

# Leveraging ITM 6.x System P Agents to monitor an AIX Virtualized environment

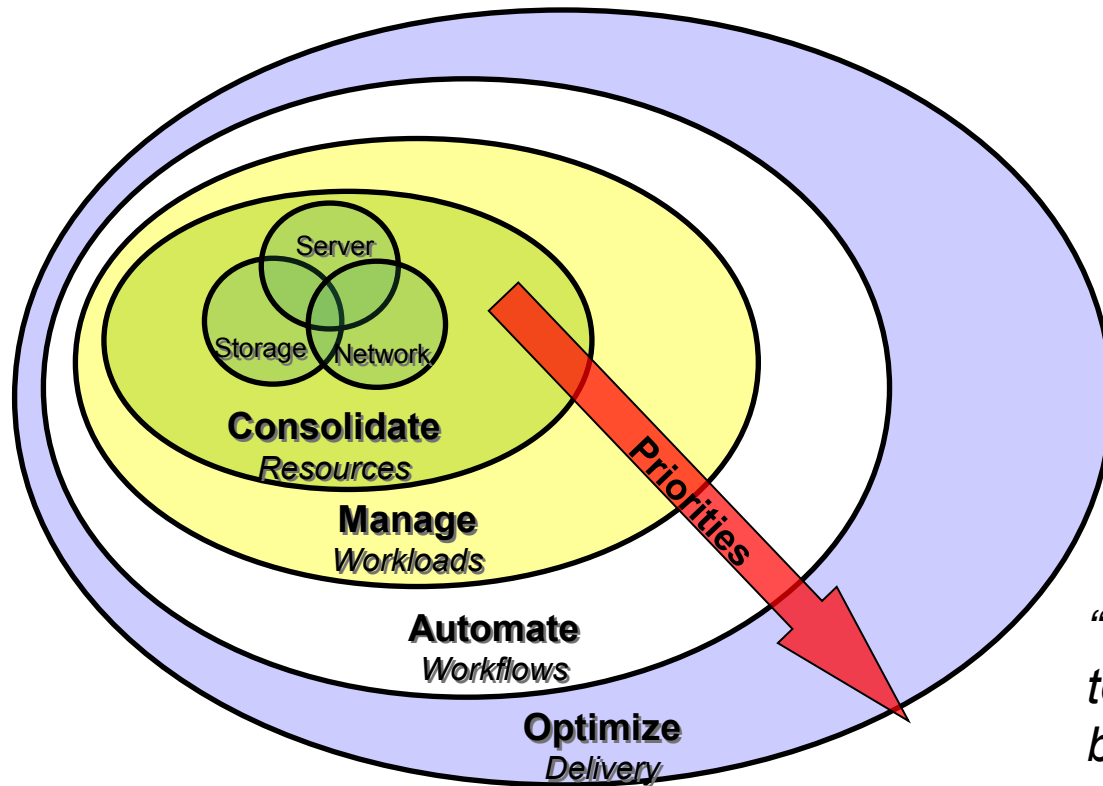
- Archana Raghavan

# AGENDA

- Infrastructure Challenges – Need for monitoring
- Overview of IBM Tivoli Monitoring
- Tivoli Monitoring for Power Systems
- Deployment Tips
- Questions

