

How to Perform Problem Determination and Analysis in a Virtualized Enterprise - An Example with Discussion

Bob Neill and Ernie Gilman
IBM

August 8th, 2011
Session Number 10119

Agenda



- Challenges of Virtualization Monitoring
- Discovery and Business Service Management
- End-to-end monitoring/tracking of transactions
- Drill down into applications, virtual machines, hypervisors and OSES (z/OS, z/VM and VMware)
- How to perform problem determination and analysis for virtualized environments
- Differences in problem determination and analysis across hypervisors.
- z/OS and z/VM Monitoring and Tuning:
 - Overall System Health. z/VM has a System Health workspace, analogous to what z/OS gives with the z/OS Management Console. Points to look at (on z/VM
 - LPAR usage Allows you to look at the utilization of all your LPARS
 - Processor usage Within an LPAR, are the processors being utilized effectively
 - Sizing a Linux Guest. What is the right size?

Monitoring Challenges

- **Guest OS goes offline...should you alert on the offline system?**
 - Tie into your change control process or self-service catalog to know when a guest OS should be offline
- **How do you view Performance, Availability & Capacity Planning Reports when resources are constantly changing?**
- **Too many snapshots using up disk space**
 - Set limits on the number of snapshots and monitor
- **Machine has been offline for too long**
 - Setup process to delete the image with approval from the owner
- **Guest OS is extremely low utilization for a long time**
 - Setup process to contact the owner and find out whether resource is still in use
- **Monitor VMs to see whether they are using all of their allocated resources**
- **Resources and monitoring topologies are constantly changing due to new VMs, VMotion/Partition Mobility**
- **Capacity Planning and what-if scenarios**

Anticipating Cloud/Virtualization Management Challenges

The Virtualization Space is Growing Fast

Consolidation not only bring benefits, but new challenges.

- **Virtualized environments are heterogeneous**
 - IBM Power Systems
 - System z
 - Vmware
 - Microsoft Virtual Server/Hyper-V
 - Solaris
 - Citrix
 - XenServer
 - KVM
- **Requires key performance management insights into virtualized environments**
 - Overall resource utilization of servers?
 - Resources allocated per VM?
 - Resource utilization per VM, and how can I optimize it?
- **Cloud introduces additional complexity**
 - Dynamic provisioning
 - Dynamic de-provisioning
 - Capacity on-demand



Service Availability and Performance Management



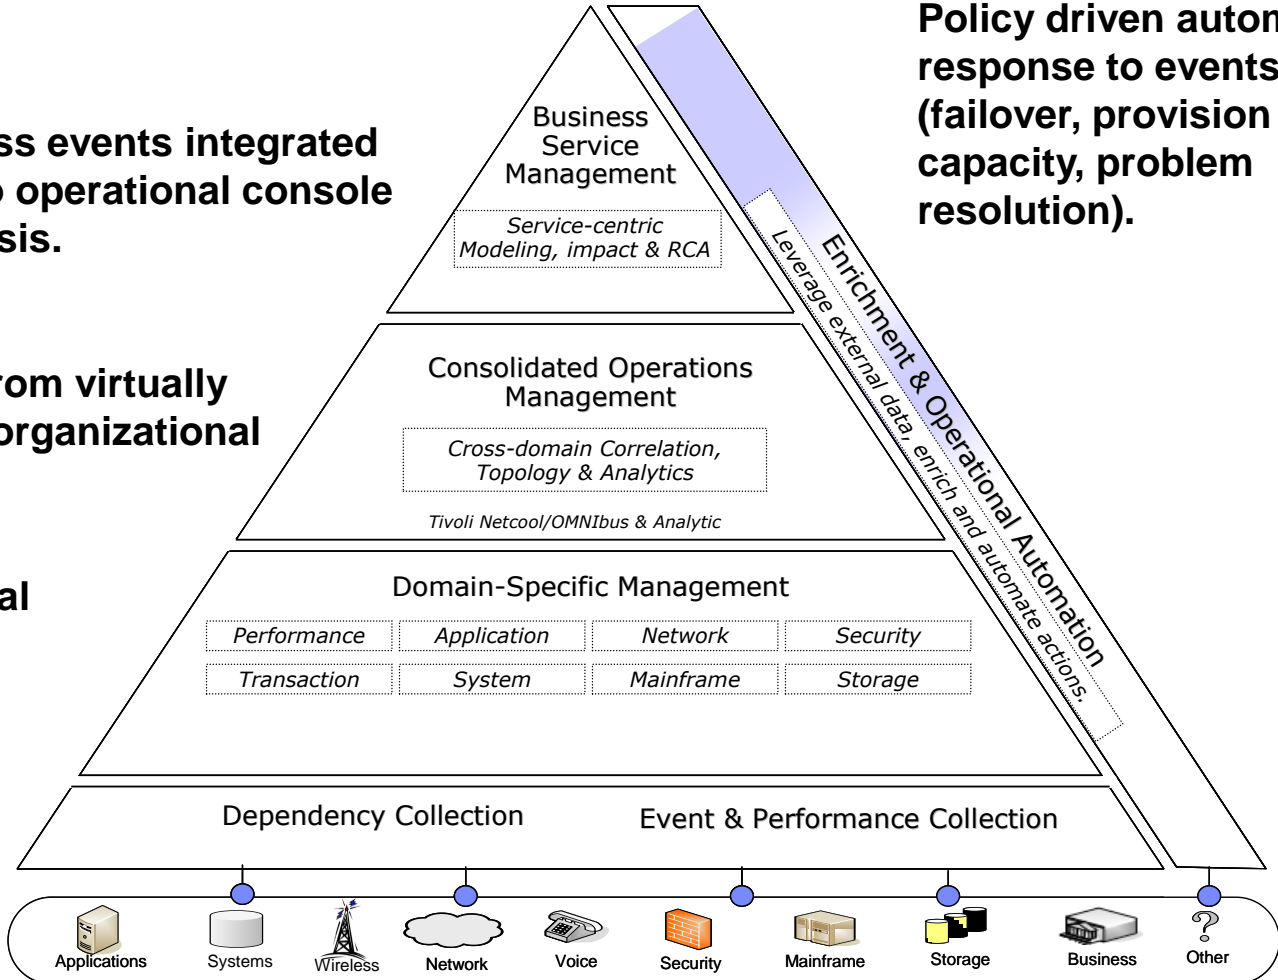
IT Operational and business events integrated in one console. Launch to operational console for quick root cause analysis.

Real-time data access, from virtually any data source across organizational boundaries

Consolidated operational view of performance & availability

Complete coverage of over 1000 device types

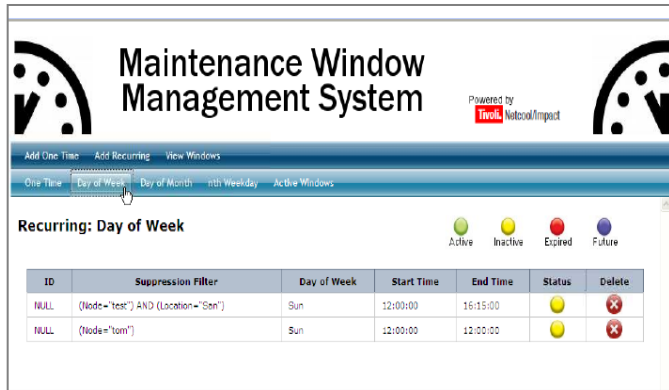
Policy driven automated response to events (failover, provision new capacity, problem resolution).



Note: All layers are inclusive of distributed and mainframe.



