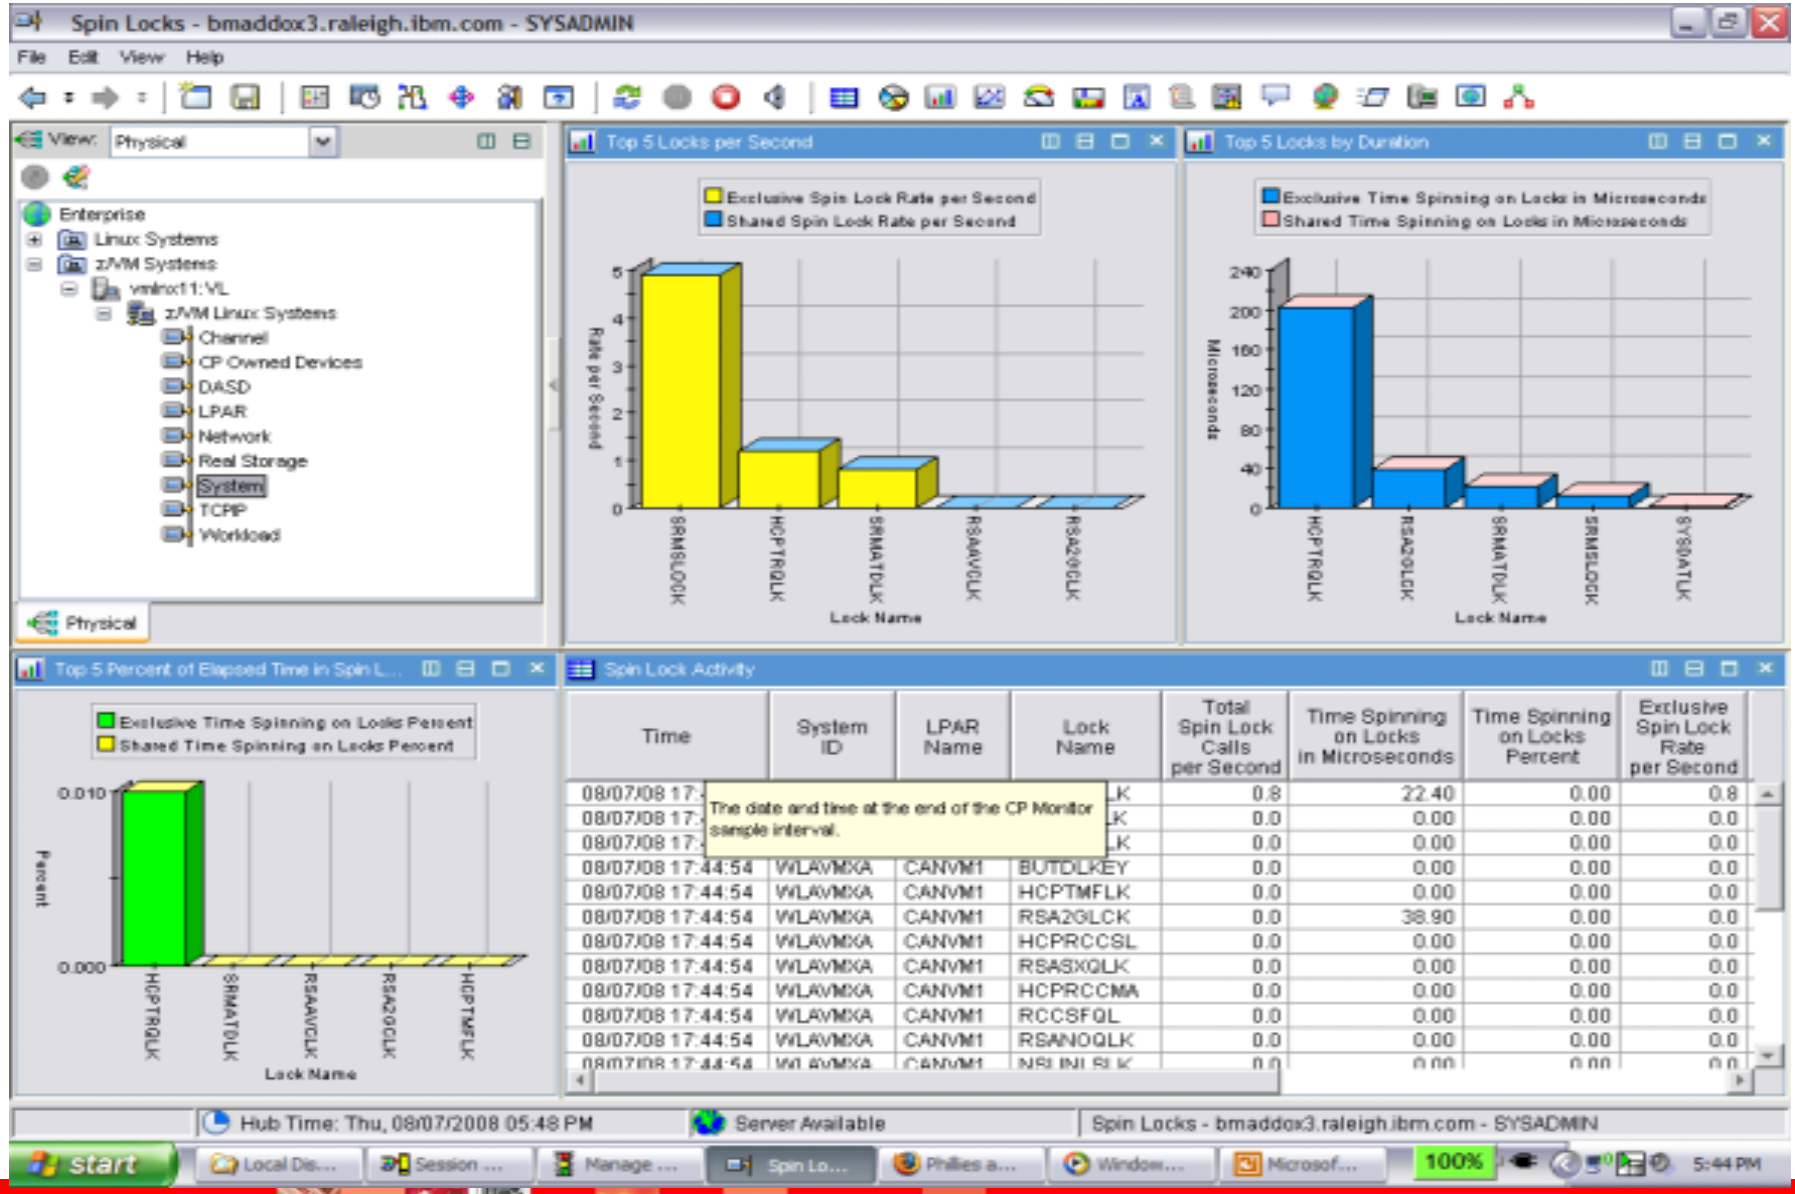

Spin Locks Workspace

- The name of the data space
- The virtual device number of the VDISK
- The number of locked data space pages
- The number of links to the virtual disk
- The number of slots occupied on auxiliary storage
- The number of data space pages moved from central storage to expanded storage, per second
- The number of data space pages read from DASD, per second
- The number of data space pages stolen per second
- The number of data space pages moved from expanded storage to central storage, per second
- The number of data space pages migrated from expanded storage to DASD, per second
- The number of data space pages written to DASD, per second
- The number of data space pages resident in central storage
- The user ID of the owner of the VDISK
- The size of the VDISK,
- The virtual I/O rate, per second, to the VDISK
- The number of XSTORE blocks occupied by the data space

Virtual Disk Workspace

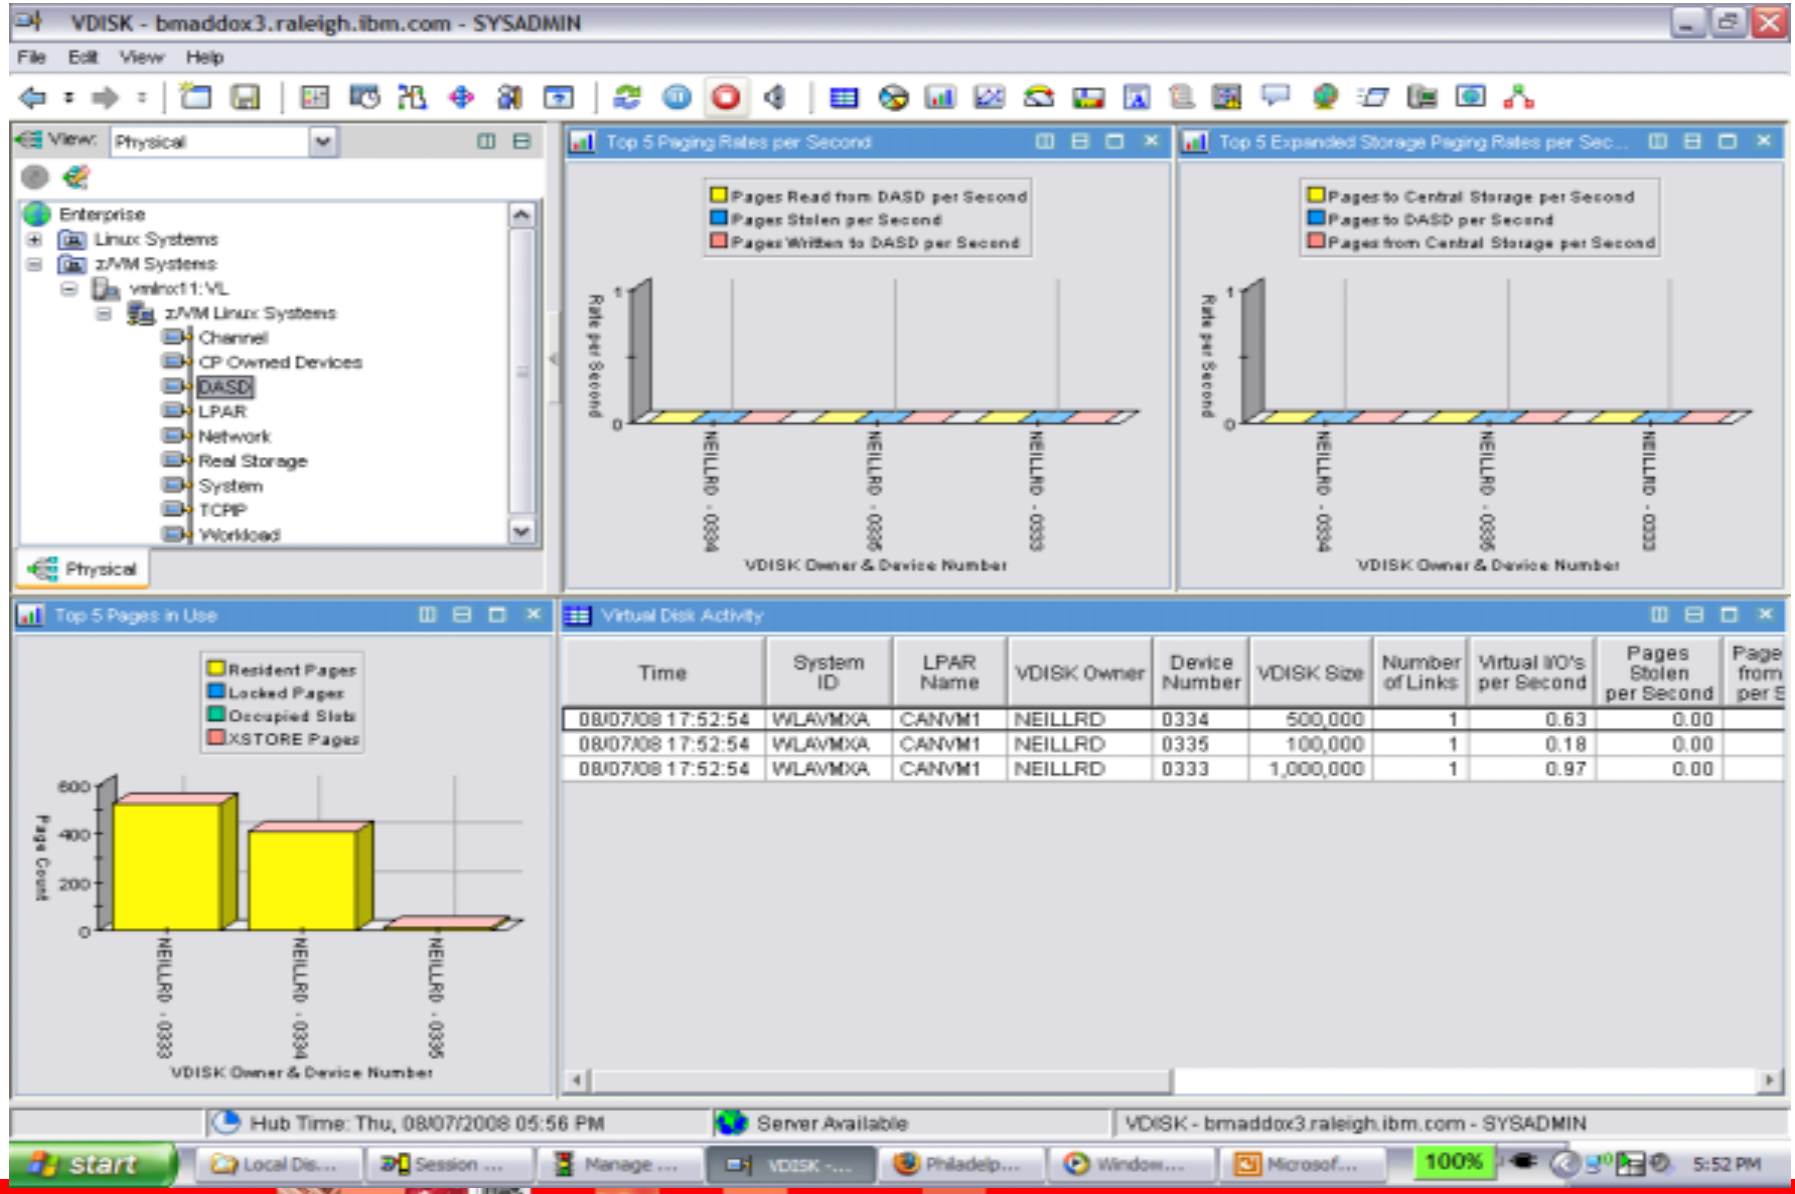

Virtual Disk Workspace

- The name of the data space
- The virtual device number of the VDISK
- The number of locked data space pages
- The number of links to the virtual disk.
- The number of slots occupied on auxiliary storage
- The number of data space pages moved from central storage to expanded storage, per second
- The number of data space pages read from DASD, per second
- The number of data space pages stolen per second
- The number of data space pages moved from expanded storage to central storage, per second
- The number of data space pages migrated from expanded storage to DASD, per second
- The number of data space pages written to DASD, per second
- The number of data space pages resident in central storage
- The user ID of the owner of the VDISK
- The size of the VDISK
- The virtual I/O rate, per second, to the VDISK
- The number of XSTORE blocks occupied by the data space

Linux on zSeries Primary Workspaces

- **Linux OS**
- **Capacity Usage**
- **Disk Usage**
- **File Information**
- **Network**
- **Process**
- **System Information**
- **Users**