# Virtualization Journey

*Addressing challenges to achieve value*



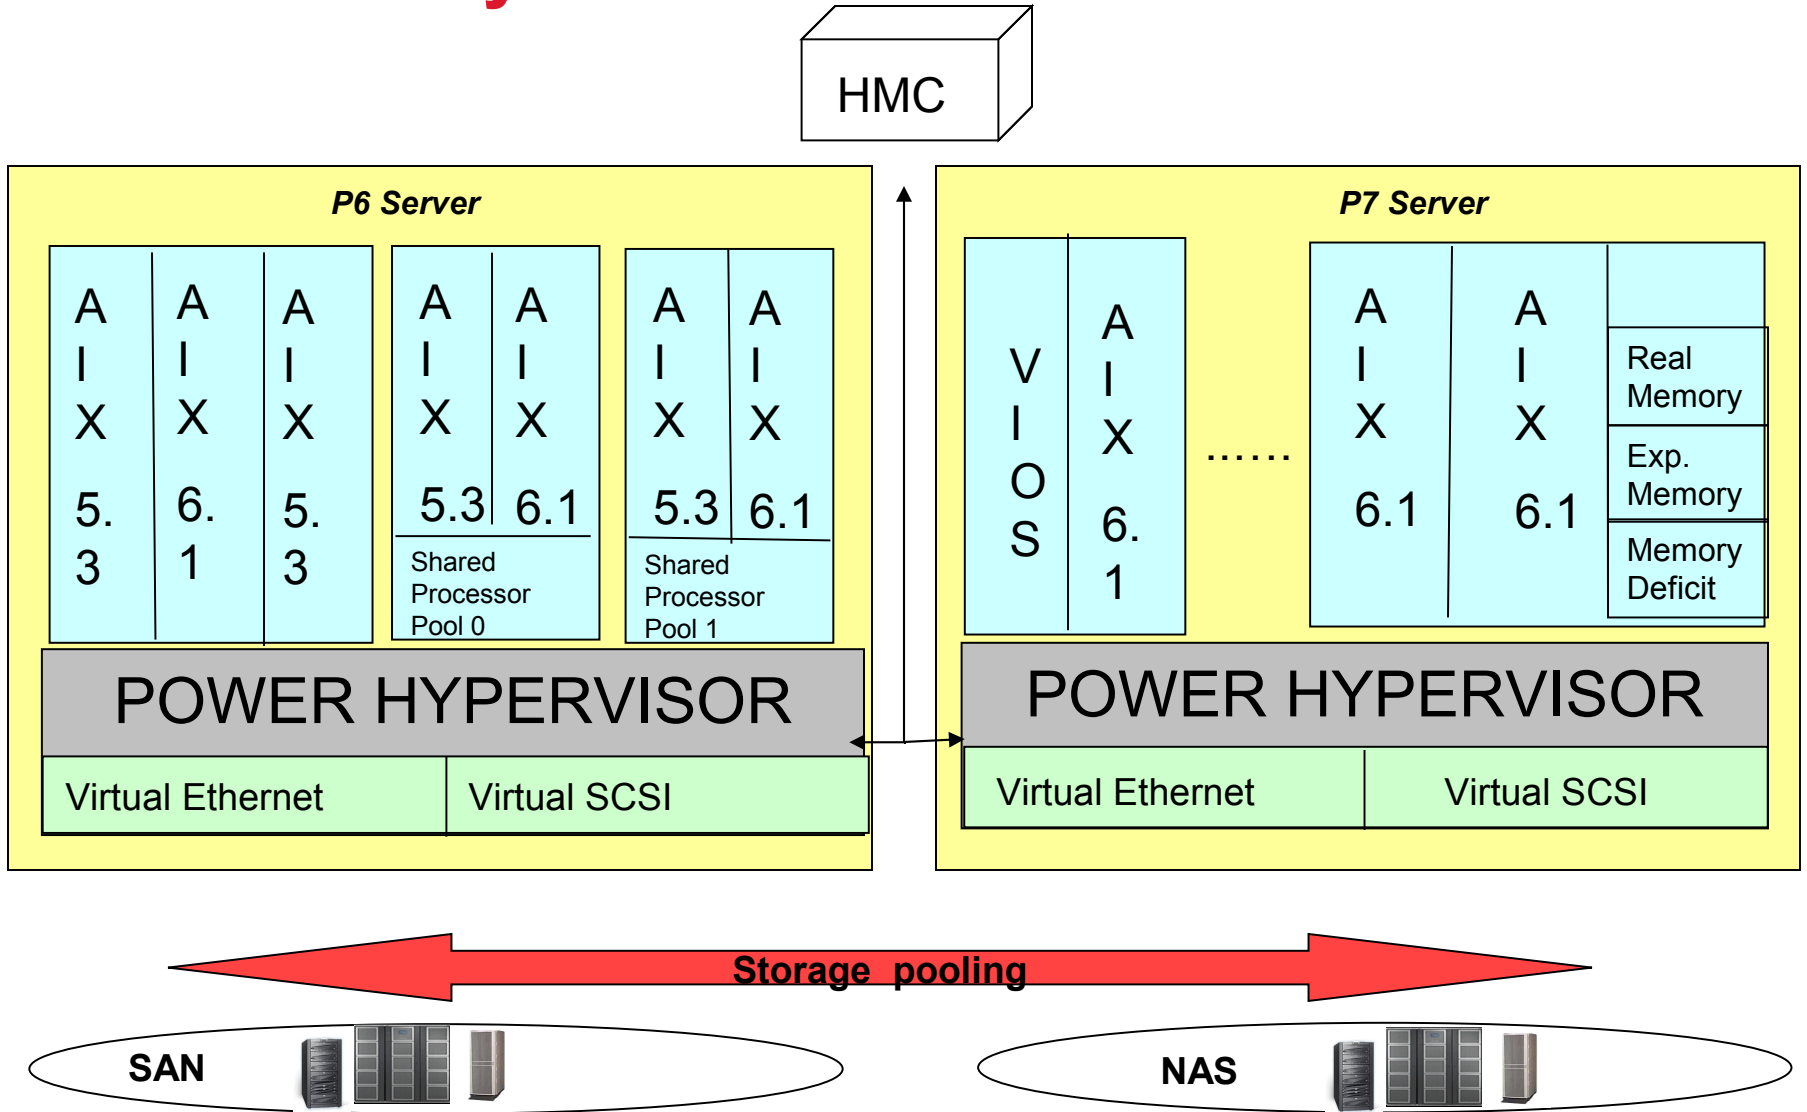
*“I need to **consolidate** to reduce sprawl & CAPEX.”*