Tivoli Netcool/Impact v5.1.1 Highlights



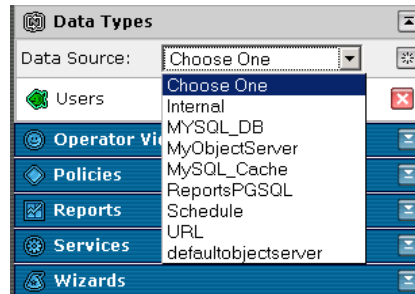
Maintenance Window Management System
Powered by Tivoli Netcool/Impact

Recurring: Day of Week

ID	Suppression Filter	Day of Week	Start Time	End Time	Status	Delete
NULL	(Node="test") AND (Location="San")	Sun	12:00:00	16:15:00	Active	✖
NULL	(Node="tom")	Sun	12:00:00	12:00:00	Inactive	✖

- Use for Maintenance Windows
- Enrich data with Business App.
- Enrich with contact information
- Enrich with SLA requirements
 - Priority 1, 2, 3
- With enriched data, easier to build reports

```
01 log("SimpleEventEnrichment received event with Node " + @Node);
02 //Look for the matching node in customer's database
03 OrgNodes = GetByFilter("oracle_swordfish_node", "NAME=" + EventContainer.Node + "", false);
04 Num = length(OrgNodes);
05 log("SimpleEventEnrichment GetByFilter successful. Found " + Num + " dataItem(s).");
06 if (Num > 0) {
07 //Update Event with information found in customer's database
08 EventContainer.Location = Nodes[0].LOCATION;
09 EventContainer.Summary = "Node Administrator " + OrgNodes[0].Admin + ". Email " + OrgNodes[0].E
10 //Return Event to
11 ReturnEvent(
12 log("ReturnEvent successful");
13 }
14
```



Data Types

Data Source: Choose One


Users	Choose One	✖
Operator Vi	Internal	▼
Policies	MYSQL_DB	▼
Reports	MyObjectServer	▼
Services	MySQL_Cache	▼
Wizards	ReportsPGSQL	▼
	Schedule	▼
	URL	▼
	defaultobjectserver	▼

IBM Tivoli Application Dependency Discovery Manager (TADDM)



IBM Tivoli Application Dependency Discovery Manager initiates and assists planning for consolidation by providing best-of-breed discovery capabilities

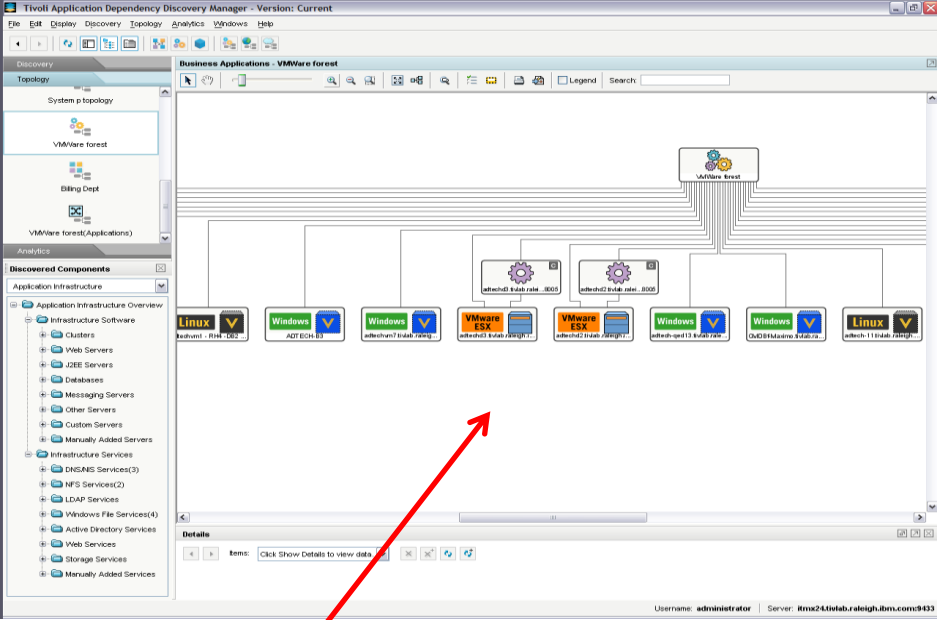
- Discovers the COMPONENTS in a Data Center Environment
- CENTRALIZES and VISUALIZES the CONFIGURATION of the Components in a Data Center Environment
- Discovers the RELATIONSHIP of the Components in a Data Center Environment
- DISCOVERS AND TRACKS THE CHANGES in a Data Center Environment



Feeds this Information to other IBM Tivoli Products

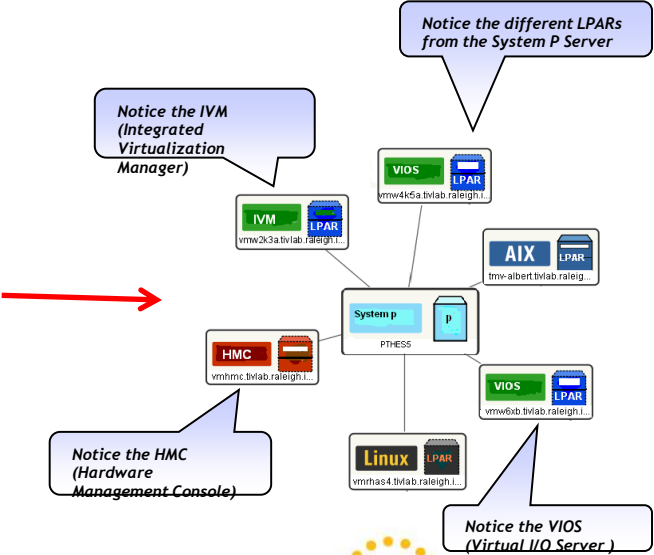
Now that I have virtualized, things seem more complex.

- Discover physical & virtual environments & their dependencies. Customer are using this for problem determination, change window planning, placement of new workloads.
- Track and report on configuration changes for quick problem isolation
- Compare configurations across like resources or against the “gold standard”
- Discover both the hypervisor and the virtual machines.



VMware

Power Systems






Tivoli Application Dependency Discovery Manager reduces incident isolation time by 50% & reduces application roll backs by 10-30%.



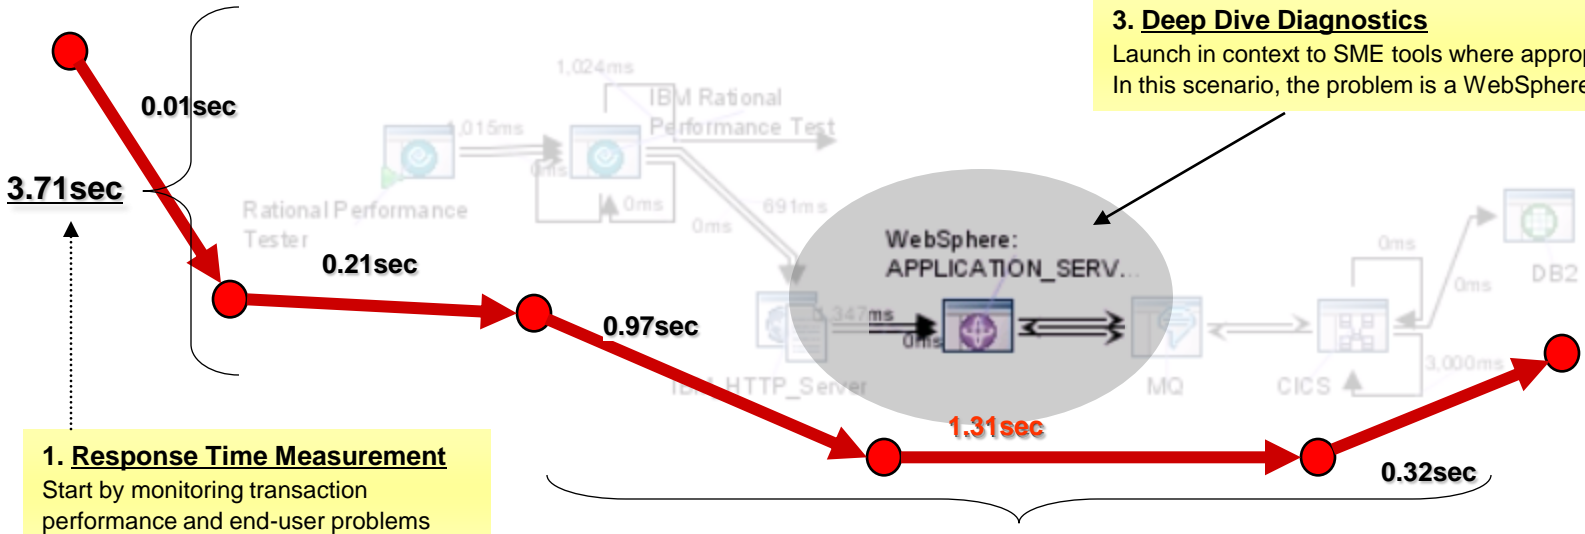
Composite Application Management and Resource Monitoring