Linux OS

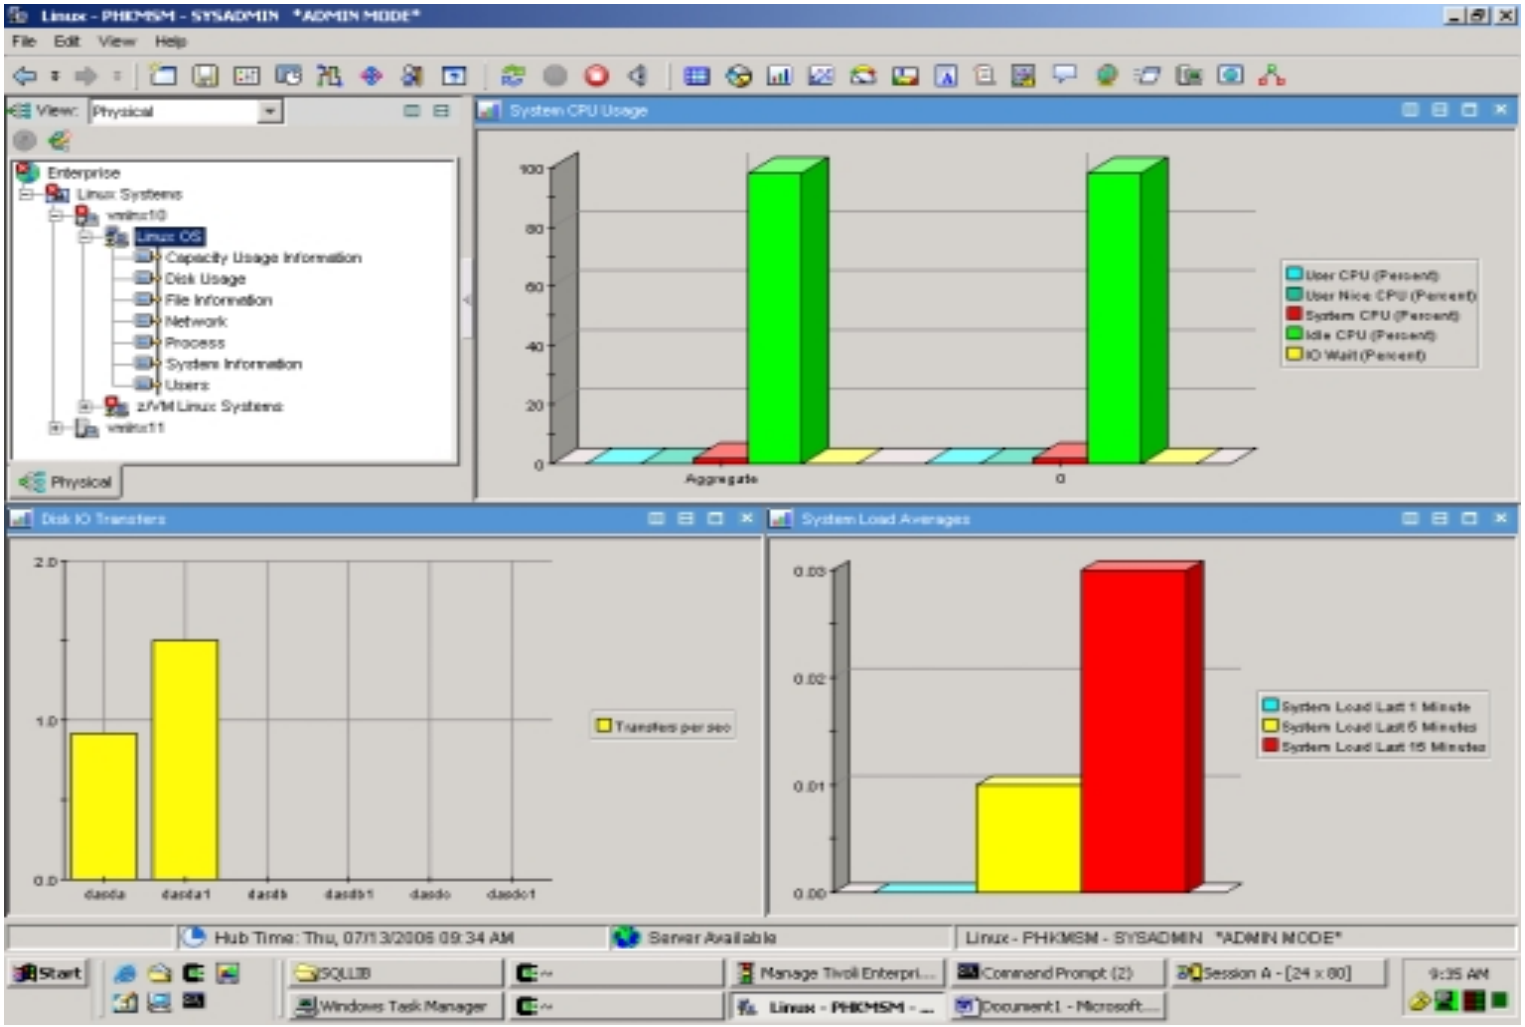

Linux OS Workspace

- **Linux OS Workspace**
  - This workspace displays overall system data for the Linux system

- **Linux OS**

- **Description**
  - Bar chart showing the percentages of CPU usage, by user CPU, user nice, system, and idle categories
  - Bar chart showing the number of transfers per second that were issued to each device
  - Bar chart showing the load on the system's processor during the previous one, five, and fifteen minutes

Capacity Usage

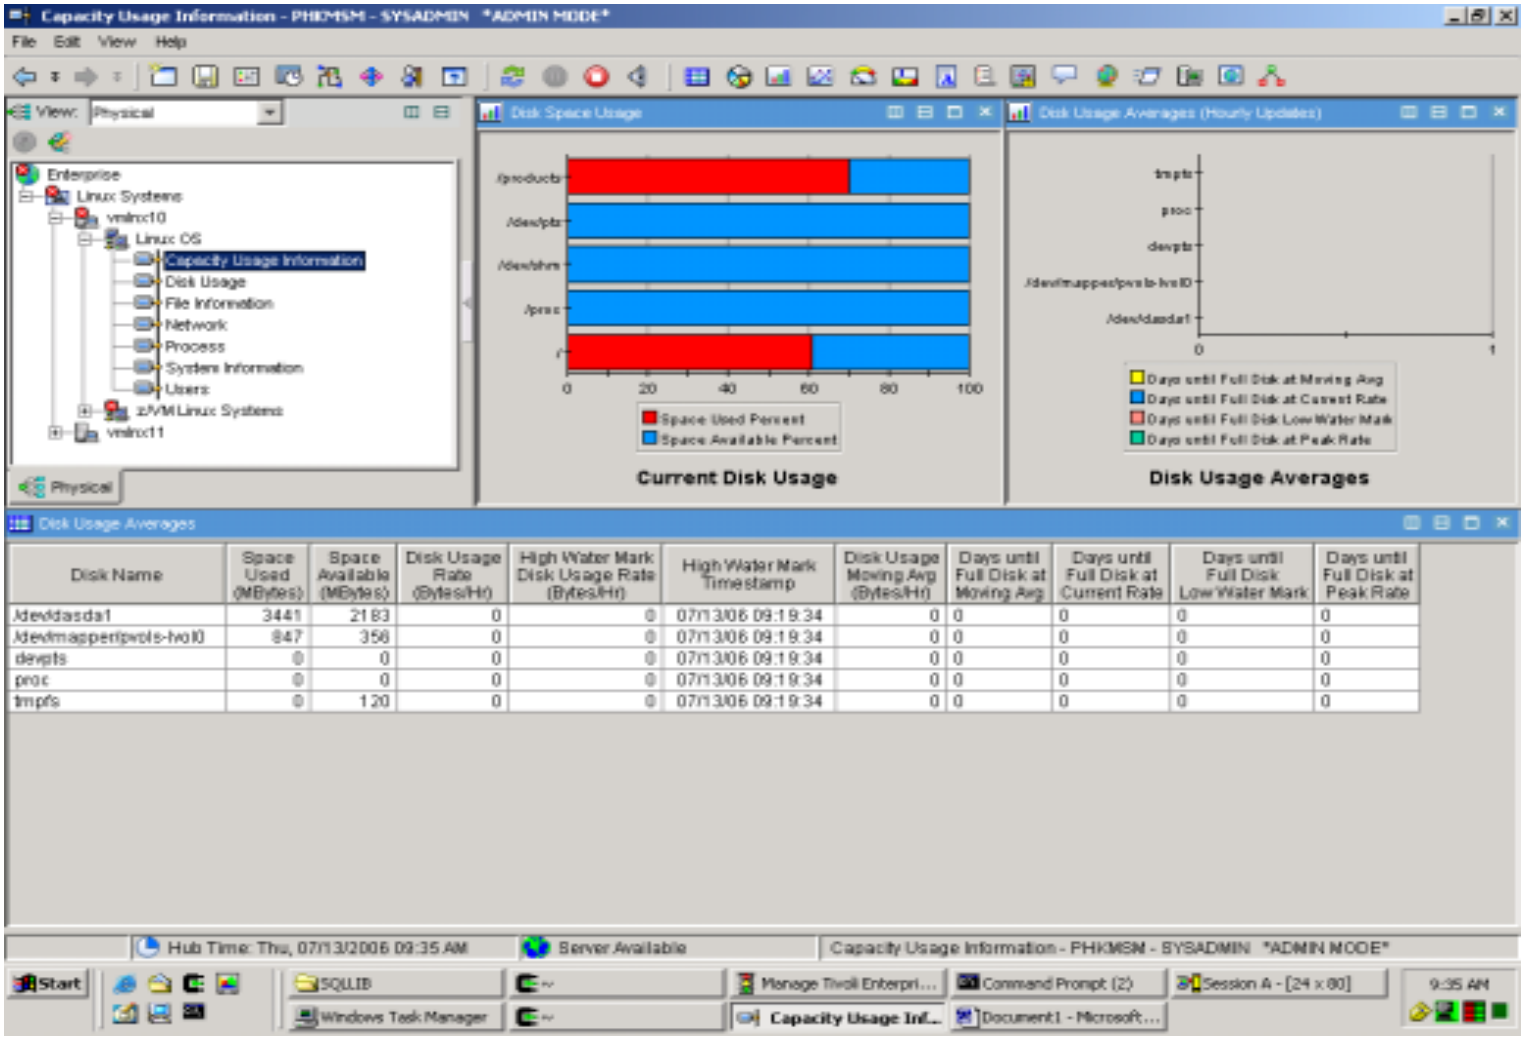

Capacity Usage Workspace

- **Capacity Usage Information Workspace**
  - This workspace displays the health of your system by providing CPU, disk, and swap space usage statistics.

- **Description**
  - The name of the physical disk partition where the file system is mounted.
  - The amount of disk space currently in use on a file system, expressed in megabytes.
  - The amount of unused space currently available to non-superusers on a file system, expressed in megabytes.
  - The bytes per hour of disk usage over the last sample period.
  - The bytes per hour rate that represents the high water mark of disk usage.
  - The date and time that the disk usage reaches a high water mark
  - The bytes per hour of disk usage averaged over all previous samples.
  - The number of days until the disk is full based on the moving average rate of disk usage.
  - The number of days until the disk is full based on the current rate of disk usage.
  - The number of days until the disk is full based on the disk usage rate that represents the low water mark.
  - The number of days until the disk is full based on the peak rate of disk usage.

CPU Averages

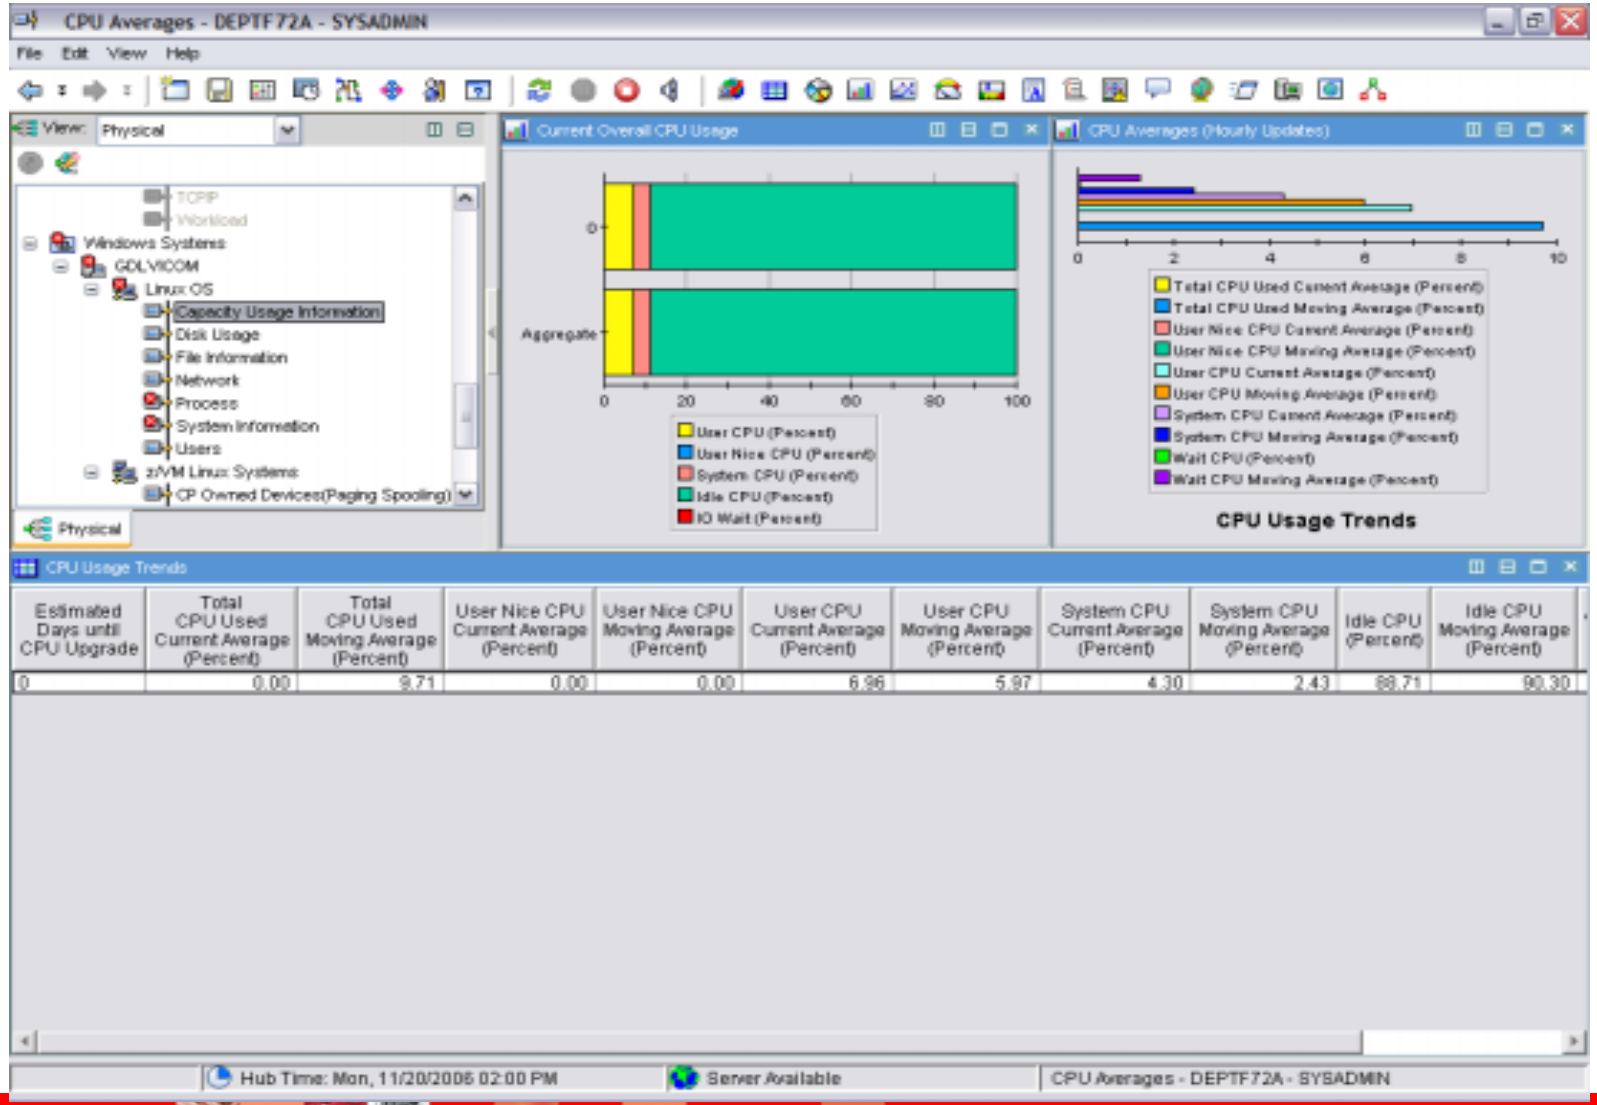

CPU Averages Workspace

- **CPU Averages Workspace**
  - This workspace displays the average CPU usage and shows trends that might indicate whether the usage is increasing
- **Description**
  - The number of days until CPU Usage Moving average hits 100% rate.
  - The current average of CPU usage, expressed as a percentage.
  - The moving average of CPU usage, expressed as a percentage
  - The current average of the user nice CPU time, expressed as a percentage
  - The moving average of the user nice CPU time, expressed as a percentage.
  - The current average of the user CPU time, expressed as a percentage.
  - The moving average of the user CPU time, expressed as a percentage.
  - The current average of the system CPU time, expressed as a percentage
  - The moving average of the system CPU time, expressed as a percentage.
  - The current average of the system's idle CPU time, expressed as a percentage.
  - The moving average of the system's idle CPU time, expressed as a percentage.
  - The current average of the wait CPU time, expressed as a percentage.
  - The moving current average of the wait CPU time, expressed as a percentage.