*“I need to **manage** my workloads to rapidly respond to changes in capacity requirements and to ensure high availability of this environment.”*

*“I need to **automate** my processes to consistently & quickly respond to business requirements.”*

*“I need to **optimize** to create value & improve IT’s ability to rapidly support business innovation.”*

# AIX/Power Systems Architecture: Internals



# IBM Tivoli Monitoring - Overview

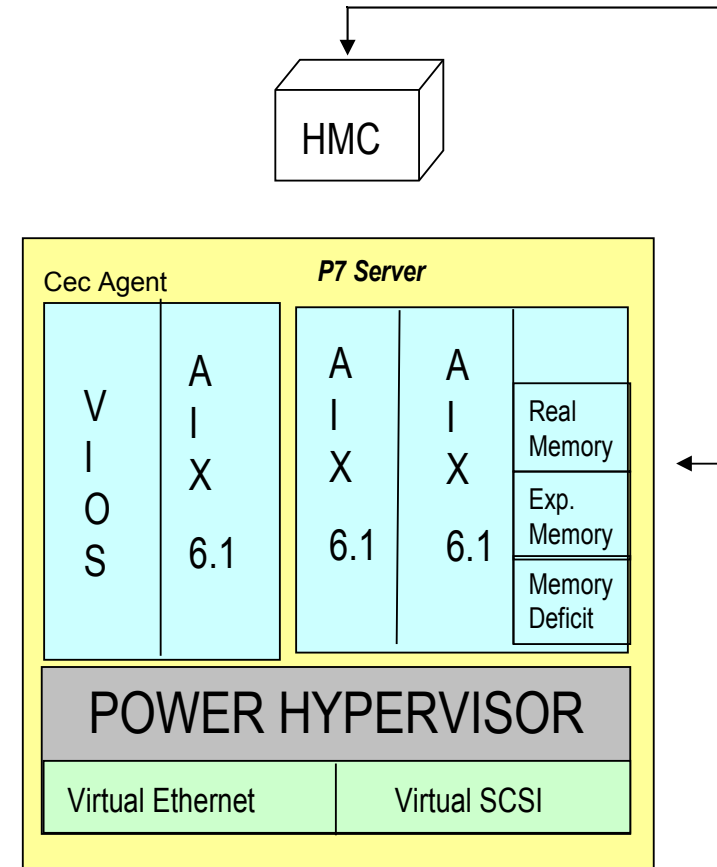
- Monitors Power Systems
  - Software / Hardware errors
  - Resource utilization (CPU, memory, storage space etc)
  - I/O performance
  - Processes
- Provides a framework to also monitor
  - AIX Operating Systems
  - Key applications/databases (SAP, DB2, Oracle)
  - Other Virtualized environments (z/VM, VMware, Hyper-V)
- Single console with easy to use displays
- Customizable workspaces
- Real-time & predictive alerts for problems identified
- Automated actions for problems
- Collection of historical data for reporting & analysis
- Tools to build customized solutions (Universal Agent & Agent Builder)

# Tivoli Monitoring for Power Systems - Overview

- Tivoli Monitoring for Power Systems includes:
  - CEC agent
  - AIX agent
  - HMC agent
  - VIOS agent
  - Log Alert agent

# CEC AGENT

- Benefit
  - Helps determine how efficiently the resources of a frame are being utilized
  - Helps determine if LPAR resources are over/under utilized against their entitlement
- Overview of the entire frame
  - Inventory of resources such as CPU, memory
  - Resource allocation to individual LPARS, Pools
  - Information regarding number of LPARS, processor, CEC modes
- Integrates with System Director
  - Contextually launch to make configuration changes, etc.
- Pre-installed on the VIOS



# Multiple Shared Processor Pools (MSPP)

CEC Resource Inventory - boson1.raleigh.ibm.com - SYSADMIN

File Edit View Help

Navigator View: Physical

- User
  - Status
  - h17p02
    - VIOS Premium
    - h17p40
      - AIX Premium
      - CEC Base
        - CEC Resources
        - CEC Utilization
      - PK
      - HMC Base
      - UNIX OS
      - npiv-ft-vios3

Physical

CEC Resource Inventory

Name	Number of Partitions	CPU Total	CPU Allocated	CPU Unallocated	CPU Allocated Pct	CPU Unallocated Pct	CPU Shar
h17	48	16.0	13.6	2.4	85	15	14.0