- Monitor application response to ensure business expectations are met
- Understand transaction flows over complex topologies
- Monitor infrastructure performance and availability
- Diagnose application performance issues
- Increase application availability and customer satisfaction
- Improve MTTR and MTBF



End-to-End Monitoring, Tracking and Diagnosis



1. Response Time Measurement
Start by monitoring transaction performance and end-user problems

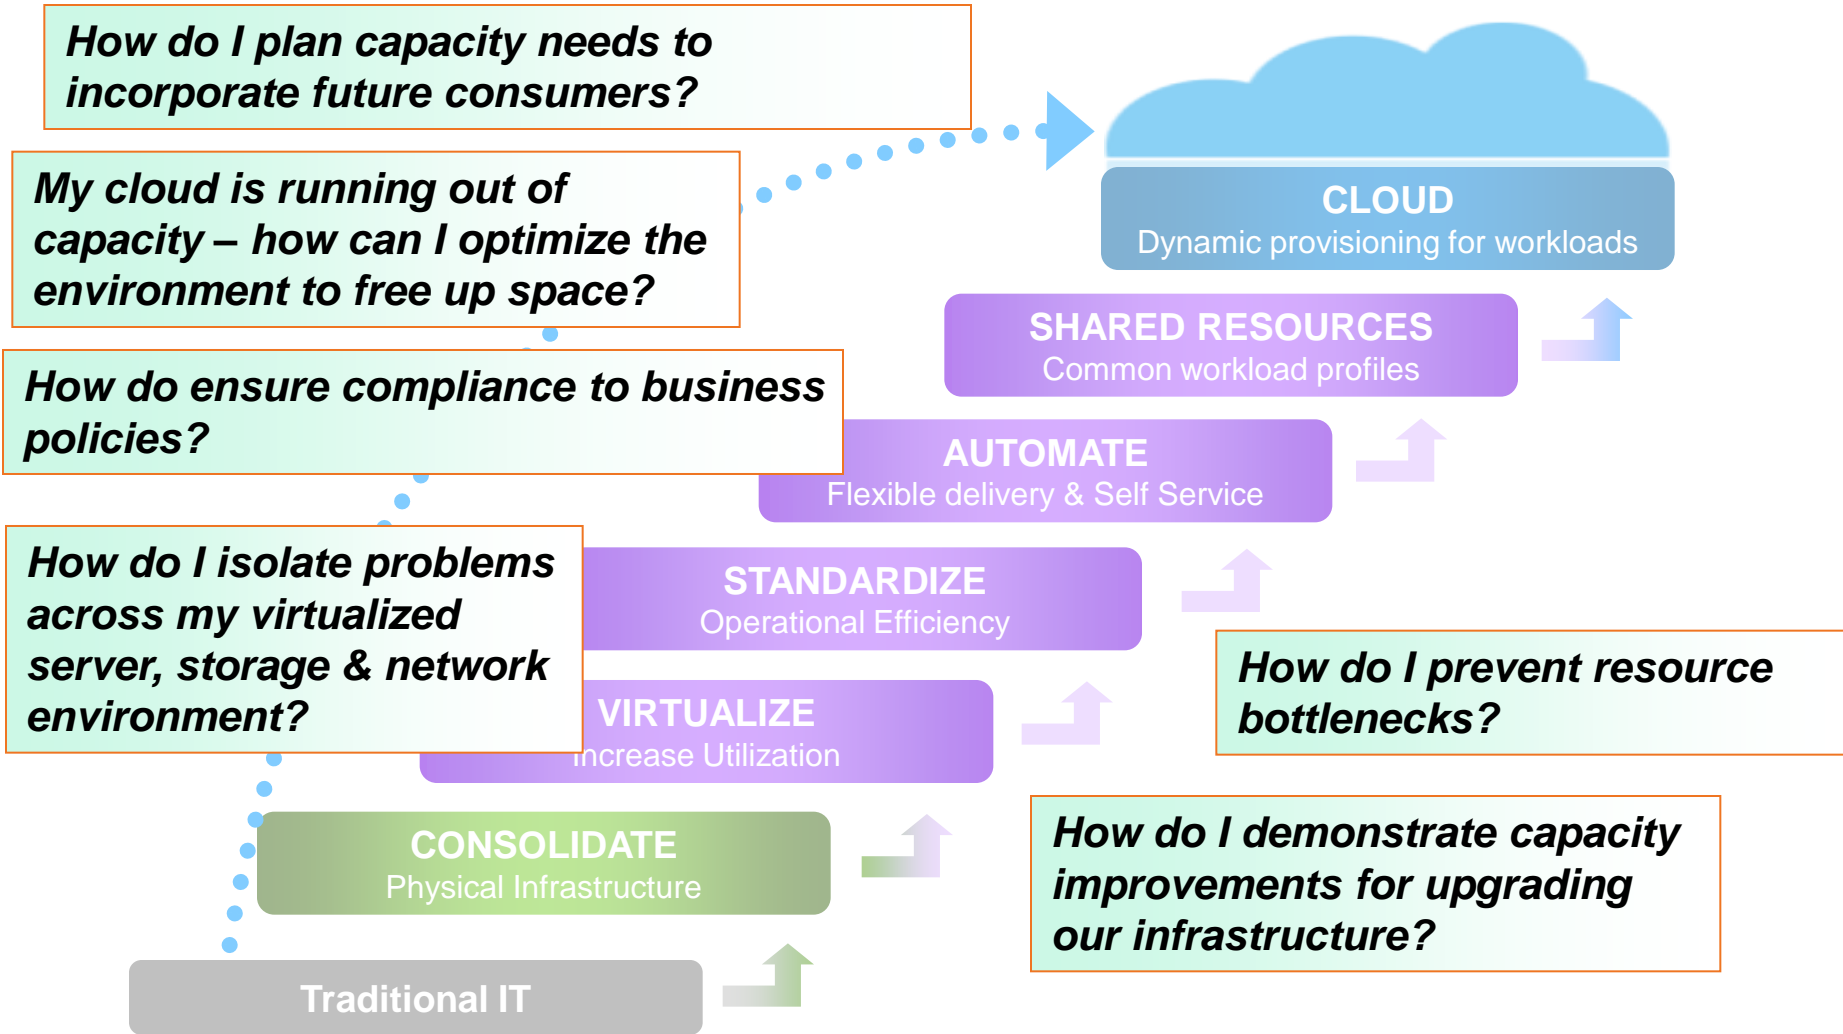
2. Transaction Tracking
Correlate data from app server, MQ, CICS, IMS, custom instrumentation, etc. to show topology and isolate problems

3. Deep Dive Diagnostics
Launch in context to SME tools where appropriate. In this scenario, the problem is a WebSphere JEE memory leak.