Virtual Memory Usage Trends

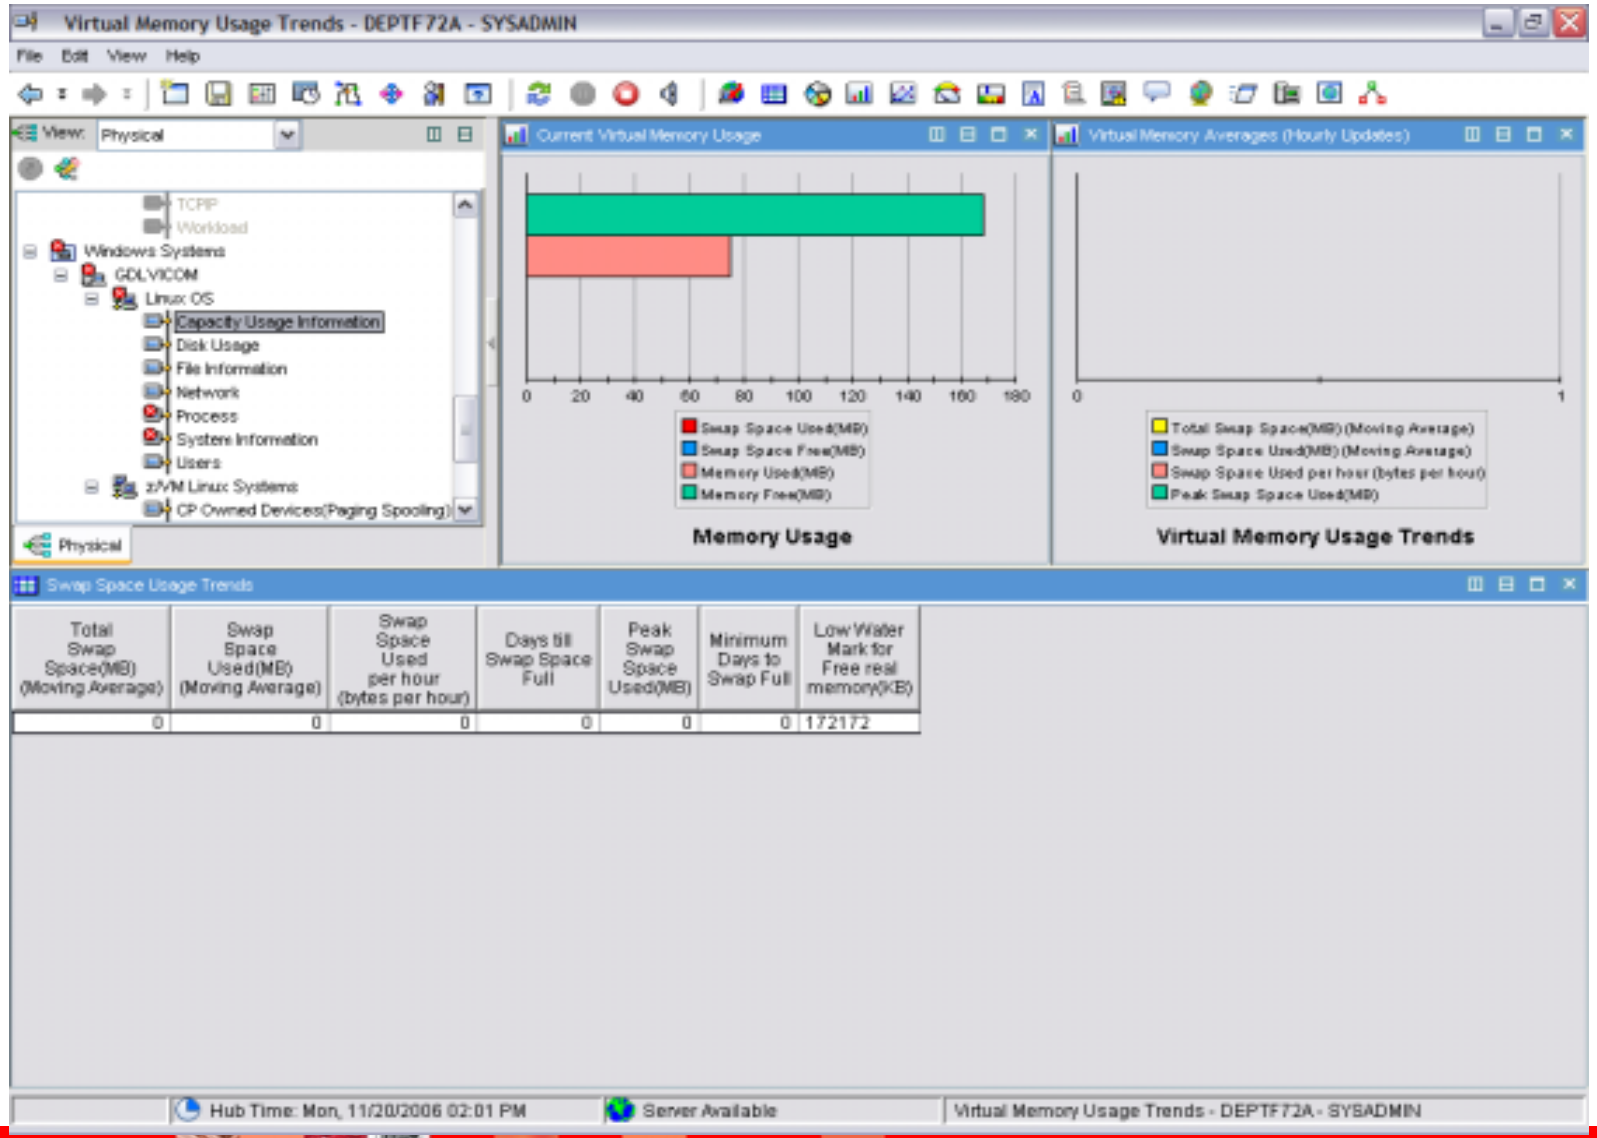

Virtual Memory Usage Trends Workspace

- **Virtual Memory Usage Trends Workspace**
  - This workspace displays information about current memory usage and usage and swap space usage trends
- **Description**
  - The moving average of total swap space, expressed in megabytes.
  - The moving average of swap space used, expressed in megabytes
  - The swap space usage rate, expressed in bytes per hour.
  - The predicted number of days until swap space is completely used (moving average).
  - The minimum number of days until swap space is completely used (peak rate based).
  - The lowest level that free real memory has reached, expressed in kilobytes.

Disk Usage

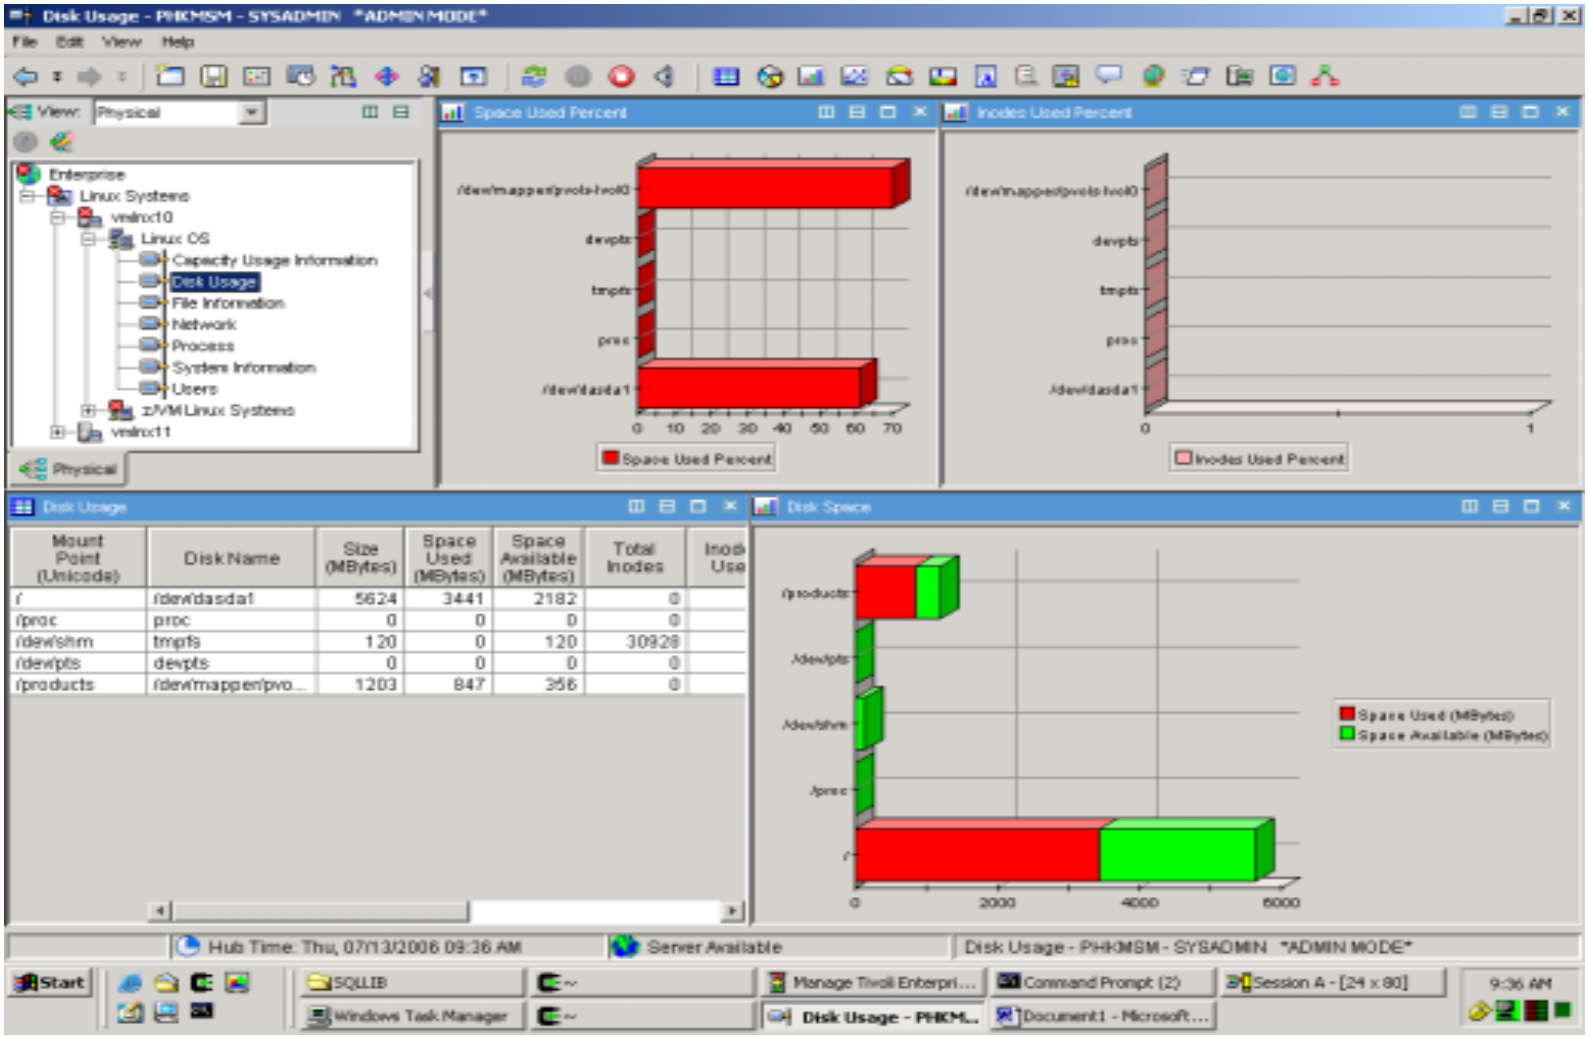

Disk Usage Workspace

- **Disk Usage Workspace**
  - This workspace displays information about the health of storage space within your monitored systems
- **Description**
  - The path name of the directory to which a file system is mounted. This is the virtual name for the directory.
  - The name of the physical disk partition where the file system is mounted
  - The total size of a file system, expressed in megabytes.
  - The amount of disk space currently in use on a file system, expressed in megabytes
  - The amount of unused space currently available to non-superusers on a file system, expressed in megabytes
  - The number of inodes allocated on a file system
  - The number of inodes currently allocated to files on the file system.
  - The number of inodes currently available on your file system.
  - The space currently used on the file system, expressed as a percentage of the sum of used and available space.
  - The percentage of inodes currently allocated to files, calculated by dividing the Inodes Used value by the Total Inodes value.
  - The file system type, such as hsfs, nfs, tmpfs, and ufs.
  - The amount of unused space currently available to non-superusers on a file system, expressed as a percentage.

File Information

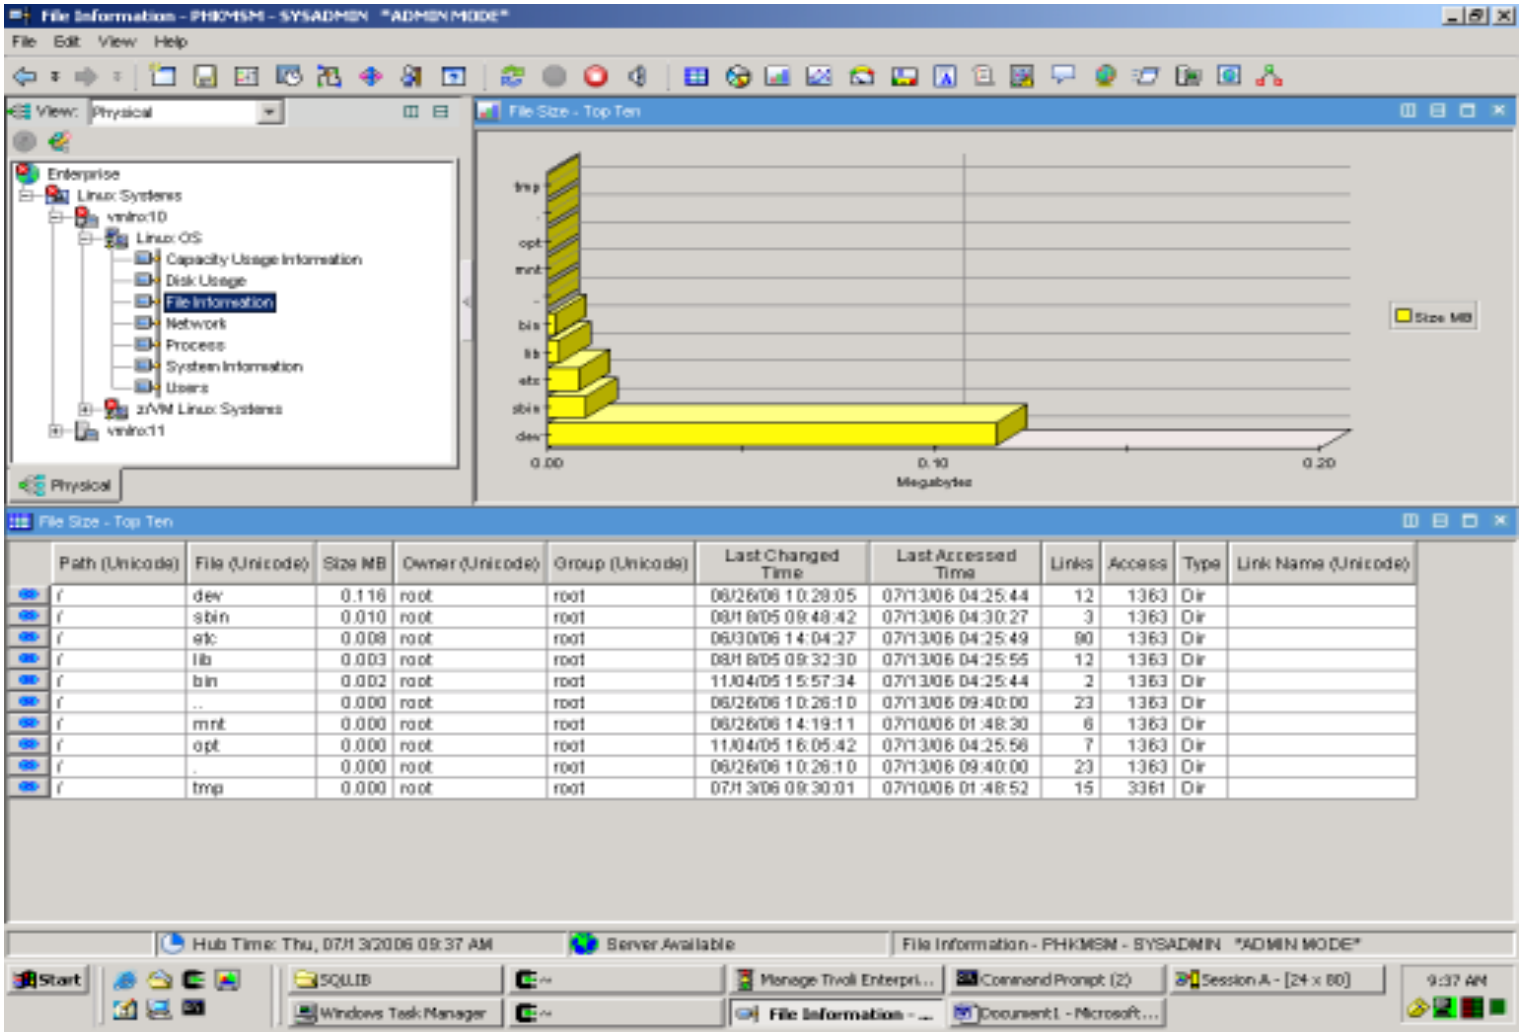

File Information Workspace

- **File Information Workspace**
  - This workspace displays information about the top ten files in size on your system
- **Description**
  - The path name of the directory to which a file system is mounted. This is the virtual name for the directory.
  - The name of the file
  - The size of the file in megabytes
  - The owner of the file
  - The group that the file belongs to
  - The date of the last time the file was changed
  - The date of the last time the file was accessed
  - Then number of links to the file
  - The access permissions for the file
  - Whether the file is a file or a directory
  - The link name of the file