CEC LPAR Metrics

Page: 1 of 4

Name	ID	State	Environment	PoolID	Entitlement	CPU Allocated Pct	Memory Allocated MB	Memory Allocated Pct	Capped Mode	Shared Mode	Machine ID
h17p48	48	Not Activated	aixlinux	0	0.00	0	0	0	uncapped	dedicated	0
h17p47	47	Running	aixlinux	0	0.00	0	3	0	uncapped	shared	0
h17p46	46	Not Activated	aixlinux	0	0.00	0	0	0	uncapped	dedicated	0
h17p45	45	Not Activated	aixlinux	0	0.00	0	0	0	uncapped	dedicated	0
h17p44	44	Running	aixlinux	0	0.20	1	1024	1	uncapped	shared	E6964B4C
h17p43	43	Running	aixlinux	0	0.20	1	1024	1	uncapped	shared	E6964B4C
h17p42	42	Not Activated	aixlinux	0	0.00	0	0	0	uncapped	dedicated	0
h17p41	41	Not Activated	aixlinux	0	0.00	0	0	0	uncapped	dedicated	0
h17p40	40	Running	aixlinux	1	0.20	1	1024	1	uncapped	shared	E6964B4C
h17p39	39	Running	aixlinux	0	0.30	1	1024	1	capped	shared	E6964B4C
h17p38	38	Running	aixlinux	0	0.20	1	1024	1	uncapped	shared	0
h17p37	37	Running	aixlinux	0	0.20	1	1024	1	uncapped	shared	E6964B4C
h17p36	36	Not Activated	aixlinux	0	0.00	0	0	0	uncapped	dedicated	0

Hub Time: Tue, 05/05/2009 01:53 PM Server Available CEC Resource Inventory - boson1.raleigh.ibm.com - SYSADMIN



# Integration with IBM Systems Director

The image displays two overlapping windows from the IBM Systems Director console. The top window, titled 'KPK\_CEC\_avg\_mem\_usage\_Warn - lepton.raleigh.ibm.com - SYSADMIN \*ADMIN MODE\*', shows a 'Take Action' dialog box on the left and a table of system metrics on the right. The table is divided into 'Initial Situation Values' and 'Current Situation Values'. Both tables show a value of 3 for 'UnMonitored Active LPARs' and 72% for 'Total Monitored Memory Used Pct'. The bottom window, titled 'IBM Systems Director - lepton.raleigh.ibm.com - SYSADMIN \*ADMIN MODE\*', shows the main console interface with a tree view on the left and a 'Welcome Itmuser' page on the right. The tree view highlights the 'KPK\_CEC\_avg\_m' resource. The right pane shows a table of resources for 'IBM 9131 S2A 101E61G (Computer System)' with columns for Name, Access, Problems, Compliance, and Changed. The table lists several resources (10-1E61G, soa-a-vios, soa-a-1, soa-a-2, soa-a-3, soa-a-4, soa-a-5) all with 'OK' status.

Initial Situation Values										
	UnMonitored Active LPARs	Total Monitored Memory Used Pct	Node	Timestamp	Monitored Shared LPARs	Machine ID	Monitored Dedicated LPARs	Inst LPARs	Inst LPARs	Inst LPARs
	3	72	bluecloud02:PK	09/01/09 07:46:33	3	C9991D4C	1	0	0	0

Current Situation Values										
	UnMonitored Active LPARs	Total Monitored Memory Used Pct	Node	Timestamp	Monitored Shared LPARs	Machine ID	Monitored Dedicated LPARs	Inst LPARs	Inst LPARs	Inst LPARs
	3	72	bluecloud02:PK	09/01/09 07:46:33	3	C9991D4C	1	0	0	0

Navigate Resources						
Select	Name	Access	Problems	Compliance	Changed	...
<input type="checkbox"/>	10-1E61G	OK	OK	OK	Sep 1, 2009 ...	
<input type="checkbox"/>	soa-a-vios	OK	OK	OK	Sep 1, 2009 ...	
<input type="checkbox"/>	soa-a-1	OK	OK	OK	Sep 1, 2009 ...	
<input type="checkbox"/>	soa-a-2	OK	OK	OK	Sep 1, 2009 ...	
<input type="checkbox"/>	soa-a-3	OK	OK	OK	Sep 1, 2009 ...	
<input type="checkbox"/>	soa-a-4	OK	OK	OK	Sep 1, 2009 ...	
<input type="checkbox"/>	soa-a-5	OK	OK	OK	Sep 1, 2009 ...	

# ACTIVE MEMORY EXPANSION (V 6.2.2 new)

The screenshot displays the 'Active Memory Expansion' management console. The interface includes a Navigator pane on the left showing a tree view of system resources, with 'Memory' selected. A context menu is open over the 'Memory' node, with 'Active Memory Expansion' checked. The main area is divided into three panes: 'Effective Memory Layout' (a 3D bar chart showing Uncompressed Data Size MB, Compressed Data Size MB, and Deficit Memory MB), 'Expansion Factors' (a 2D bar chart comparing Target Memory Expansion Factor and Current Memory Expansion Factor), and a summary table.