Transaction Root Cause Analysis

<p>1. Sense End User Experience and alert on threshold violation</p>	<p>2. Isolate by measuring performance data against baseline through entire infrastructure</p>	<p>3. Diagnose and repair through launch-in-context into deep-dive diagnostics</p>
---	---	---

Challenges with Managing a Virtual Environment



Virtualization Monitoring

- Hypervisors such as z/VM, VMware and Hyper-V only provide visibility into what they manage
- ITM provides in ONE tool:
 - Dynamic thresholding
 - Capacity planning
 - Reporting for performance trends
 - Monitoring across servers, storage – phys & virt view
 - Predictive trending
 - Integrated monitoring for Virtualization, Applications, Middleware, and more.
 - Response Time is the most important metric for a business application whether it's running on physical servers or VMs
 - End-to-End Transaction Monitoring
 - Storage
 - Network Monitoring

Key Metrics – Physical & Virtual

- **Server Problems**
 - **VM, Host, Cluster CPU, Memory Utilization**
 - **VMware CPU % ready**
 - **VMware Memory over commit (Active Memory/Physical Memory)**
 - **VMware Swapping/ballooning**
 - **VM Swapping**
 - **Absence of processes, growth of process resource consumption**
 - **Host server failures**
 - **Unallocated resources**
 - **Storage and Networking problems (see next 2 slides)**
 - **Pool utilization on Power Systems**

Physical & Virtual Storage & Network Problems



Storage Problems

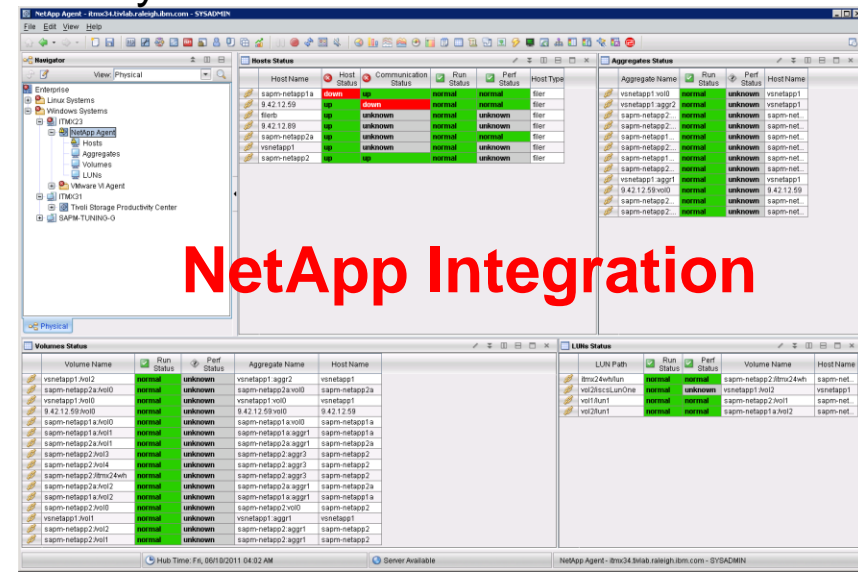
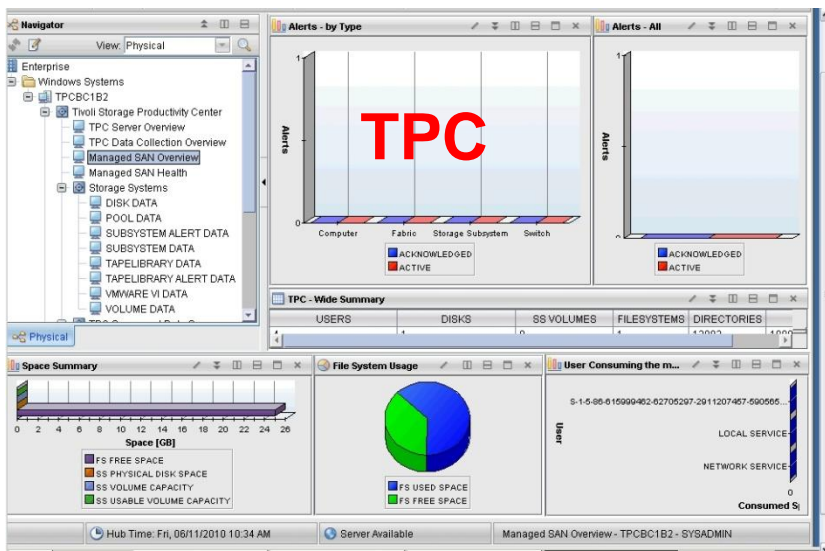
- Data store issues
 - May be caused by insufficient space - # of VMs, storage usage growth, utilization of datastore, changes in configuration
- Storage Latency/response time issues
 - May be caused by too many VMs associated to same LUN (Volume), HBA bottlenecks, backend storage issues (disk, etc.)

Network Problems

- Network response time problem
 - May be caused by too many VMs sharing NICs without necessary throughput capacity, changes in configuration
- Network connection problem
 - May be caused by unavailability of connection to physical network (physical switch / port)

Storage Considerations


- Storage Monitoring is critical.
 - VCenter only provides a portion of the picture. It does not show you full utilization of the storage volumes and storage subsystems
 - VCenter does not provide information about whether space is available to allocation more resources.
- Adaptive Monitoring for Disk I/O and Latency which can be very difficult to define a threshold.
- Monitor Datastores for Disk Latency including queue, device, and kernel latency
- For NetApp Storage, use the NetApp Agent
- For other storage, use TPC (Tivoli Storage Productivity Center) or other monitoring tools to ensure the storage is healthy and has low latency.



Monitor Storage Efficiency Utilization with Tivoli Storage Productivity Center





- View capacity utilization by computer, virtual machine or storage system
- File System and database storage utilization details
- Identify wasted space on volumes based on age, file type, or any other user defined filter



Performance Analyzer for:

- Vmware and Power Systems (CPU Trends, Disk Utilization, Memory, Network) out of the box
- Recommend setting up Analytic Task for:
 - VMware Cluster s
 - VMware VM CPU Percent Ready
 - Other hypervisors such as z/VM, Hyper-V, etc.



Forecast Status			
System Name	Confidence	Strength	Number Of Samples
TestWinXP-7	48	1	89
TestWinXP-8	83	3	89
TestWinXP-9	87	3	89
TestWinXP-10	90	3	89
TestWinXP-4	100	3	89
TestWinXP-5	89	3	89
TestWinXP-6	86	3	89
TestWin2003-2	82	3	89
TestWin2003-3	89	3	89
TestWin2003-4	73	3	89

7 Day Forecast	
System Name	Data
TestWinXP-7	829
TestWinXP-8	3400
TestWinXP-9	3642
TestWinXP-10	4370
TestWinXP-4	2318
TestWinXP-5	2206
TestWinXP-6	925
TestWin2003-2	5094
TestWin2003-3	3430
TestWin2003-4	2519

30 Day Forecast	
System Name	Data
TestWinXP-7	987
TestWinXP-8	4231
TestWinXP-9	4484
TestWinXP-10	5395
TestWinXP-4	2870
TestWinXP-5	2718
TestWinXP-6	1151
TestWin2003-2	6185
TestWin2003-3	4229
TestWin2003-4	3135