File Information

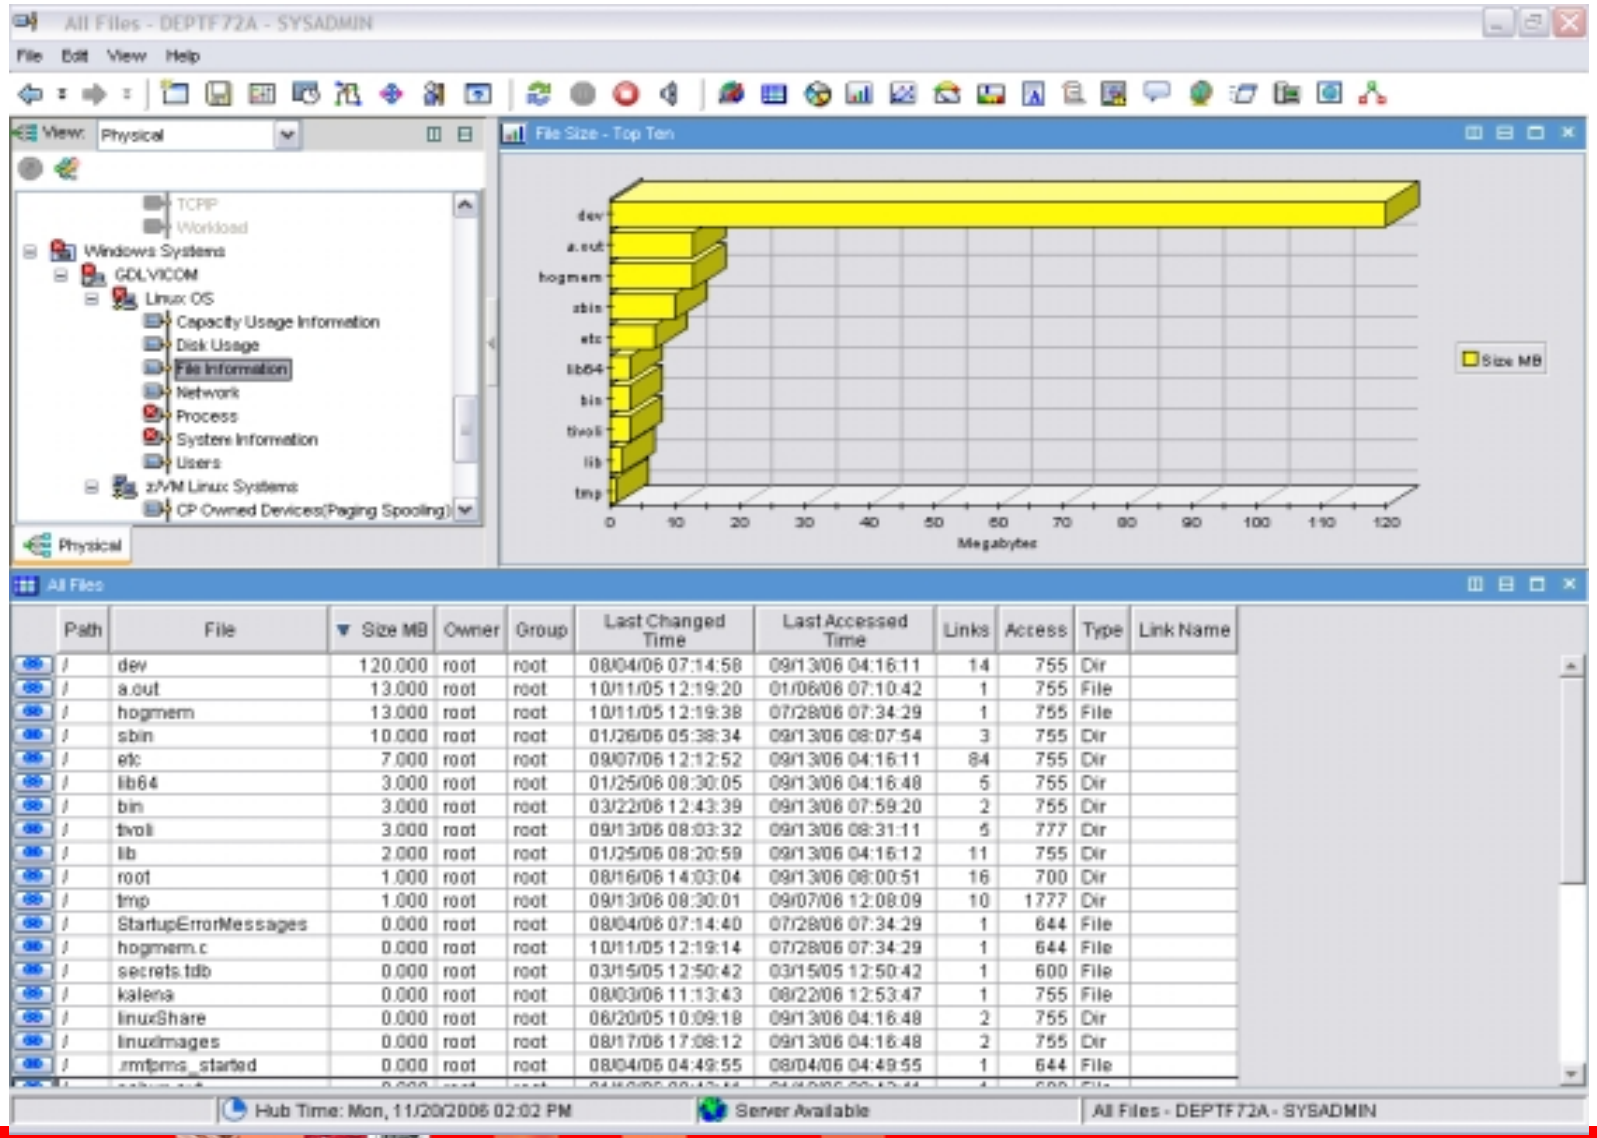

All Files Information Workspace

- **All Files Information Workspace**
  - This workspace displays information about all the files on your system
- **Description**
  - The path name of the directory to which a file system is mounted. This is the virtual name for the directory.
  - The name of the file
  - The size of the file in megabytes
  - The owner of the file
  - The group that the file belongs to
  - The date of the last time the file was changed
  - The date of the last time the file was accessed
  - Then number of links to the file
  - The access permissions for the file
  - Whether the file is a file or a directory
  - The link name of the file

Network

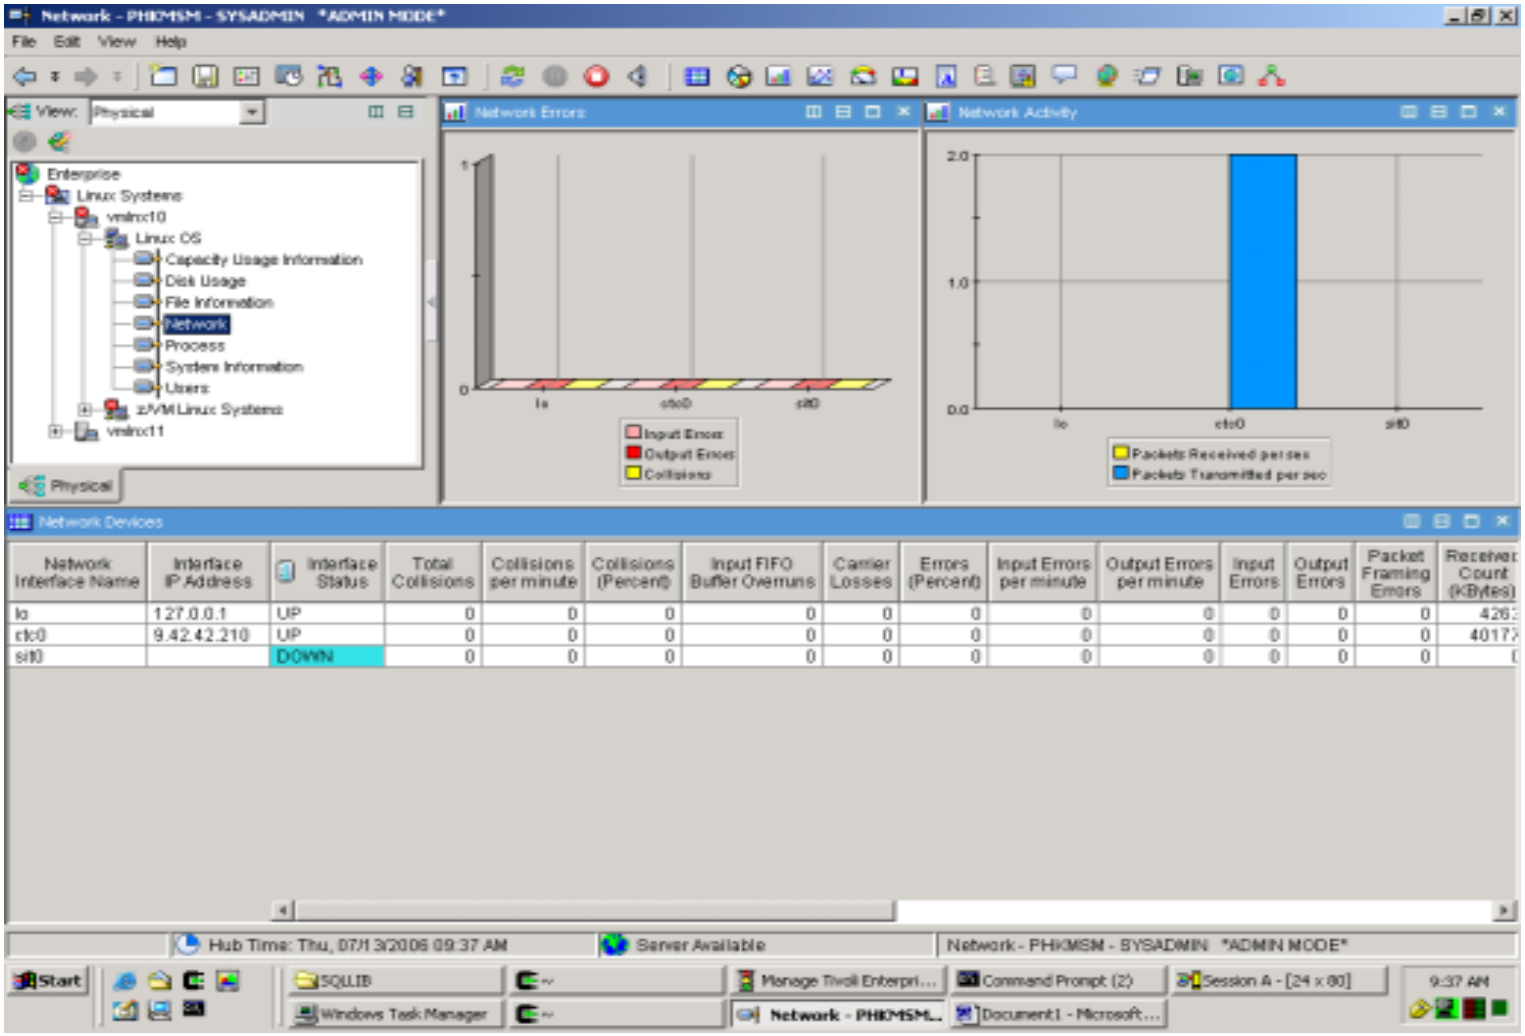

Network Workspace

- **Network Workspace**
  - This workspace displays information about the network components within your monitored systems
- **Description**
  - The Dynamic Name Server (DNS) entry associated with the IP address of the network interface.
  - The Internet Protocol (IP) address of the network interface.
  - An indication of whether or not a network interface is currently available
  - The number of times during the sampling period that a packet transmitted by the network interface collided with another packet.
  - The number of times a packet collided with another packet per minute.
  - Of the total number of packets transmitted in this sample period, the percentage involved in a collision.
  - The number of input FIFO buffer overruns that occurred during the sampling period.
  - The number of carrier losses that occurred in the interface
  - Of the total number of packets received and transmitted, the percentage that were in error during this sample period.
  - The number of packets with errors received per minute by the interface.
  - The number of packet transmission errors per minute during the monitoring interval.
  - The number of packets received with errors in the interface.
  - The number of packets packet transmission errors in the interface.
  - The number of packet framing errors that occurred in the interface.
  - The number of packets received by the interface during the sampling period
  - The number of bytes received per second by the interface.
  - The number of kilobytes transmitted by an interface since boot time.
  - The number of bytes received per second by the interface.
  - The number of packets received by the interface during the sampling period.
  - The number of packets received per second by the interface.
  - The number of input packets dropped by the device driver.
  - The number of output packets dropped by the device driver.
  - The number of output FIFO buffer overruns that occurred during the sampling period.
  - The number of packets transmitted by the interface during the sampling period.
  - The number of packets transmitted per second by the interface.
  - The maximum packet size (in bytes) for the specified network interface