**AME Summary Table:**

AME Mode	True Memory Size MB	Compressed Memory Free Pct	Target Memory Expansion Factor	Current Memory Expansion Factor	Compression Ratio	CPU Used Pct	Deficit Memory MB		
Enabled	8192.00	16384.00	17.61	8.00	2.00	2.00	6.31	51.32	0.00

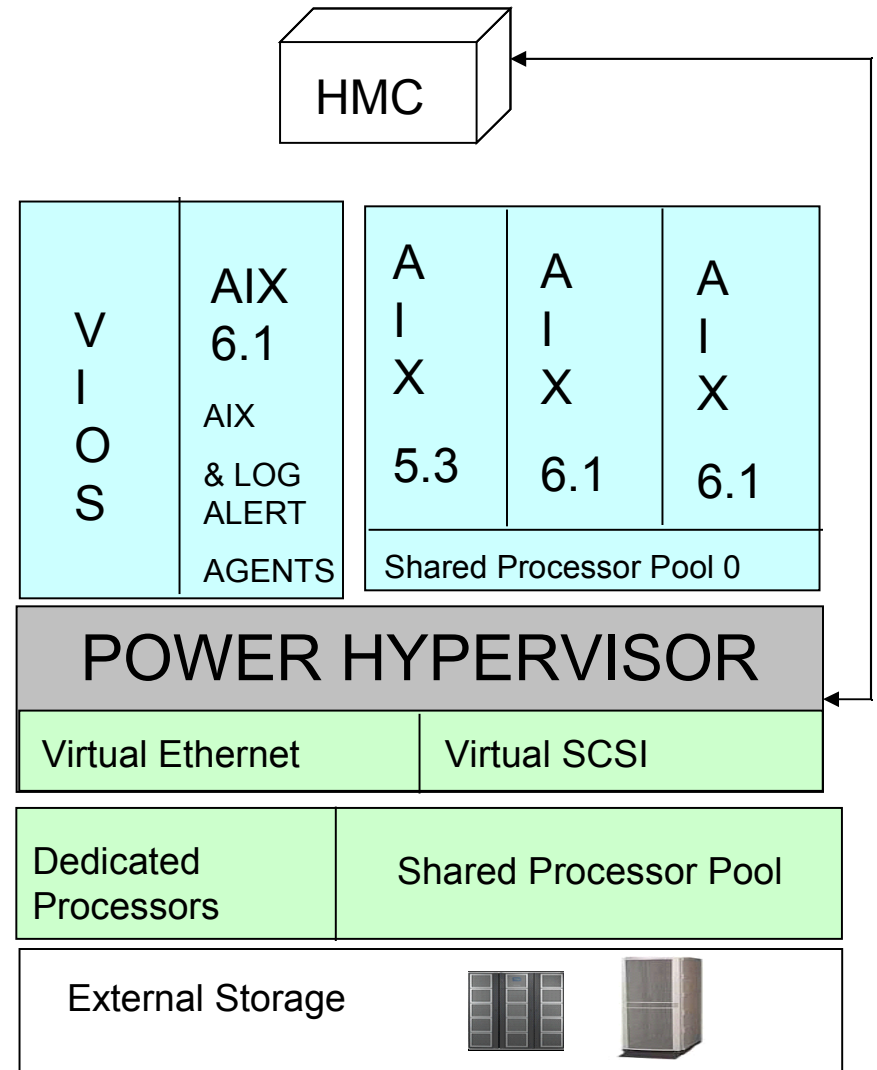
**AME Details Table:**

Target Compressed Memory Size MB	Max Compressed Memory Size MB	Minimum Uncompressed Memory Size MB	Compressed Memory Size MB	Compressed Memory In Use MB	Compressed Memory In use Pct	Compressed Memory Pct	Compressed Memory Page Ins	Compressed Memory Page Outs	Compressed Number Of Working Pages	Uncc Mem
1545.00	4997.00	0.00	1558.00	1443.00	92.00	17.61	65638	65582	2329371	6633

At the bottom of the console, there are status indicators: 'Hub Time: Sun, 03/21/2010 05:23 PM', 'Server Available', and the window title 'Active Memory Expansion - boson3 - SYSADMIN'.

# Monitoring the AIX Operating System

- Benefits
  - Insight into resource utilization on LPARs and WPARs
  - Which processes are using the most resource?
  - Alerts based on errors in logs
- AIX Premium Agent
  - CPU / memory usage & load
  - Processes & their cpu usage
  - Resource consumption by individual WPARs
  - Storage & Network usage
- Tivoli Monitoring for Unix Logs
  - Monitor system and user logs
  - Customizable alerts based on messages in the logs
  - Hardware alerts
- Install agents in each LPAR you wish to monitor



# WPAR Supplementary Workspace links

The screenshot displays the WPAR Summary application interface for the system h17lp39.austin.ibm.com. The interface includes a navigation tree on the left, two charts showing resource utilization, and a table of WPAR configurations.