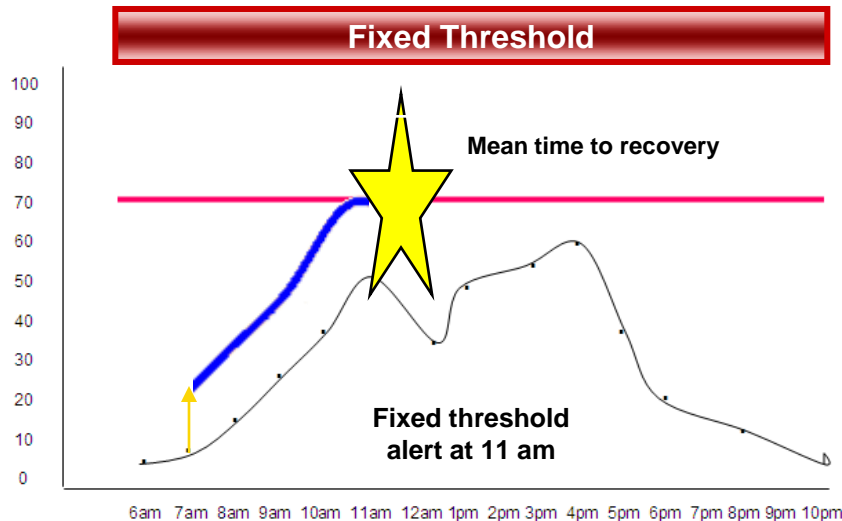
90 Day Forecast	
System Name	Data
TestWinXP-7	1398
TestWinXP-8	6397
TestWinXP-9	6682
TestWinXP-10	8068
TestWinXP-4	4310
TestWinXP-5	4054
TestWinXP-6	1741
TestWin2003-2	9032
TestWin2003-3	6311
TestWin2003-4	4740

Incident Avoidance - Dynamic Thresholds

Dynamic Thresholds can calculate baseline values using one of several statistical functions based on historical data from the Tivoli Data Warehouse and/or agents.

This allows tracking deviations **from the norm** as predictors of future problems.

- No automated approach to define
- No warning of abnormal behaviors prior to peak periods
- No flexibility in the monitoring environment
- Automated definitions with + or – variations
- **Proactive** warning for abnormal behavior occurring before peak periods or during non-peak periods



Historical Collection Recommendations

- **Historical Collection Best Practices paper:**
<https://www.ibm.com/developerworks/wikis/display/tivolimonitoring/Historical+Collections+Best+Practices+in+Tivoli+Monitoring+6.2.2>
- **TCR Reports for:**
 - Vmware
 - Power Systems
 - Hyper-V
 - KVM
- **Complete Warehouse Planning Spreadsheet:**
 - <http://www-01.ibm.com/software/brandcatalog/ismlibrary/details?catalog.label=1TW10TM1Y>
- **6.2.2 Offers Granular Warehousing (configuration per Agent)**
 - Filter out CD/DVD and Floppy Drive Data
 - For Data Stores are shared across multiple ESX servers, consider filtering
- Performance Analyzer uses Summarized Data
- Adaptive Monitoring/Dynamic Thresholding uses detailed data
- Out of the box TCR reports are written for Summarized data
- **For Cloud Environments, build reports that show change in addition to performance data**

Leverage TCR Reports



Tivoli Integrated Portal - Mozilla Firefox: IBM Edition

File Edit View History Bookmarks Tools Help

ibm.com https://itmx21a.tivlab.raleigh.ibm.com:16311/ibm/console/login.do?action=secure

Tivoli Integrated Portal

Tivoli View: All tasks Welcome cesar Help | Communities | Logout IBM

VMware Cluster Dashboard | Server Details | **Common Reporting**

Work with reports

Viewer - VMware VI Balanced and Unbalanced Clusters

cesar About

Keep this version Add this report

Tivoli software

VMware VI Balanced and Unbalanced Clusters

Date filter: Last 7 days
Start Date: May 27, 2011
Data Center: RTP_SAPM
Cluster(s): ALL
Shift Periods: All Shifts
Vacation Periods: All days

End Date: Jun 3, 2011 11:59:59 PM

CPU Utilization

Cluster Name	Avg CPU Utilization
BladeCenter_Cluster_32bit	~45
BladeCenter_Cluster_64bit	~20
Development_Cluster	~70
Test_Cluster	~20
XSeries_Cluster	~5

Memory Utilization

Cluster Name	Avg Memory Utilization
BladeCenter_Cluster_32bit	~50
BladeCenter_Cluster_64bit	~75
Development_Cluster	~45
Test_Cluster	~45
XSeries_Cluster	~15

Statistical Summary:
Statistical Min: 35%
Mean: 45%
Statistical Max: 60%

Some clusters have very high utilization, some have very low – need to balance clusters

Datstore Space Utilization

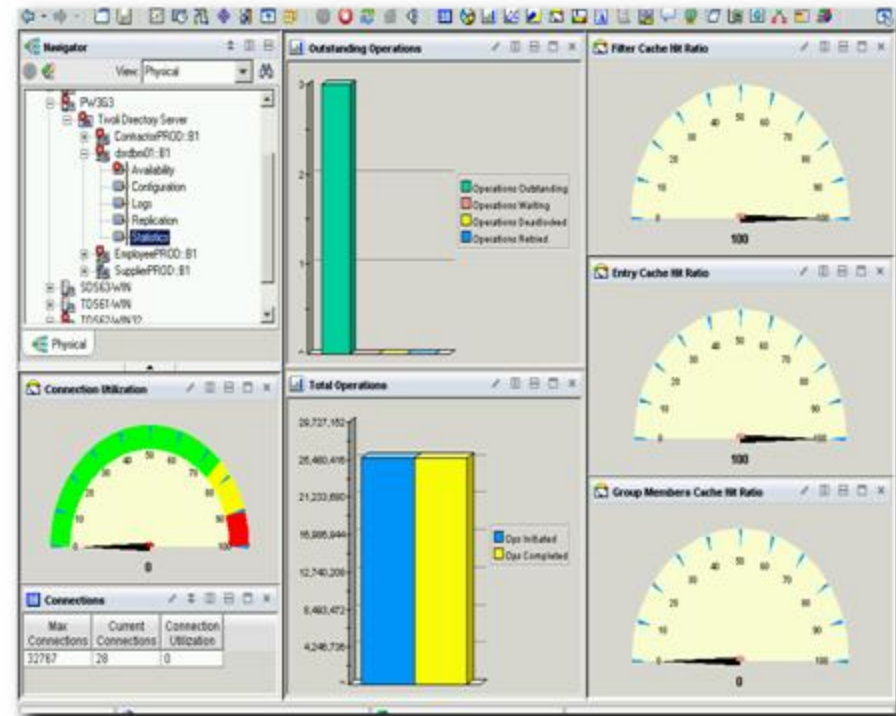
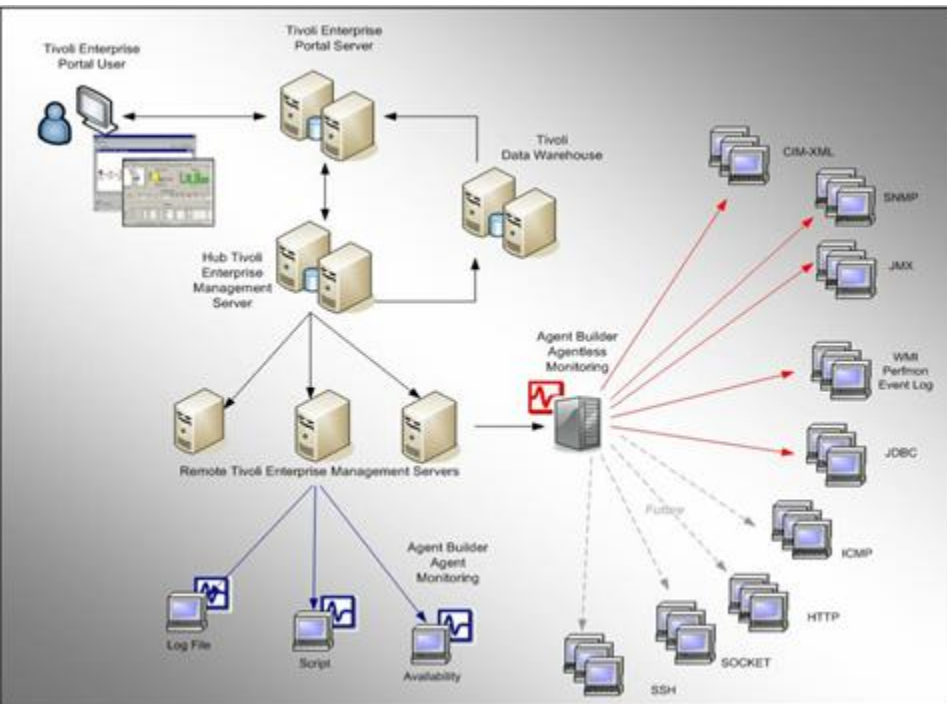
Cluster Name	Avg Datstore Utilization
BladeCenter_Cluster_32bit	~95
BladeCenter_Cluster_64bit	~95
Development_Cluster	~98
Test_Cluster	~95
XSeries_Cluster	~95