Sockets Information

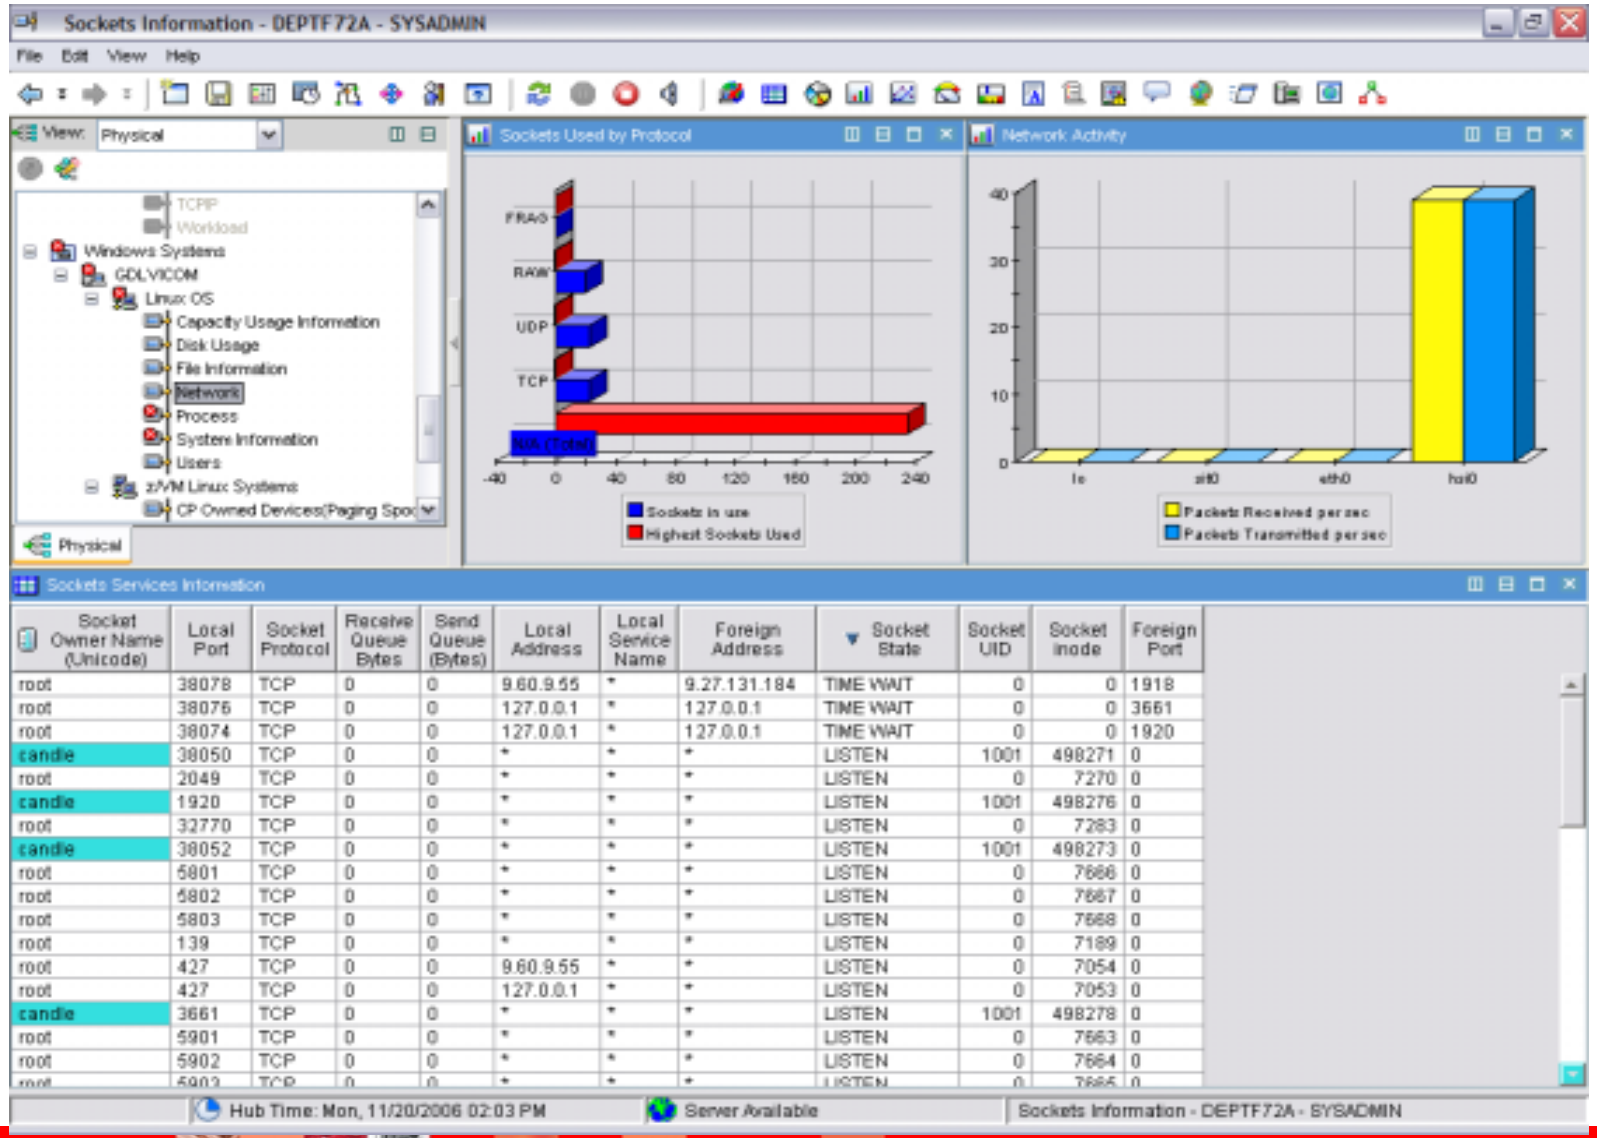

Socket Owner Name (Unicode)	Local Port	Socket Protocol	Receive Queue Bytes	Send Queue (Bytes)	Local Address	Local Service Name	Foreign Address	Socket State	Socket UID	Socket inode	Foreign Port
root	38078	TCP	0	0	9.60.9.55	*	9.27.131.184	TIME WAIT	0	0	1918
root	38076	TCP	0	0	127.0.0.1	*	127.0.0.1	TIME WAIT	0	0	3661
root	38074	TCP	0	0	127.0.0.1	*	127.0.0.1	TIME WAIT	0	0	1920
candle	38050	TCP	0	0	*	*	*	LISTEN	1001	498271	0
root	2049	TCP	0	0	*	*	*	LISTEN	0	7270	0
candle	1920	TCP	0	0	*	*	*	LISTEN	1001	498276	0
root	32770	TCP	0	0	*	*	*	LISTEN	0	7283	0
candle	38052	TCP	0	0	*	*	*	LISTEN	1001	498273	0
root	5801	TCP	0	0	*	*	*	LISTEN	0	7666	0
root	5802	TCP	0	0	*	*	*	LISTEN	0	7667	0
root	5803	TCP	0	0	*	*	*	LISTEN	0	7668	0
root	139	TCP	0	0	*	*	*	LISTEN	0	7189	0
root	427	TCP	0	0	9.60.9.55	*	*	LISTEN	0	7054	0
root	427	TCP	0	0	127.0.0.1	*	*	LISTEN	0	7053	0
candle	3661	TCP	0	0	*	*	*	LISTEN	1001	498278	0
root	5901	TCP	0	0	*	*	*	LISTEN	0	7663	0
root	5902	TCP	0	0	*	*	*	LISTEN	0	7664	0

Sockets Information Workspace

- **Sockets Information Workspace**
  - This workspace displays information about the socket connections within your monitored systems
- **Description**
  - The user name associated with the user ID that owns or started the socket connection.
  - The local port number.
  - Protocol used by the socket.
  - The count of bytes not copied by the user program connected to this socket.
  - The count of bytes not acknowledged by the remote host.
  - The address of the local end of the socket, presented as a dotted IP address.
  - The local port number translated to a service name from the etc/services subdirectory.
  - The address of the remote end of the socket.
  - The state of the socket.
  - The user ID of the owner of the socket.
  - The inode used by the socket.
  - The number of the foreign port.

RPC Statistics

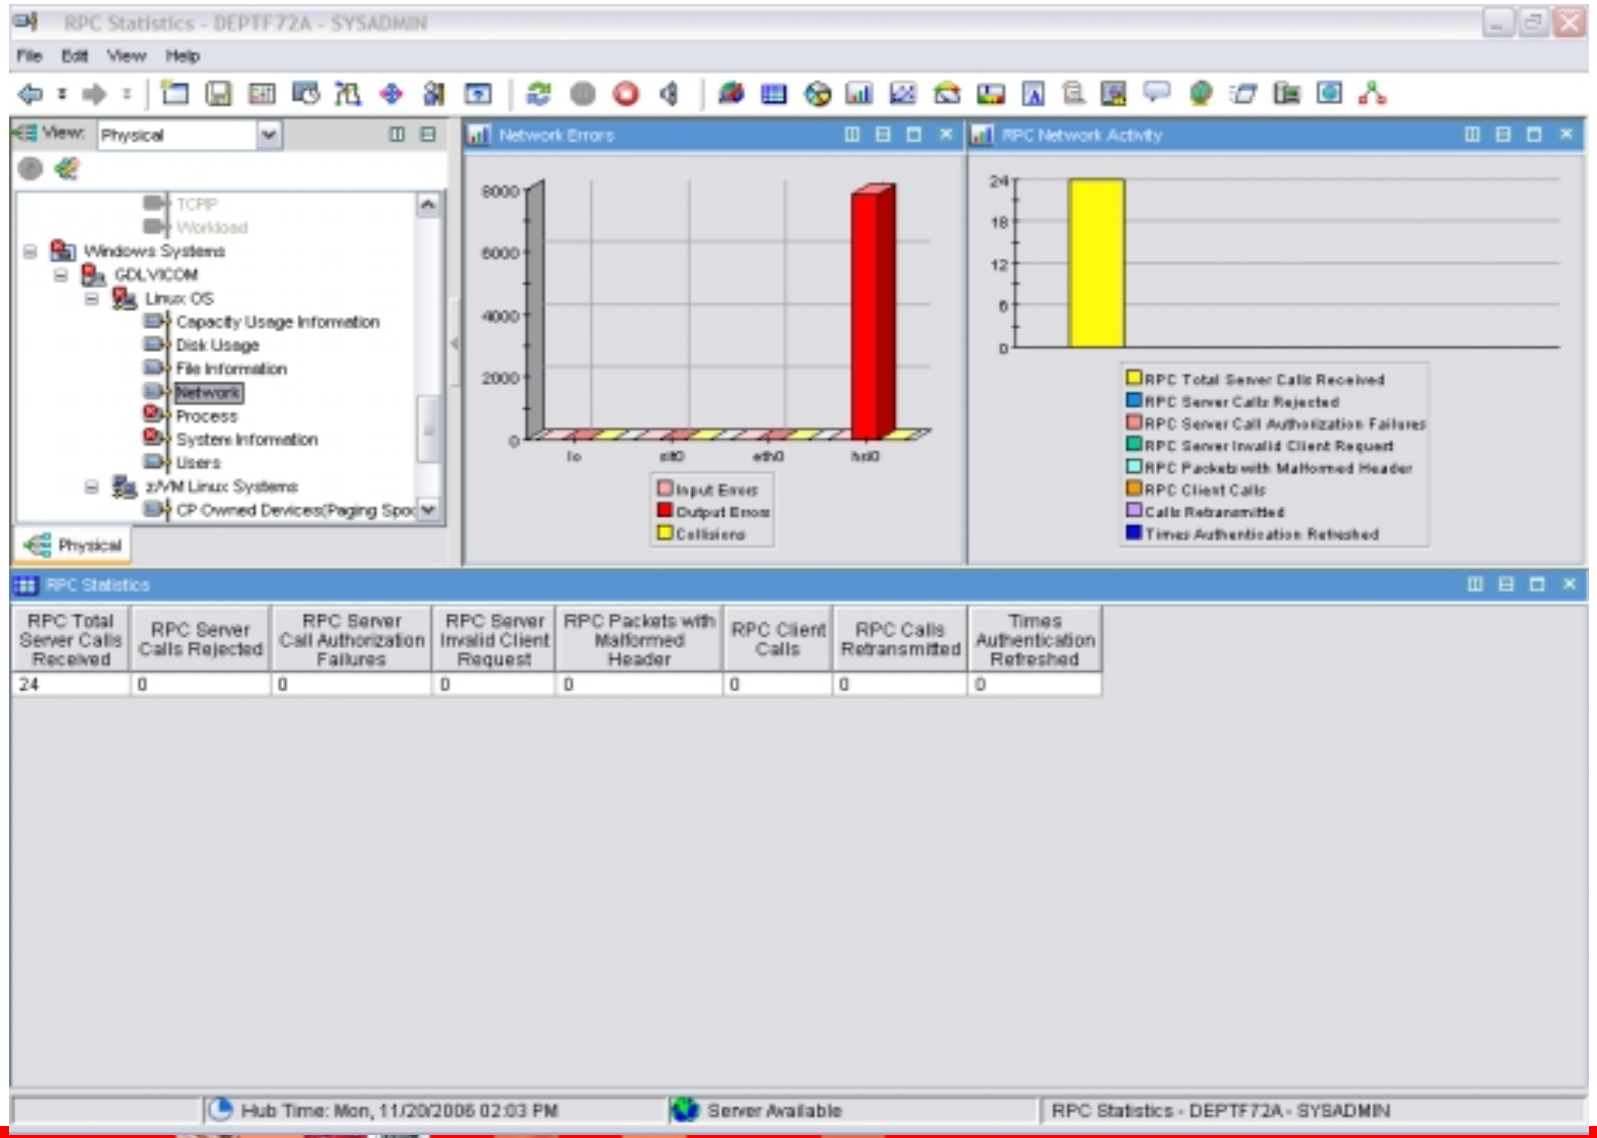

RPC Statistics Workspace

- **RPC Statistics Workspace**
  - This workspace displays statistics on the number and type of remote procedure calls being made to the server and clients
  - **Description**
    - The total number of calls made to the server (both valid and not valid).
    - The number of calls made to the server, which were rejected.
    - The number of packets that were received at the server with authorizations that were not valid.
    - The number of packets that were received at the server, which had client requests that were not valid
    - The number of packets that were received at the server with header records that were not properly formatted.
    - The number of calls to the server made by the server's clients.
    - The number of client calls that needed to be transmitted again.
    - The number of times the authentication of a client was refreshed.