**WPAR CPU Utilization by WPARs:** A 3D bar chart showing CPU utilization for h17wpar4, with a legend for LPAR CPU Consumed Pot.

**WPAR Memory Utilization by WPARs:** A 3D bar chart showing memory utilization for h17wpar4, with a legend for LPAR Memory Used Pot.

**WPAR Status and Configuration Table:**

WPAR Name	Type	State	Admin Operation	Hostname	Autostart	Shares usr Dir	IP Address	RC Is Active
h17wpar4	Custom	Active		h17wpar4	No	No		Yes
syswpar1				syswpar1	No	No		Yes
syswpar2				syswpar2	No	No		Yes

A context menu is open over the table, listing the following options:

- WPAR CPU details
- WPAR Memory details
- WPAR Network and Filesystem details
- WPAR Details
- WPAR Process Views
- Link Wizard...
- Link Anchor...

The status bar at the bottom shows: Hub Time: Mon, 02/18/2008 09:59 PM, Server Available, and WPAR Summary - h17lp39.austin.ibm.com - SYSADMIN.

# Monitor Processes

Process Views - IBM-6205SWFW - SYSADMIN

File Edit View Help

Navigator View: Physical

- AIX Premium
  - Top Resources
  - System
  - Memory
  - Process
  - Storage
  - Networking
  - User
  - WPAR

Physical

Utilization

Total Num Processes	Process Context Switches per Sec	Load Avg	Utilization Avg
96	184	15	0

Per Process Information

Process Name	Process ID	Parent Process ID	CPU Pct	Nice	User Name	IO Page Fault per Sec	Non IO Page Fault per Sec	Text Size	Resident Text Size	Resident Data Si
init	1	0	0.0	20	root	0	0	29	11	170
sched	12294	0	0.0	41	root	0	0	0	0	112
lrud	16392	0	0.0	41	root	0	0	0	0	128
vmptact	20490	0	0.0	41	root	0	0	0	0	112
psmd	24588	0	0.0	41	root	0	0	0	0	128
vmmd	28686	0	0.0	41	root	0	0	0	0	224
memgrdd	32784	0	0.0	41	root	0	0	0	0	112
pilegc	40980	0	0.0	41	root	0	42	0	0	160
xmgc	45078	0	0.0	41	root	0	0	0	0	112
netm	61470	0	0.0	41	root	0	0	0	0	112
gil	65568	0	0.0	41	root	0	0	0	0	240
wlmsched	69666	0	0.0	41	root	0	0	0	0	128
armtrace_kproc	73764	0	0.0	41	root	0	0	0	0	112
X	82074	127094	0.0	20	root	0	2345	3217	454	1770
j2pg	86090	0	0.0	41	root	0	0	0	0	512
syncd	90274	1	0.0	20	root	0	101759	2	2	133
random	94454	1	0.0	20	root	0	0	0	0	112