Avg Datstore Utilization: ~95%
Mean: ~95%

Leverage Agent Builder for Custom Monitoring solutions

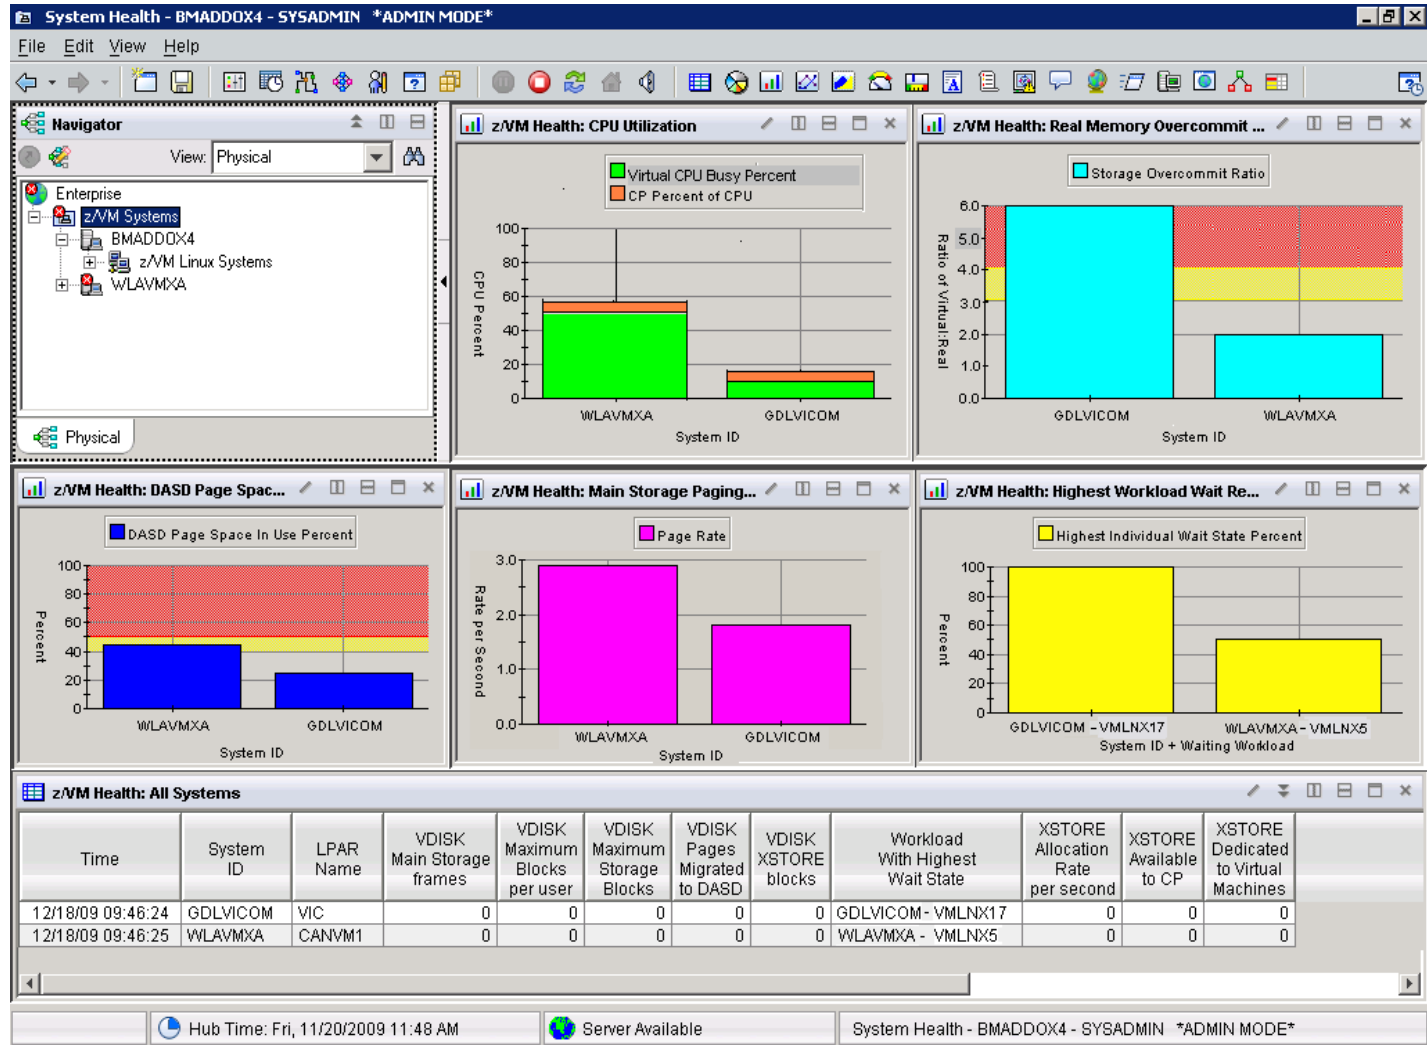
- Wide range of protocols supported for Custom Agents