NFS Statistics

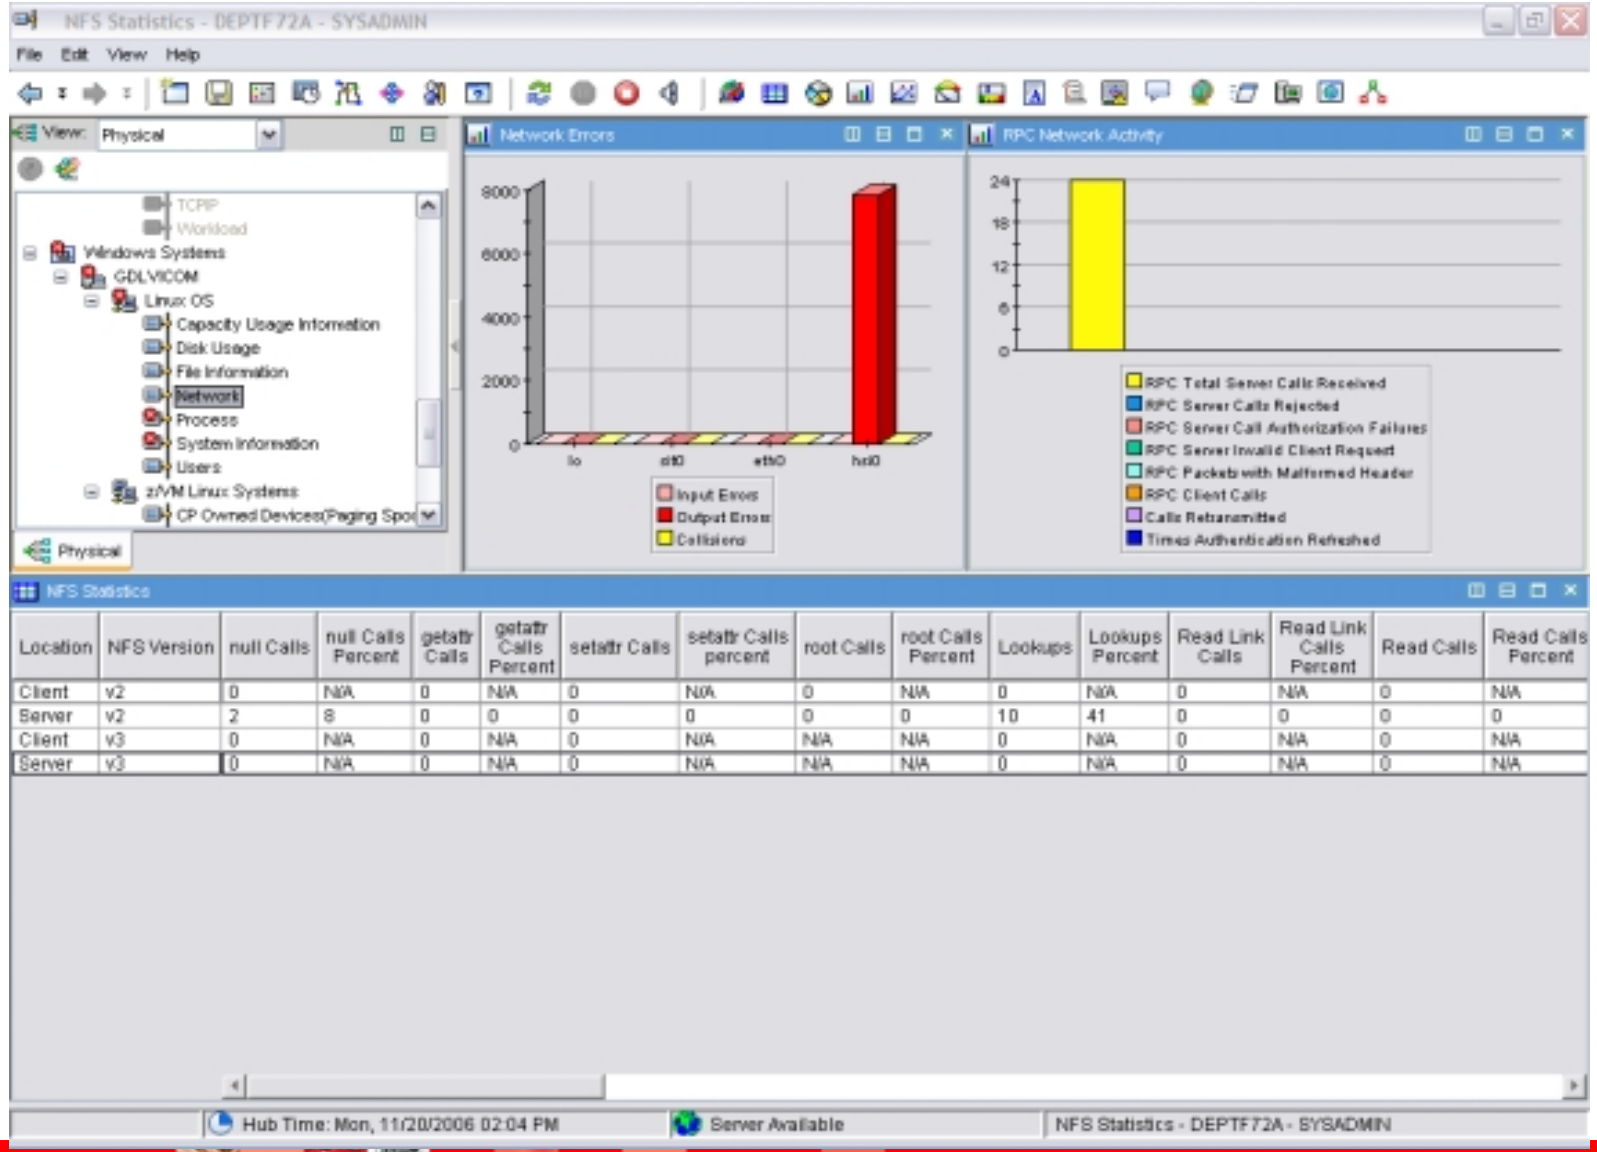

NFS Statistics Workspace

- **NFS Statistics Workspace**
  - This workspace displays statistics on the operations involving the Network File System
  - **Description**
    - The location of the origin of the call in the Network File System.
    - The software version associated with the NFS server.
    - The number of calls made to the NFS server from NFS clients which contained no data
    - Of the total number of calls made to the NFS server, the percentage that contained no data.
    - The number of calls made to the NFS server which contained a get attribute (getattr) operation.
    - Of the total number of calls made to the NFS server, the percentage that contained get attribute (getattr) operations.
    - The number of calls made to the NFS server which contained a set attribute (setattr) operation.
    - Of the total number of calls made to the NFS server, the percentage that contained a set attribute (setattr) operation.
    - The number of calls made to the NFS server which contained root calls.
    - Of the total number of calls made to the NFS server, the percentage that were root calls
    - The number of read directory plus (readdirplus) calls made to the NFS server to return the name, the file ID, attributes, and file handle.
    - The number of total calls and percentage of calls that were:
      - Lookups
      - Read link
      - Read
      - Write cache
      - Writes
      - File creates
      - Remove files
      - Rename files
      - Link
      - Symbolic link
      - Make directory
      - Remove directory
      - Read directory
      - File system statistics
      - Access
      - Make node
      - File system info
      - Pathconf
      - Commit

Process

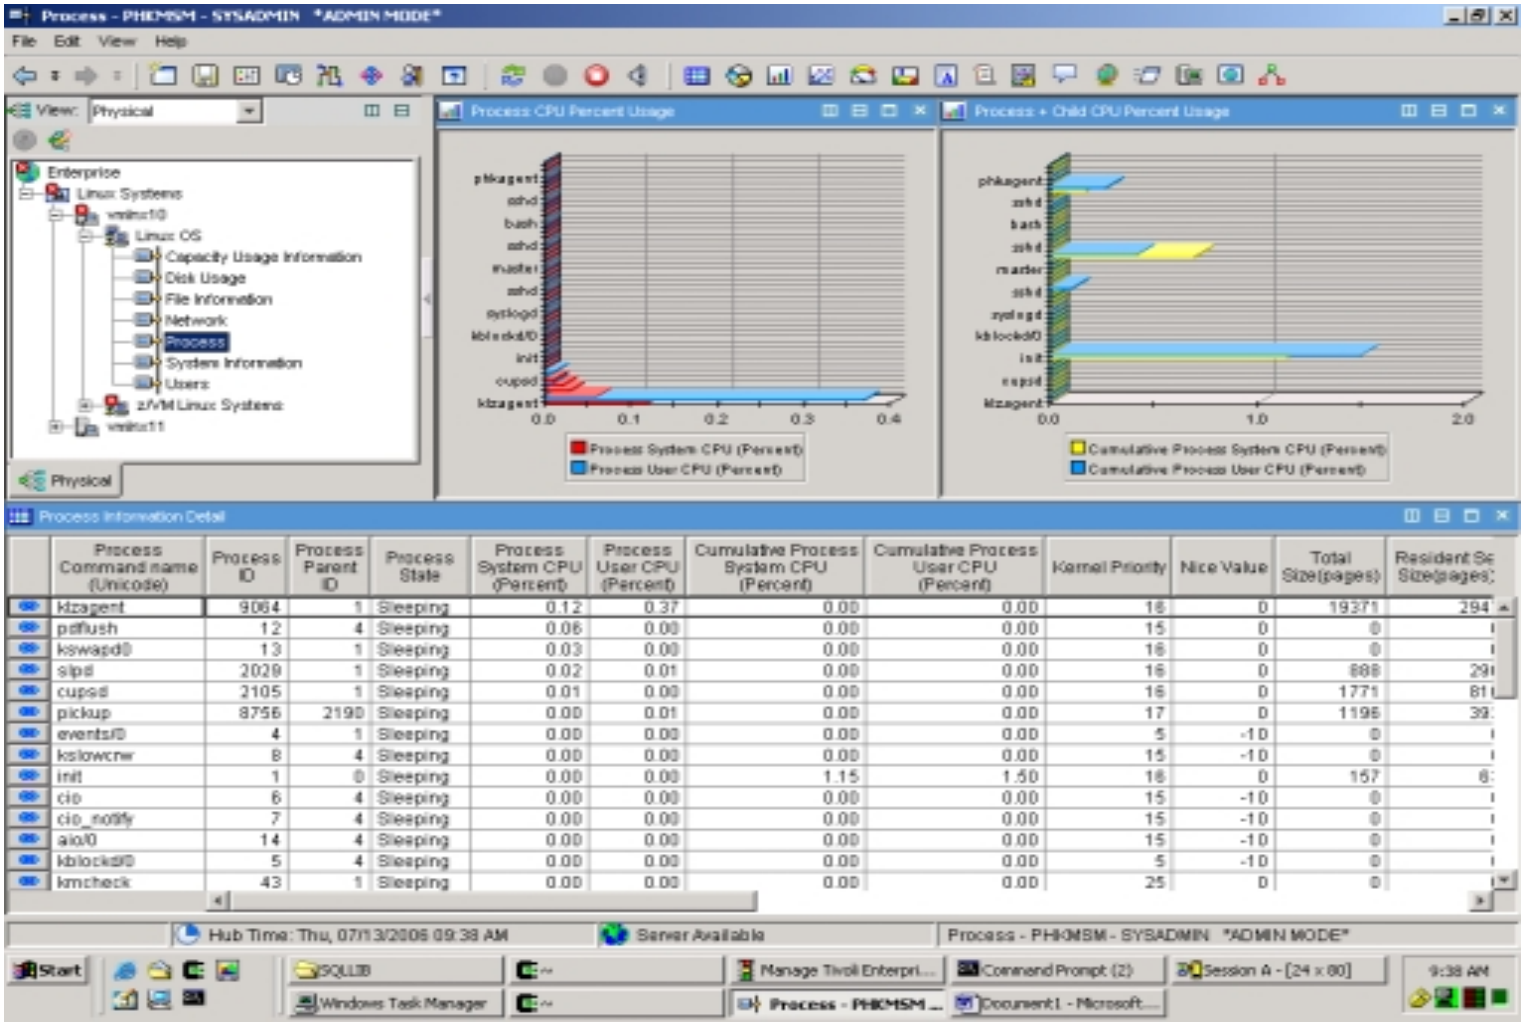

Process Workspace

- **Process Workspace**
  - This workspace displays the health of specific processes within your monitored systems
- **Description**
  - The name of the process command.
  - The identifier of the process.
  - The identifier for the parent process.
  - The state of the process (Sleeping, Disk, Running, Zombie, Trace, Dead, or N/A).
  - The percentage of CPU time spent in kernel mode by process.
  - The percentage of CPU time spent in user mode by process.
  - The percentage of cumulative CPU time spent in kernel mode by process.
  - The percentage of cumulative CPU time spent in user mode by process.
  - The kernel scheduling priority.
  - The standard Linux nice level.
  - The number of pages that the process has in real memory.
  - The number of pages the process has in real memory.
  - The number of pages of shared (mmap'd) memory.
  - The number of pages of text resident (mmap'd) memory.
  - The number of pages of shared (mmap'd) memory.
  - The size of the data set based on the number of pages.
  - Pages that have been modified (dirty) in buffer (main memory), but not yet copied to the cache
  - The data size (in kilobytes) of the virtual memory.
  - The size (in kilobytes) of locked pages of the virtual memory
  - The data size (in kilobytes) of the virtual memory.
  - The stack size (in kilobytes) of the virtual memory.
  - The executable size (in kilobytes) of the virtual memory.
  - The library size (in kilobytes) of the virtual memory.
  - The total number of minor page faults (including child processes) since the start of the process.
  - The total number of major page faults (including child processes) since the start of the process.
  - The process command line string.
  - The ID of the process CPU.
  - Of the total system CPU usage, the percentage that was user CPU usage.

Process User

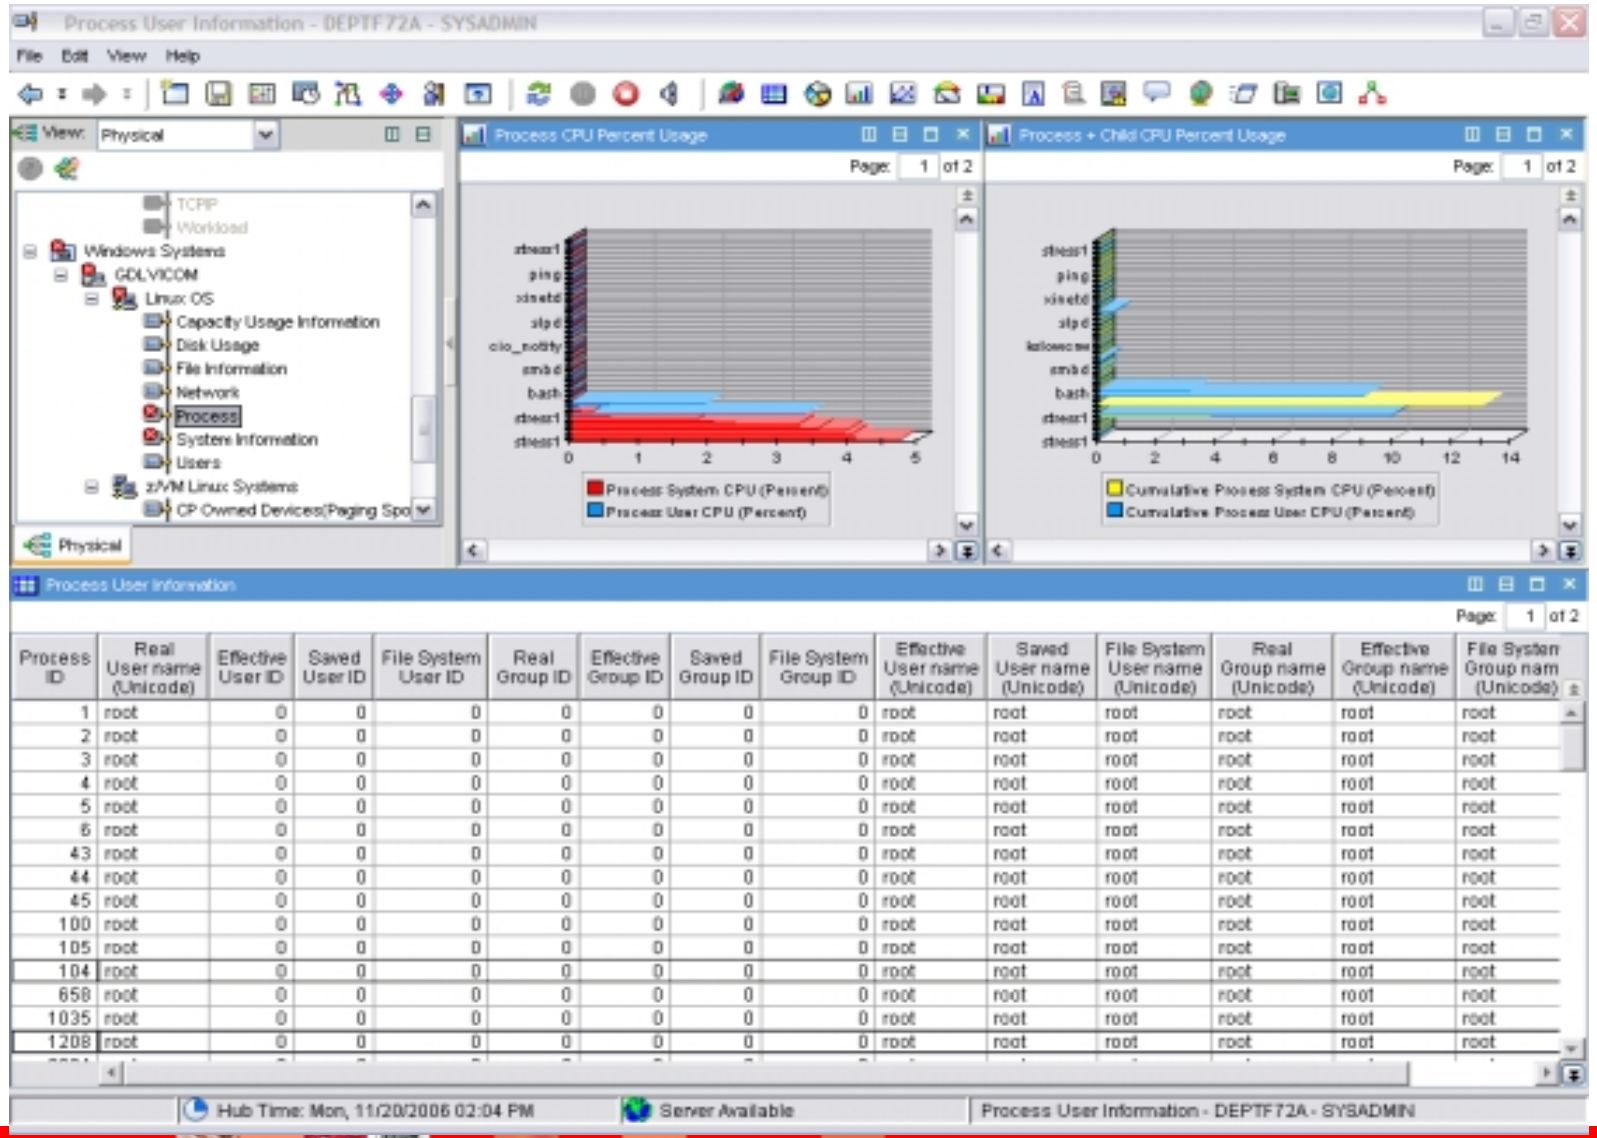

Process User Workspace

- **Process User Workspace**
  - This workspace displays process owners of your monitored Linux system and details their usage
- **Description**
  - The identifier associated with the process.
  - The identifier of the effective user.
  - The identifier of the saved user.
  - The identifier of the file system user.
  - The identifier of the real group.
  - The identifier of the effective group.
  - The identifier of the saved group.
  - The identifier of the file system group.
  - The name of the effective user.
  - The name of the saved user.
  - The name of the file system user.
  - The name of the real group.
  - The effective group name.
  - The name of the file system group.
  - The name of the saved group.