Hub Time: Sun, 01/24/2010 09:17 AM

Server Available

Process Views - IBM-6205SWFW - SYSADMIN

# AIX AGENT: LPAR MIGRATION AWARENESS (V 6.2.2 new)

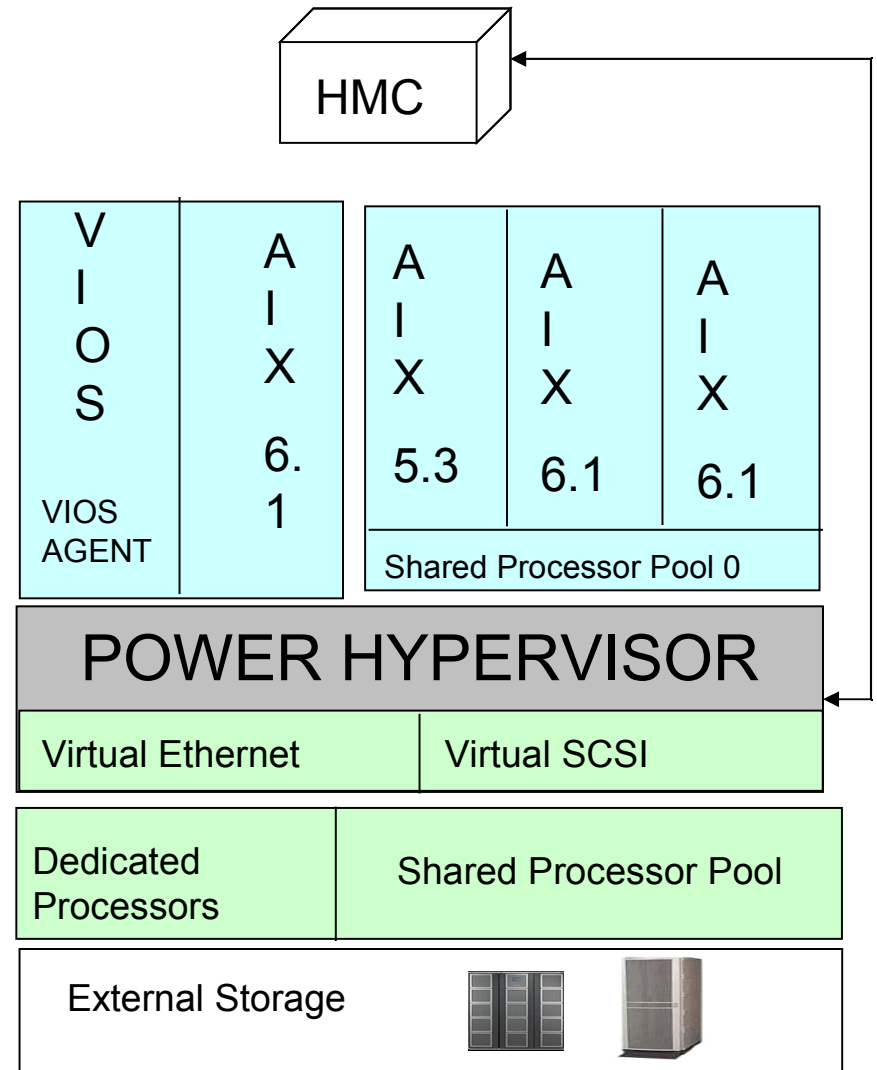
The screenshot displays the 'LPAR Information - 9.42.23.40 - SYSADMIN' console with several panels:

- Navigator:** Shows a tree view of the system hierarchy: Enterprise > UNIX Systems > p7e07 > AIX Premium > Top Resources > System.
- LPAR CPU Utilization:** A 3D pie chart showing CPU usage. Legend: User CPU Pct (yellow), System CPU Pct (blue), IO Wait CPU Pct (red), Idle CPU Pct (green).
- CPU Entitlement:** A 3D bar chart comparing 'Number of Physical CPUs' (yellow bar, ~35) and 'Entitlement' (blue bar, ~1).
- LPAR Attributes:** A table with columns: LPAR Number, SMT Mode, Shared Mode, Capped Mode, Min Memory, Online Mem, Max Memory, Capacity Weight, CPU Pool ID.
- LPAR Utilization:** A table with columns: Max Virt CPUs, Min CPU Capacity, Max CPU Capacity, CPU Capacity Increment, Machine ID, Last Machine ID, Online Mem, Max Dispatch Latency, Unallocated CPU In Pool, CPU Entitlement, Capacity Weight. A tooltip indicates: 'The previous frame hardware ID of this LPAR before it was migrated to the current frame'.
- LPAR CPUs:** A 3D bar chart showing: Number of Logical CPUs (yellow bar, ~16), Number of Virtual CPUs (blue bar, ~4), Number of Physical CPUs in Shared Pool (red bar, ~0), and Available CPUs in Pool (green bar, ~0).
- Active Memory Sharing (AMS) Pool:** A table with columns: AMS Mode, AMS Pool, AMS Physical, AMS Mem, AMS Memory, AMS Memory, AMS Memory, Hypervisor Page, Hypervisor Page.

At the bottom, the status bar shows: Hub Time: Wed, 02/24/2010 03:15 PM, Server Available, and LPAR Information - 9.42.23.40 - SYSADMIN.

# VIOS AGENT

- Benefits
  - Monitors the availability & health of the VIOS LPAR
  - Monitors the availability and performance of VIOS resources – as they relate to client LPAR
- Features
  - Mapping between physical network / storage devices and virtual devices
  - Device Status
  - CPU/Memory usage
  - Storage I/O & disk usage
  - Network Adapter throughput details
- Agent pre-installed on the VIOS



# Storage Mapping with Utilization

Storage Mappings - boson1.raleigh.ibm.com - SYSADMIN

File Edit View Help

View: Physical

Storage Mappings

VIOS Physical Adapter	Disk	LV Name	LUN ID	VSCA Slot	Client Partition ID	Client Partition Name	Client Partition State	Disk Transfers per Sec	Disk Transfers Sec Pct
0 scsi0	hdisk1		0x8100000000000000	2	4	frankvios_client	Runni...	0.5	0.0