- Extend monitoring of home-grown / custom applications quickly & easily



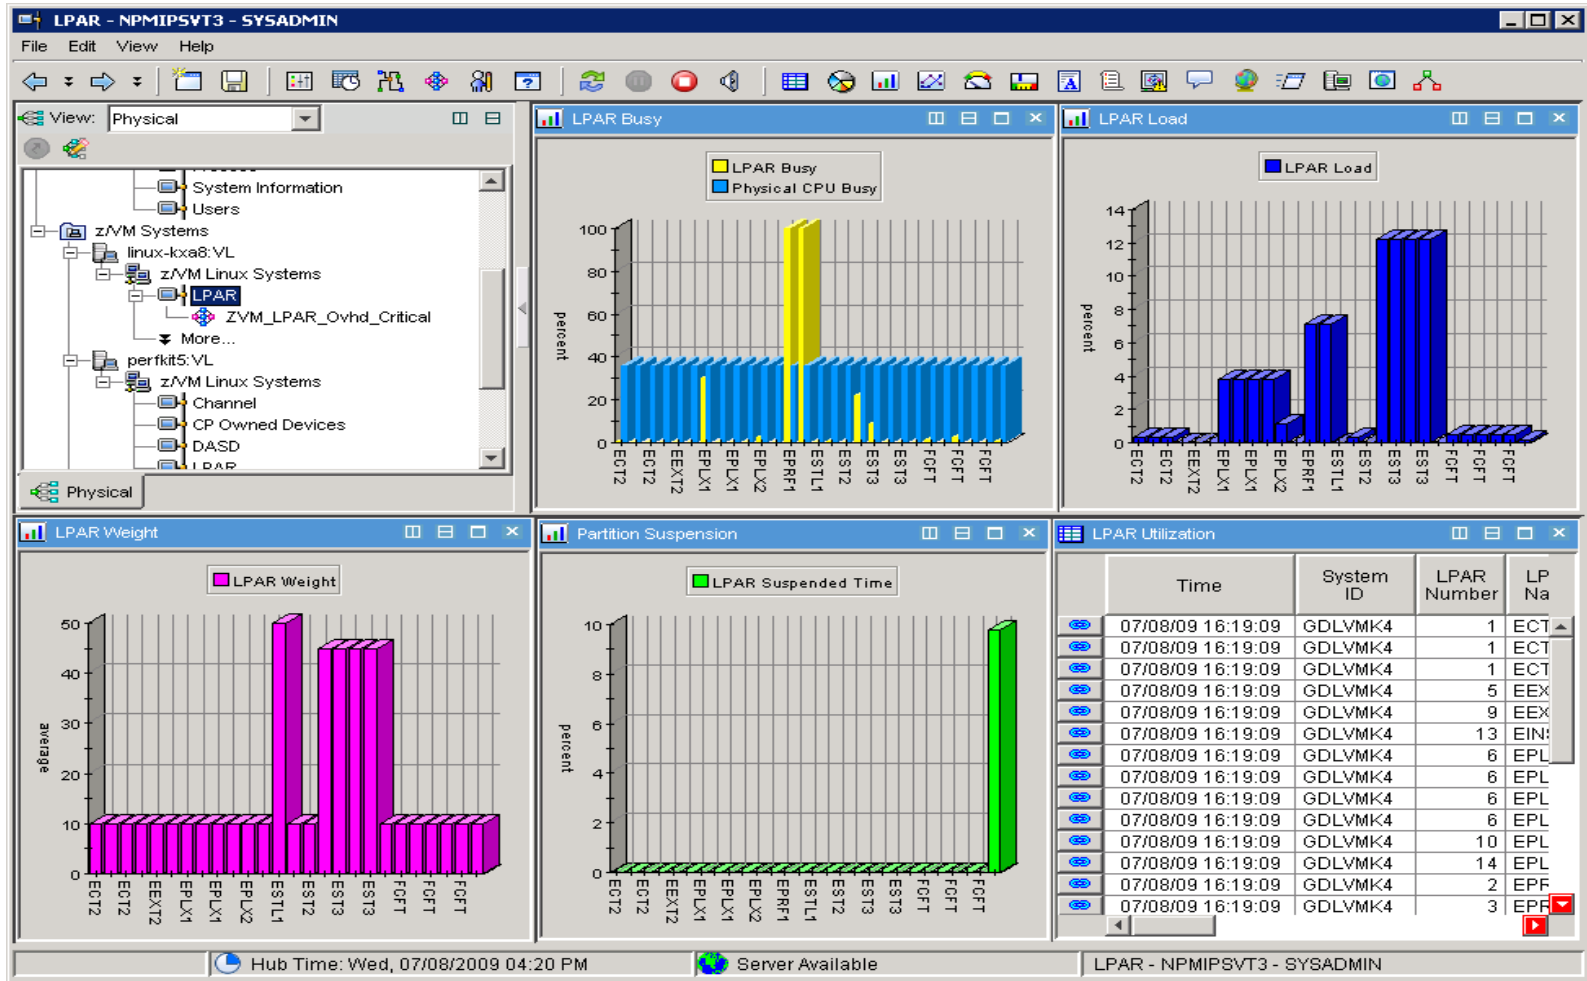
z/VM Overall System Health

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

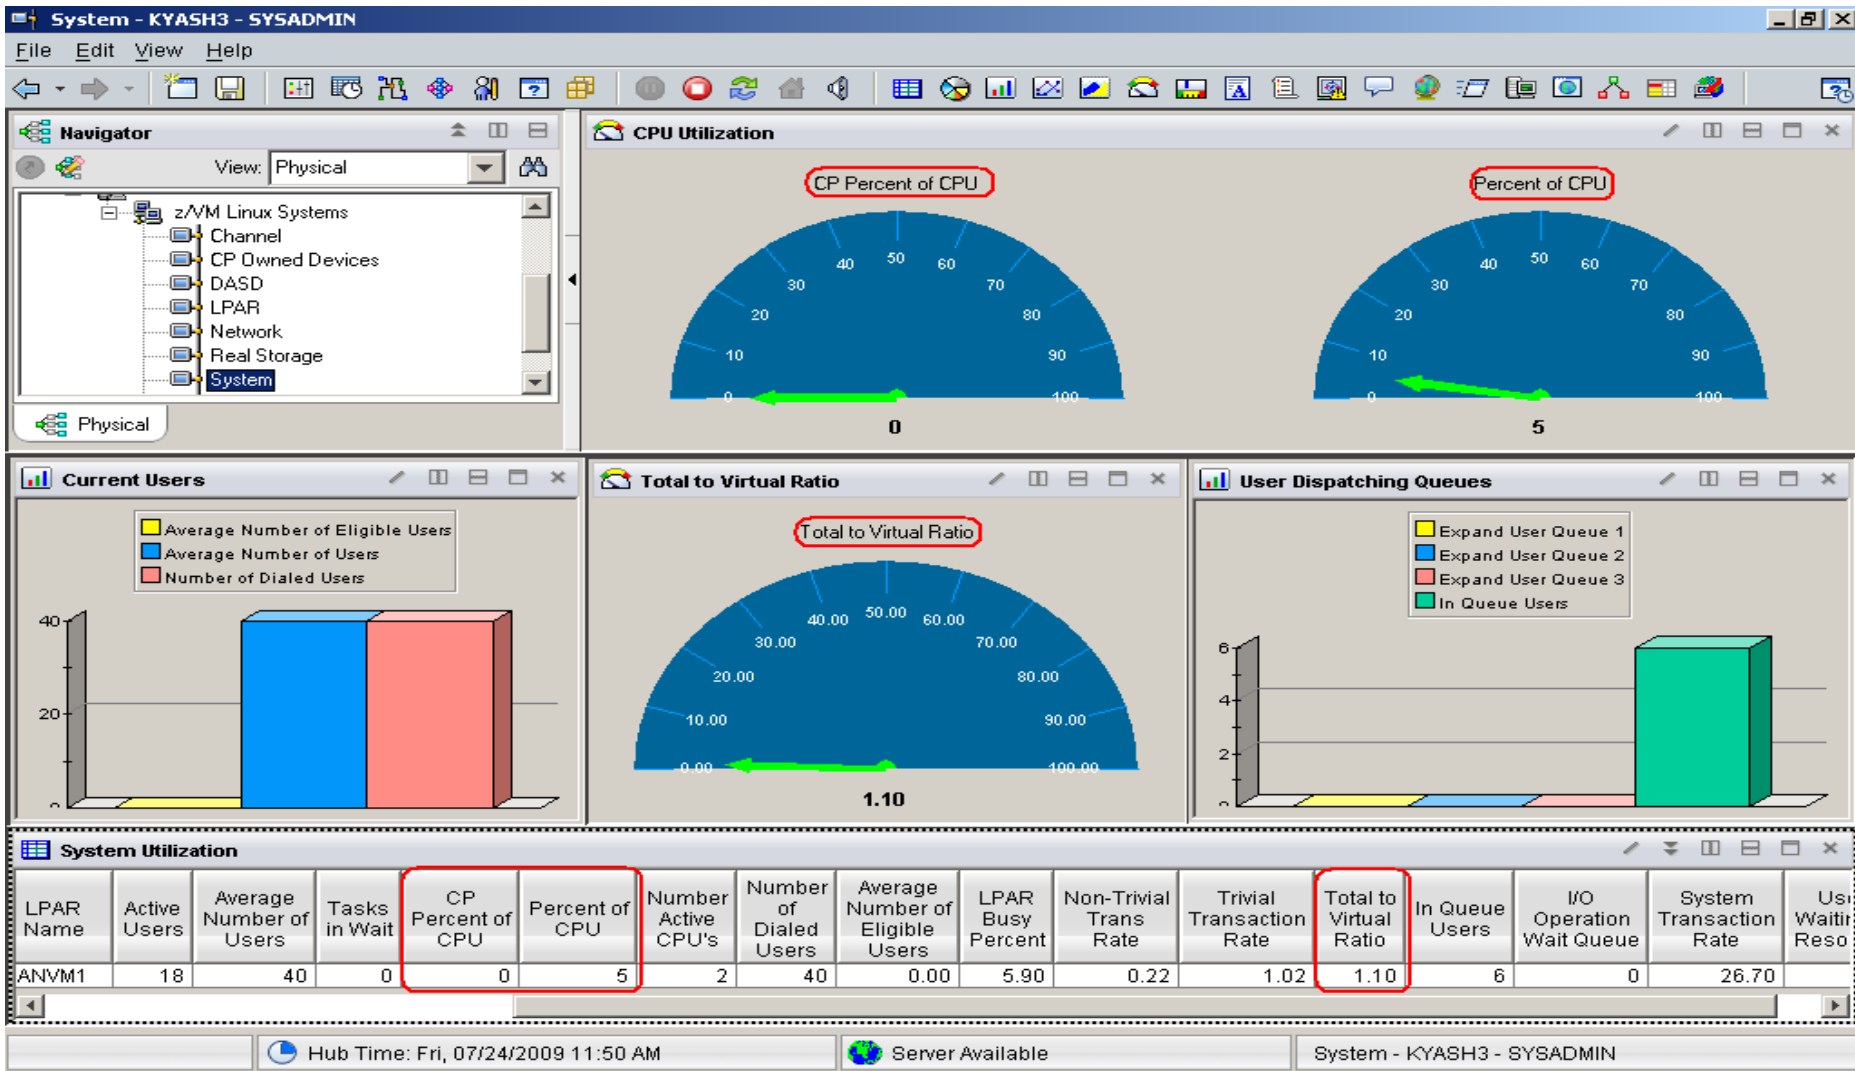


LPAR Usage

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

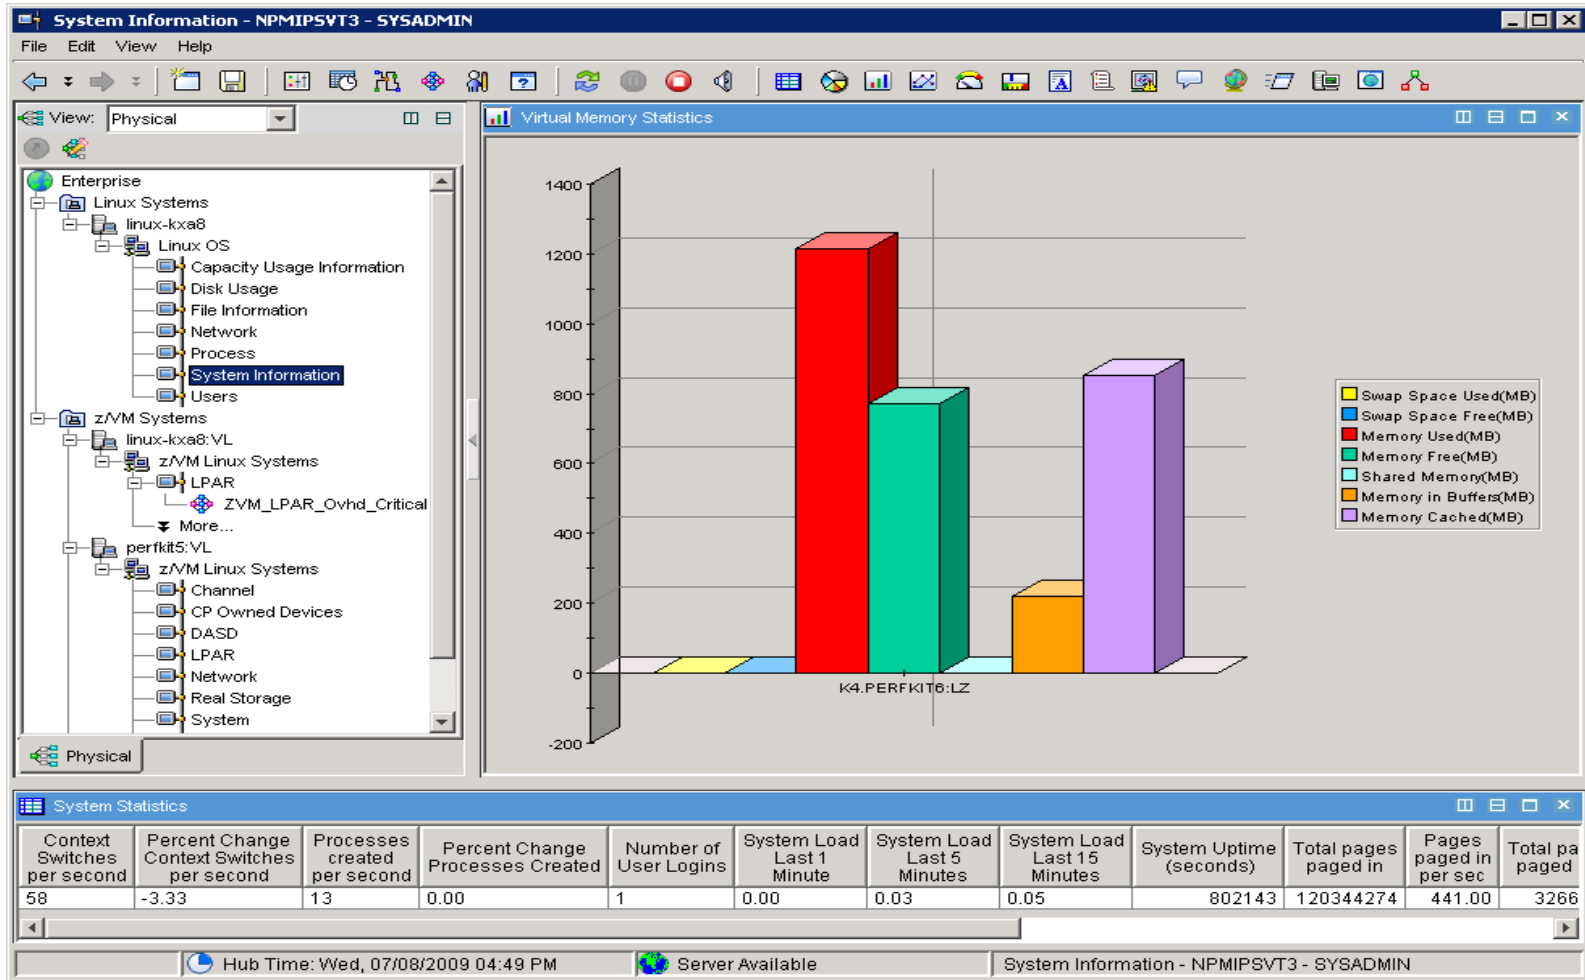


Processor Utilization



Sizing a Linux Guest

Memory usage of a particular Linux virtual machine



9916: Business Decisions for Cloud Computing

Tuesday, 11:00 AM-12:00 PM Europe 4

9476: Virtual Linux Server Disaster Recovery Planning

Wednesday, 1:30 PM-2:30 PM Oceanic 7

9917: Top 10 Tips for z/OS Network Performance Monitoring with OMEGAMON

Wednesday, 1:30 PM-2:30 PM Europe 11

9459: Cloud Computing with IBM System z

Thursday, 8:00 AM-9:00 AM Oceanic 7

9469: Managing z/VM & Linux Performance Best Practices


Thursday, 1:30 PM-2:30 PM Oceanic 8

9470: Mainframe Optimization: Making System z the Center of Enterprise Computing

Thursday, 4:30 PM-5:30 PM Oceanic 6

9308: TCP/IP Performance Management in a Virtualized Environment

Friday, 8:00 AM-9:00 AM Oceanic 7



SHARE
Technology · Connections · Results

QUESTIONS?