System Information

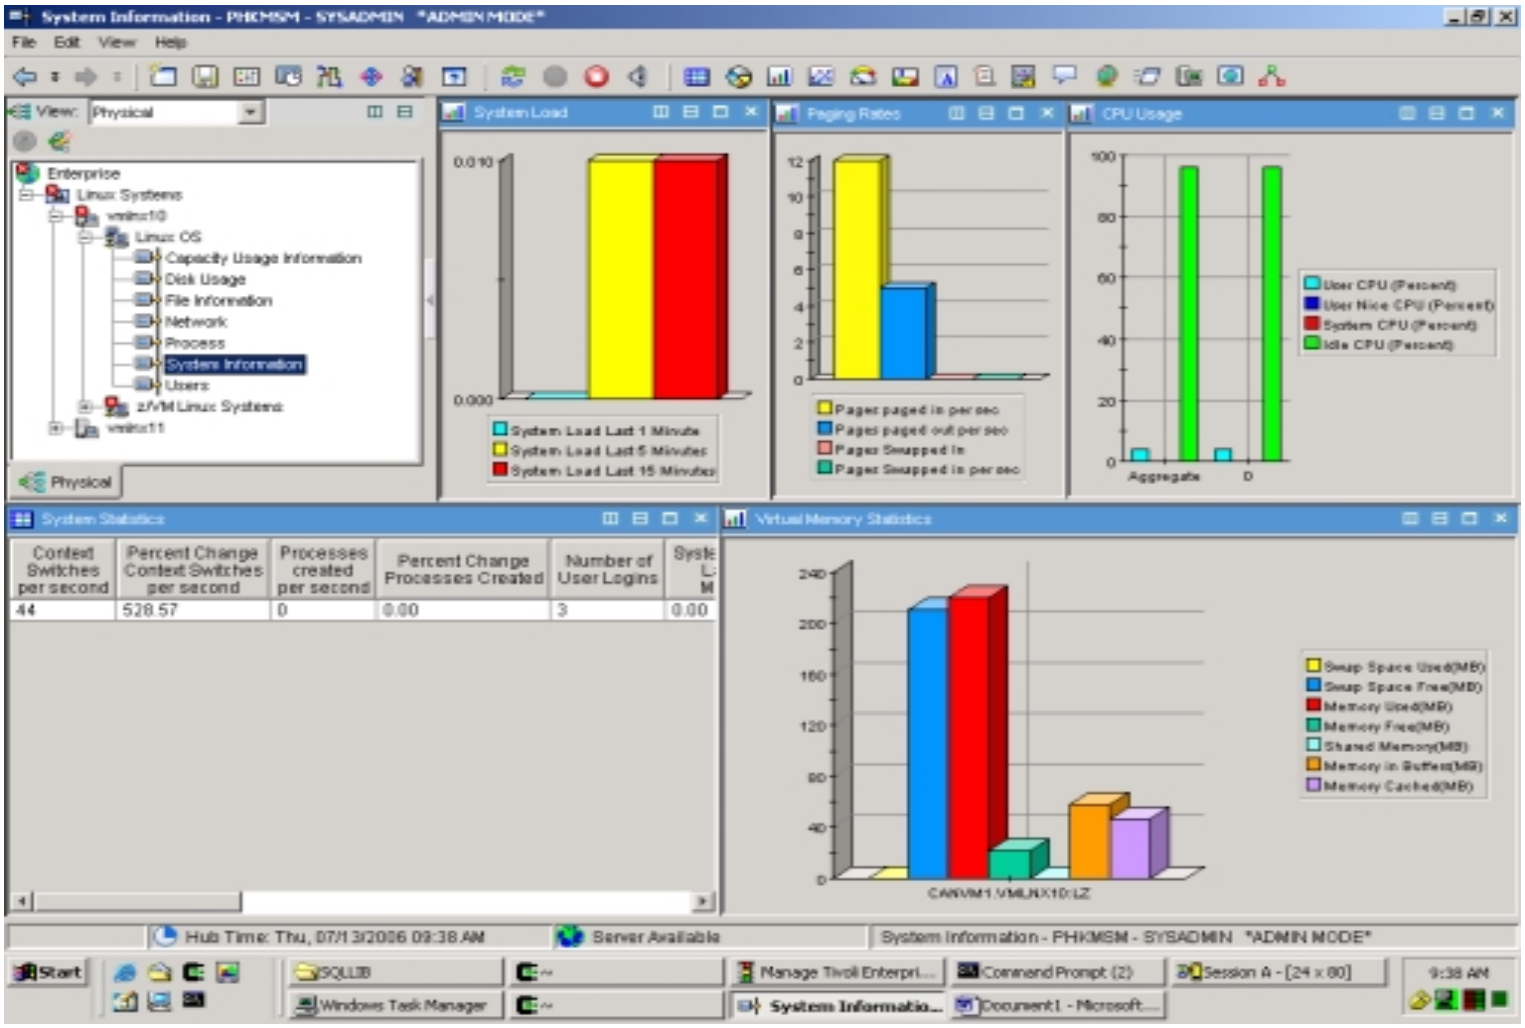

System Information Workspace

- **System Information Workspace**
  - This workspace displays data associated with CPU usage, system loads, and process creation
  - **Description**
    - The number of context switches per second.
    - The percentage change in the number of context switches per second.
    - The number of processes created per second.
    - The percentage change in the number of processes per second.
    - The current number of users logged in.
    - The load on the system for the last minute.
    - The load on the system for the last five minutes.
    - The load on the system for the last fifteen minutes.
    - The system uptime in number of seconds.
    - The total number of pages paged in.
    - The total number of pages paged in per second.
    - The total number of pages paged out.
    - The total number of pages paged out per second.
    - The total number of pages swapped in.
    - The total number of pages swapped in per second.
    - The total number of pages swapped out.
    - The total number of pages swapped out per second.
    - The total number of pages faults per second (both major and minor).
    - The total number of major faults per second.

Virtual Memory Statistics

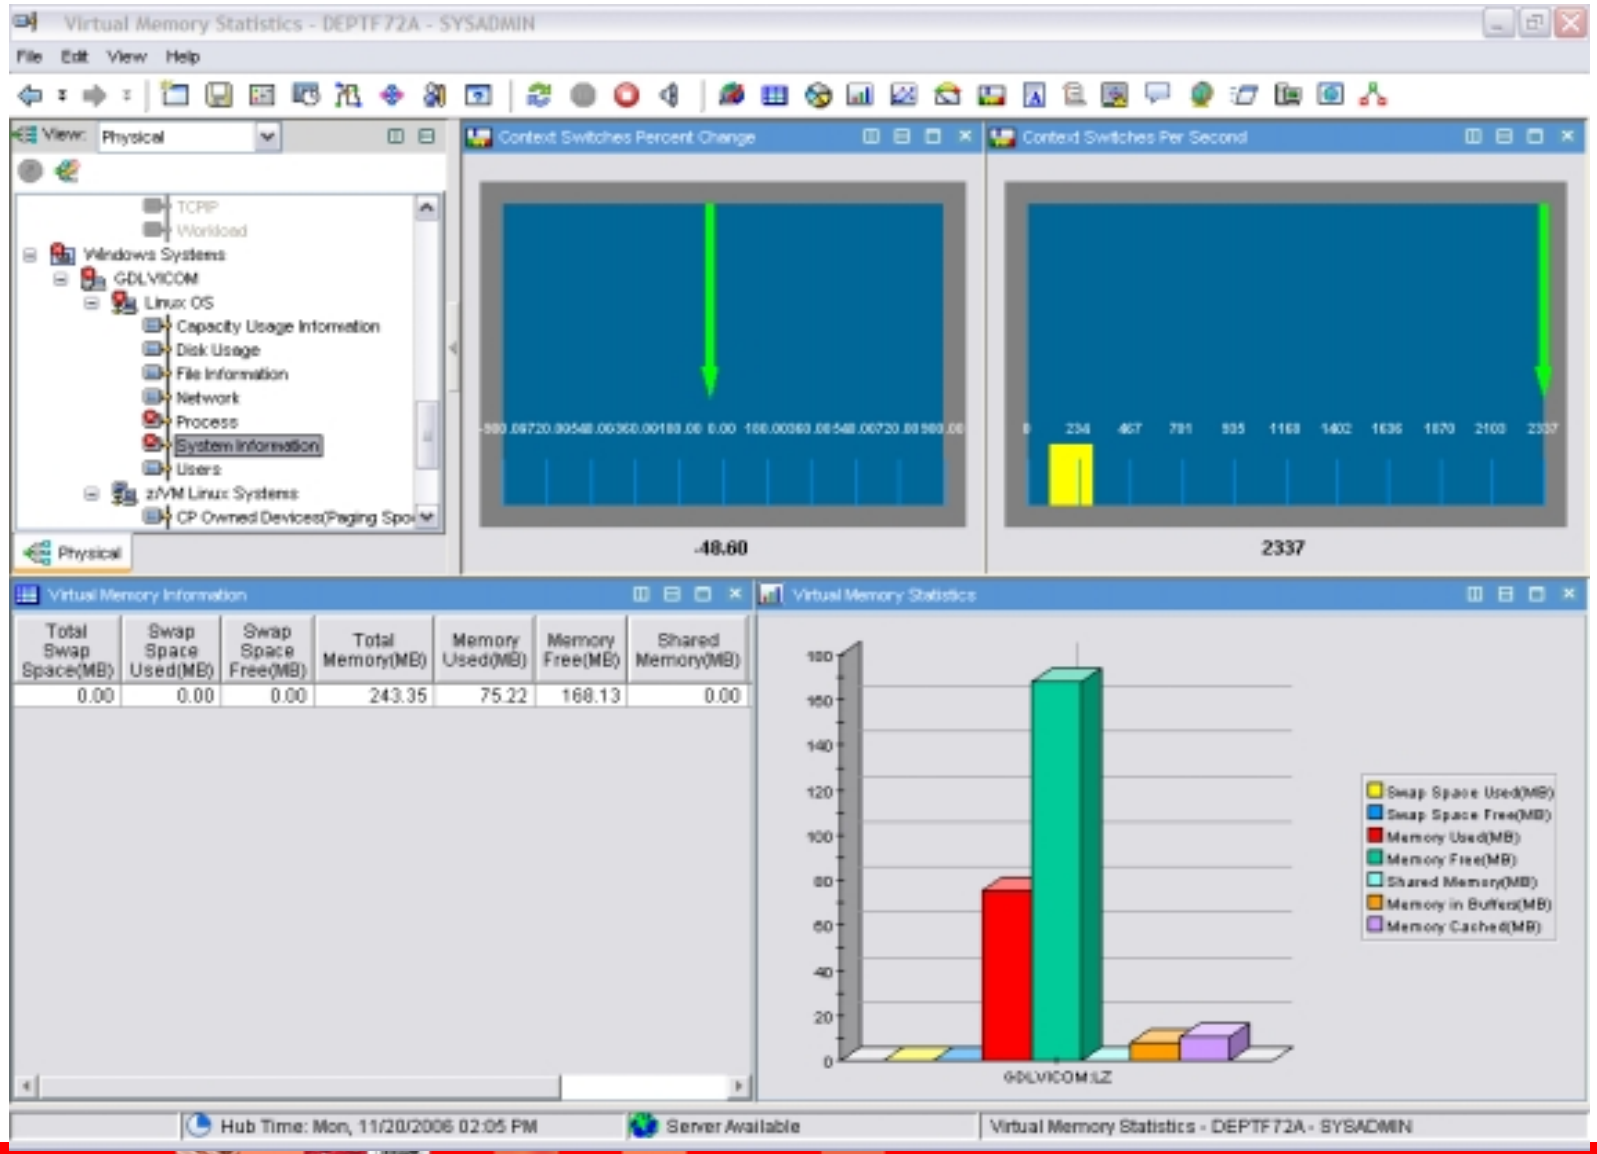

Virtual Memory Workspace

- **Virtual Memory Workspace**
  - This workspace displays data associated with memory usage.
- **Description**
  - The total size (in megabytes) of swap space.
  - The size (in megabytes) of swap space used.
  - The size (in megabytes) of swap space free.
  - The total size (in megabytes) of physical memory.
  - The size (in megabytes) of physical memory used.
  - The size (in megabytes) of physical memory free.
  - The size (in megabytes) of physical memory shared.
  - The size (in megabytes) of physical memory in buffers.
  - The size (in megabytes) of physical memory cached.

Disk I/O Rate

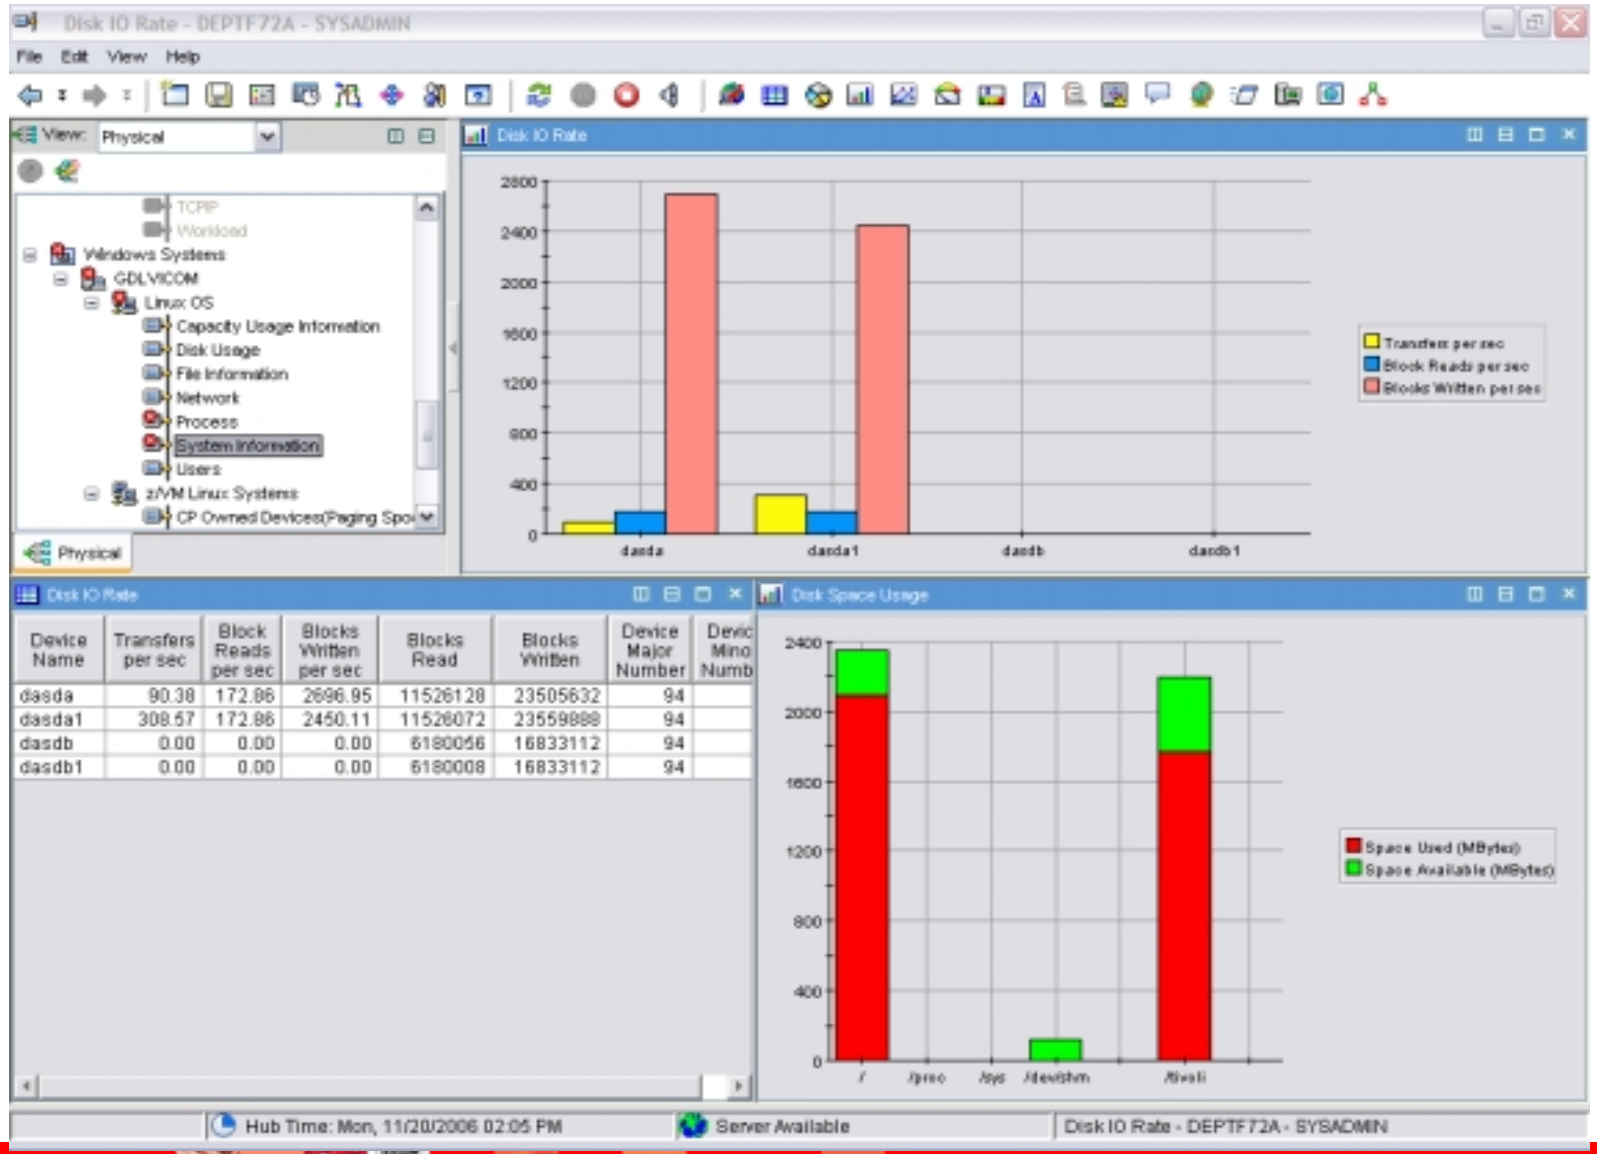

Disk I/O Rate Workspace

- **Disk I/O Rate Workspace**
  - This workspace displays input/output statistics, including the transfer rates, block read rates, and block write rates
- **Description**
  - The name of the device as it appears under the dev subdirectory.
  - The number of transfers per second that were issued to the device.
  - The amount of data read from the drive expressed in a number of blocks per second.
  - The amount of data written to the drive expressed in a number of blocks per second.
  - The total number of blocks read.
  - The total number of blocks written.
  - The major number of the device.
  - The distinctive minor number for device.

Disk I/O Extended Rate

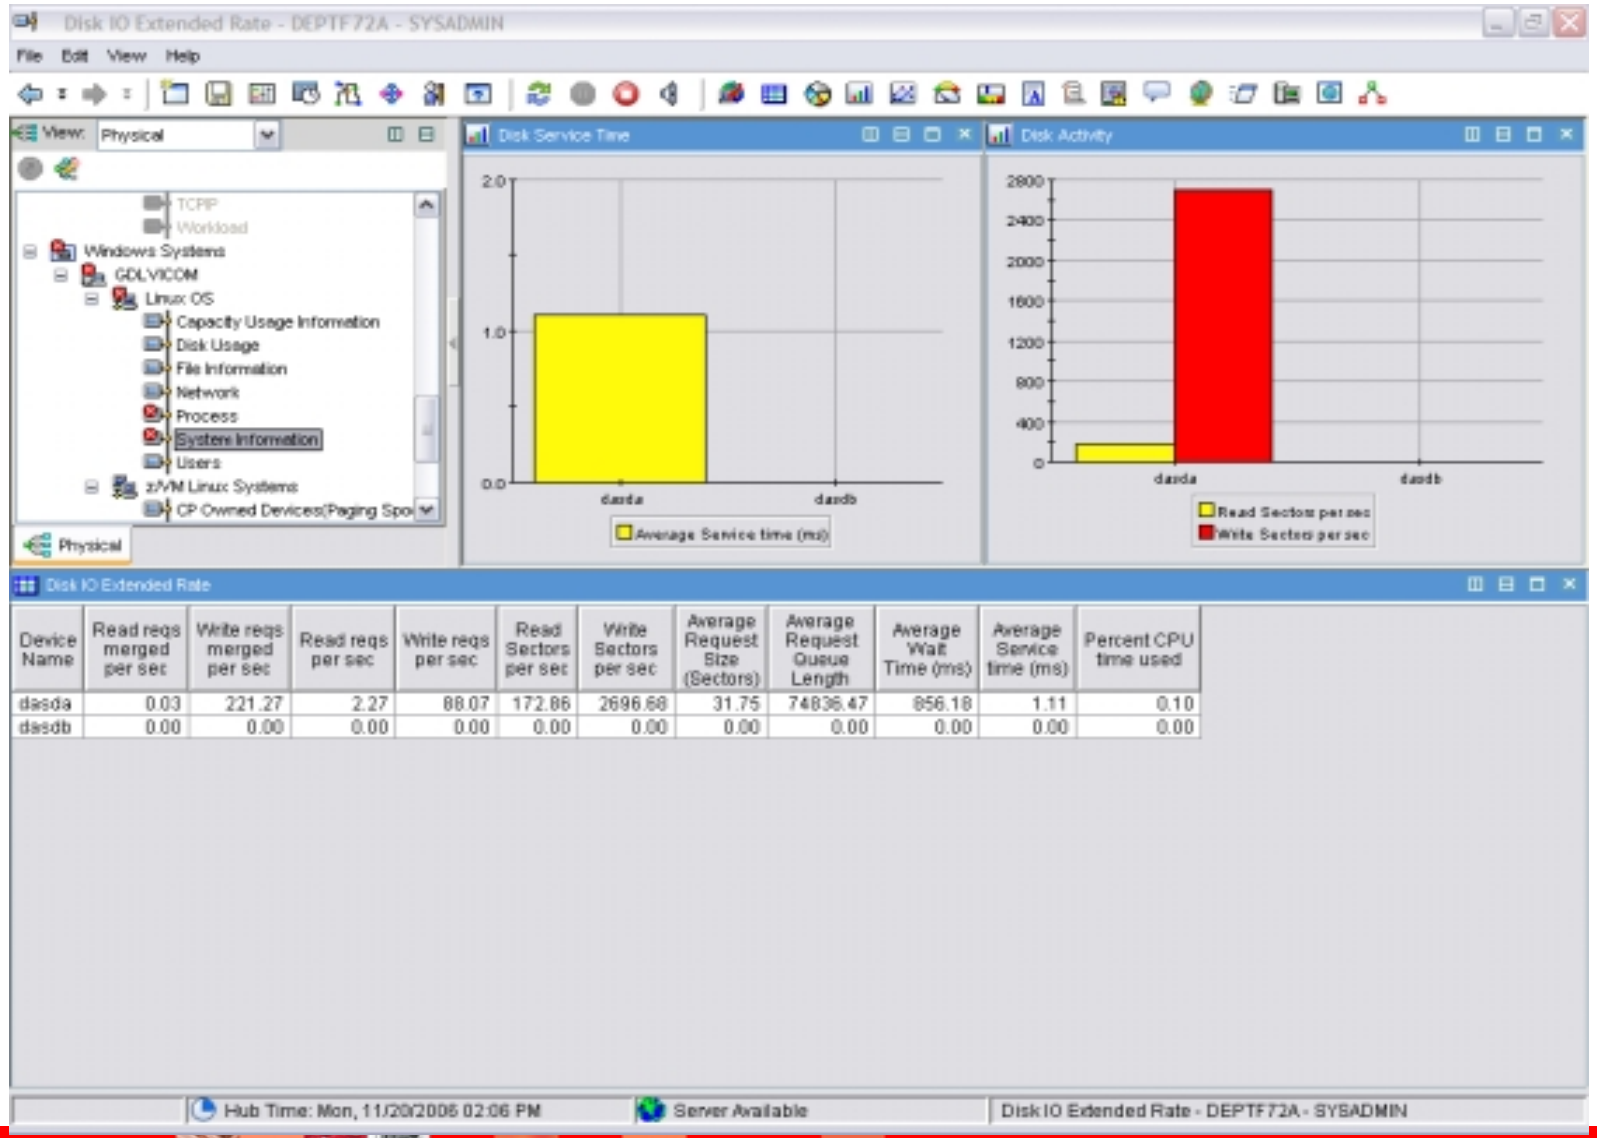

Disk I/O Extended Rate Workspace

- **Disk I/O Extended Rate Workspace**
  - This workspace displays input/output statistics and calculations
- **Description**
  - The name of the device as it appears under the dev subdirectory.
  - The number of read requests merged, per second, that were issued to the device.
  - The number of write requests merged that were issued, per second, to the device.
  - The number of read requests that were issued, per second, to the device.
  - The number of write requests that were issued, per second, to the device.
  - The number of sectors read, per second, from the device.
  - The number of sectors written to the device, per second.
  - The average size (in sectors) of the requests that were issued to the device.
  - The average queue length of the requests that were issued to the device.
  - The average time (in milliseconds) for I/O requests issued to the device to be served.
  - The average service time (in milliseconds) for I/O requests that were issued to the device.
  - Percentage of CPU time during which I/O requests were issued to the device.

System Configuration

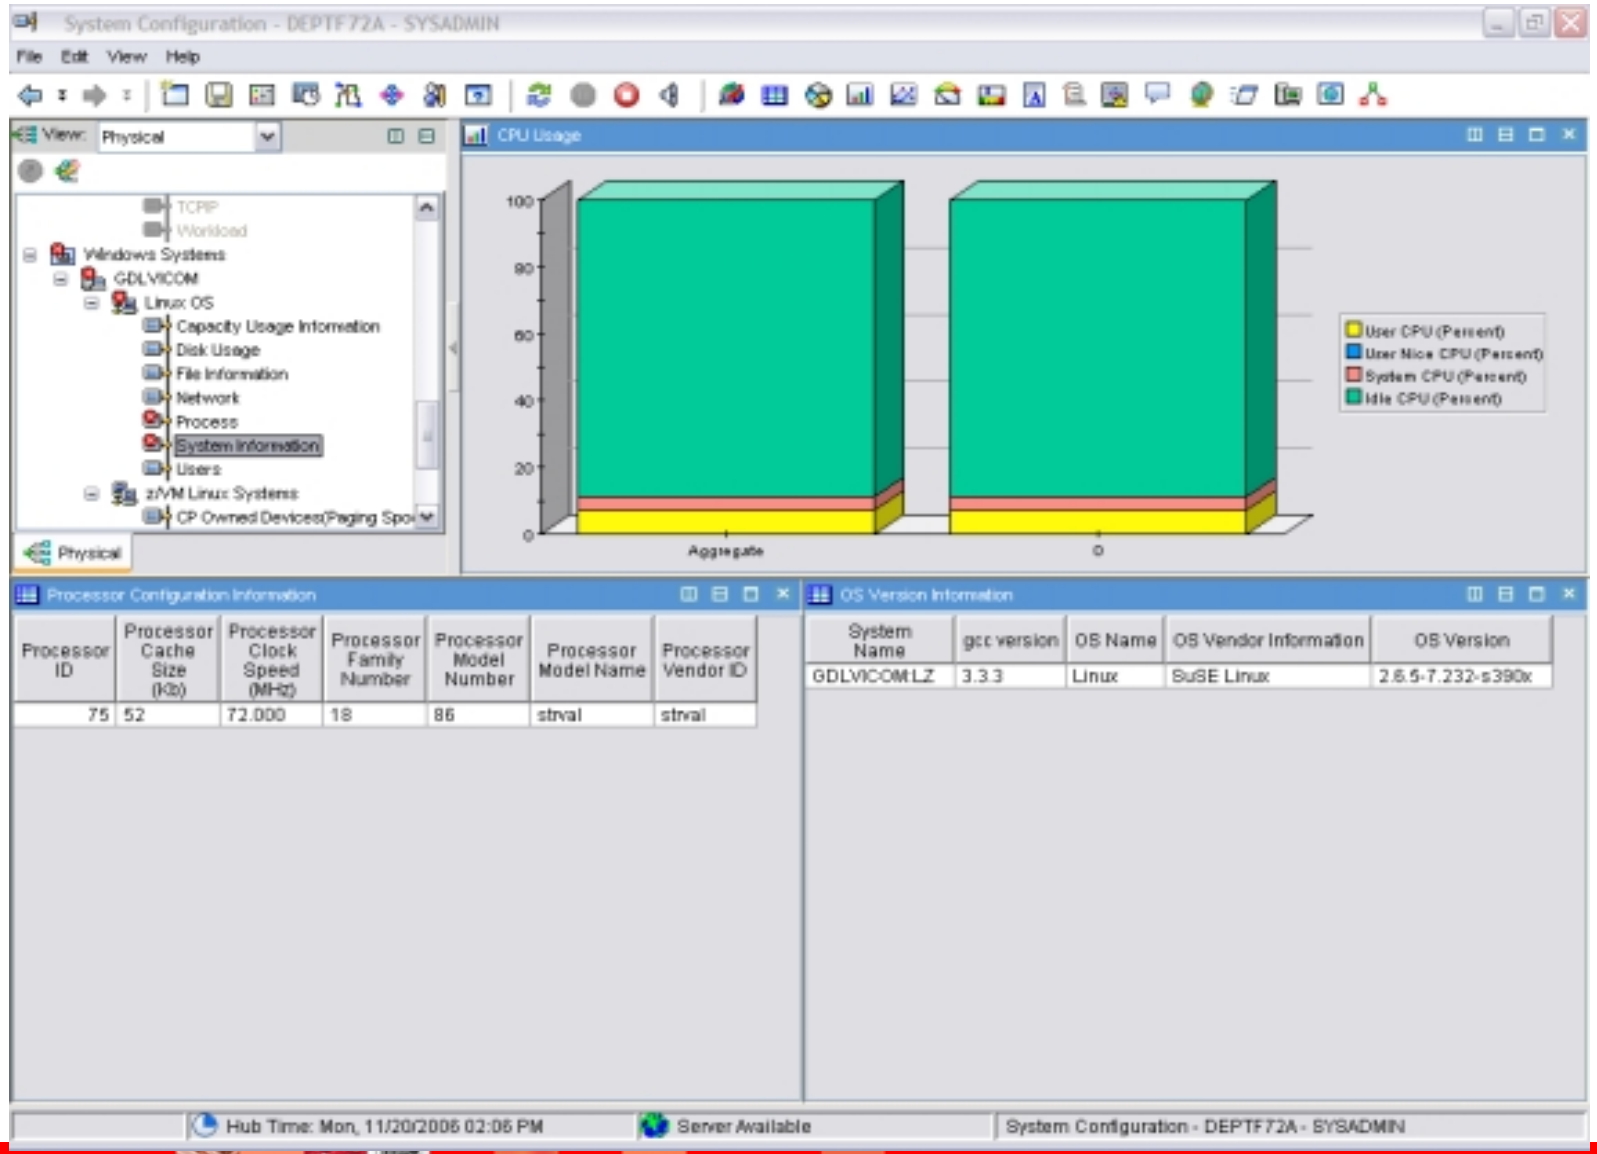

System Configuration Workspace

- **System Configuration Workspace**
  - This workspace displays information about CPU usage, the processor's configuration, and operating system level
- **Description**
  - The identification number of the processor.
  - The size of the processor cache in kilobytes.
  - The speed of the processor clock in megahertz.
  - The family number of the processor.
  - The model number of the processor.
  - The model name of the processor.
  - The identification of the processor's vendor or manufacturer.
  - The name of the host system.
  - The version of the GNU Compiler Collection (GCC) used to compile the kernel.
  - The name of the operating system.
  - The name of the operating system's vendor or manufacturer.
  - The version of the operating system.

Users