Physical

Disk Usage

Legend:   
■ Used MB   
■ Free MB

Storage Mappings Details

VIOS Name	Hostname	IP Address	Partition ID	VSSA Slot	VSSA Name	VTD Name	VIOS Physical Adapter	Disk	LV Name	LUN ID	Client Partition Name	Client Hostname	Client IP Address
frankvios	frankvios.vios.austin.ibm.com	9.3.92.178	5	3	vhos...	vtscsi0	scsi0	hdisk1		0x8100000000000000	frankvios_client		

Hub Time: Tue, 05/05/2009 01:41 PM    Server Available    Storage Mappings - boson1.raleigh.ibm.com - SYSADMIN



# Device Status

Status - lep2003e.raleigh.ibm.com - SYSADMIN \*ADMIN MODE\*

File Edit View Help

Navigator View: Physical

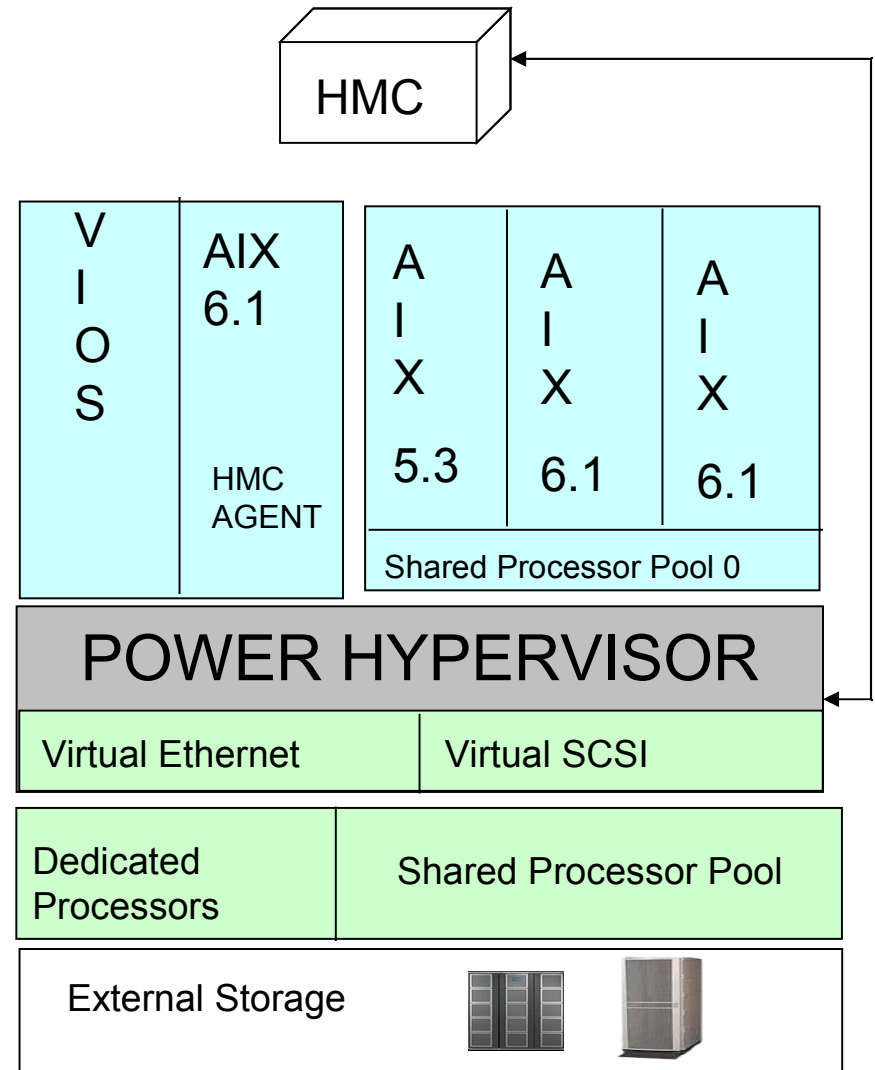
Device Status

Name	Parent	State	Class	Type
en0	inet0	Stopped	if	en
en3	inet0	Available	if	en
ent0	pci2	Available	adapter	14108902
ent1	pci2	Stopped	adapter	14108902
ent2	vio0	Available	adapter	IBM_l-lan
ent3	vio0	Available	adapter	sea
hdisk0	scsi0	Available	disk	scsd
hdisk1	scsi0	Available	disk	scsd
iscsi0		Available	driver	iscsi
io0	inet0	Available	if	lo
pager0	vbsd0	Available	adapter	pager
scsi0	pci3	Available	adapter	sym1010
vasi0	vio0	Available	adapter	IBM_VASI-1
vbsd0	vasi0	Available	adapter	vbsd
vhost0	vio0	Available	adapter	IBM_v-scsi-host
vhost1	vio0	Available	adapter	IBM_v-scsi-host
vhost2	vio0	Available	adapter	IBM_v-scsi-host
vhost3	vio0	Available	adapter	IBM_v-scsi-host
vs0	vio0	Available	adapter	hvterm1

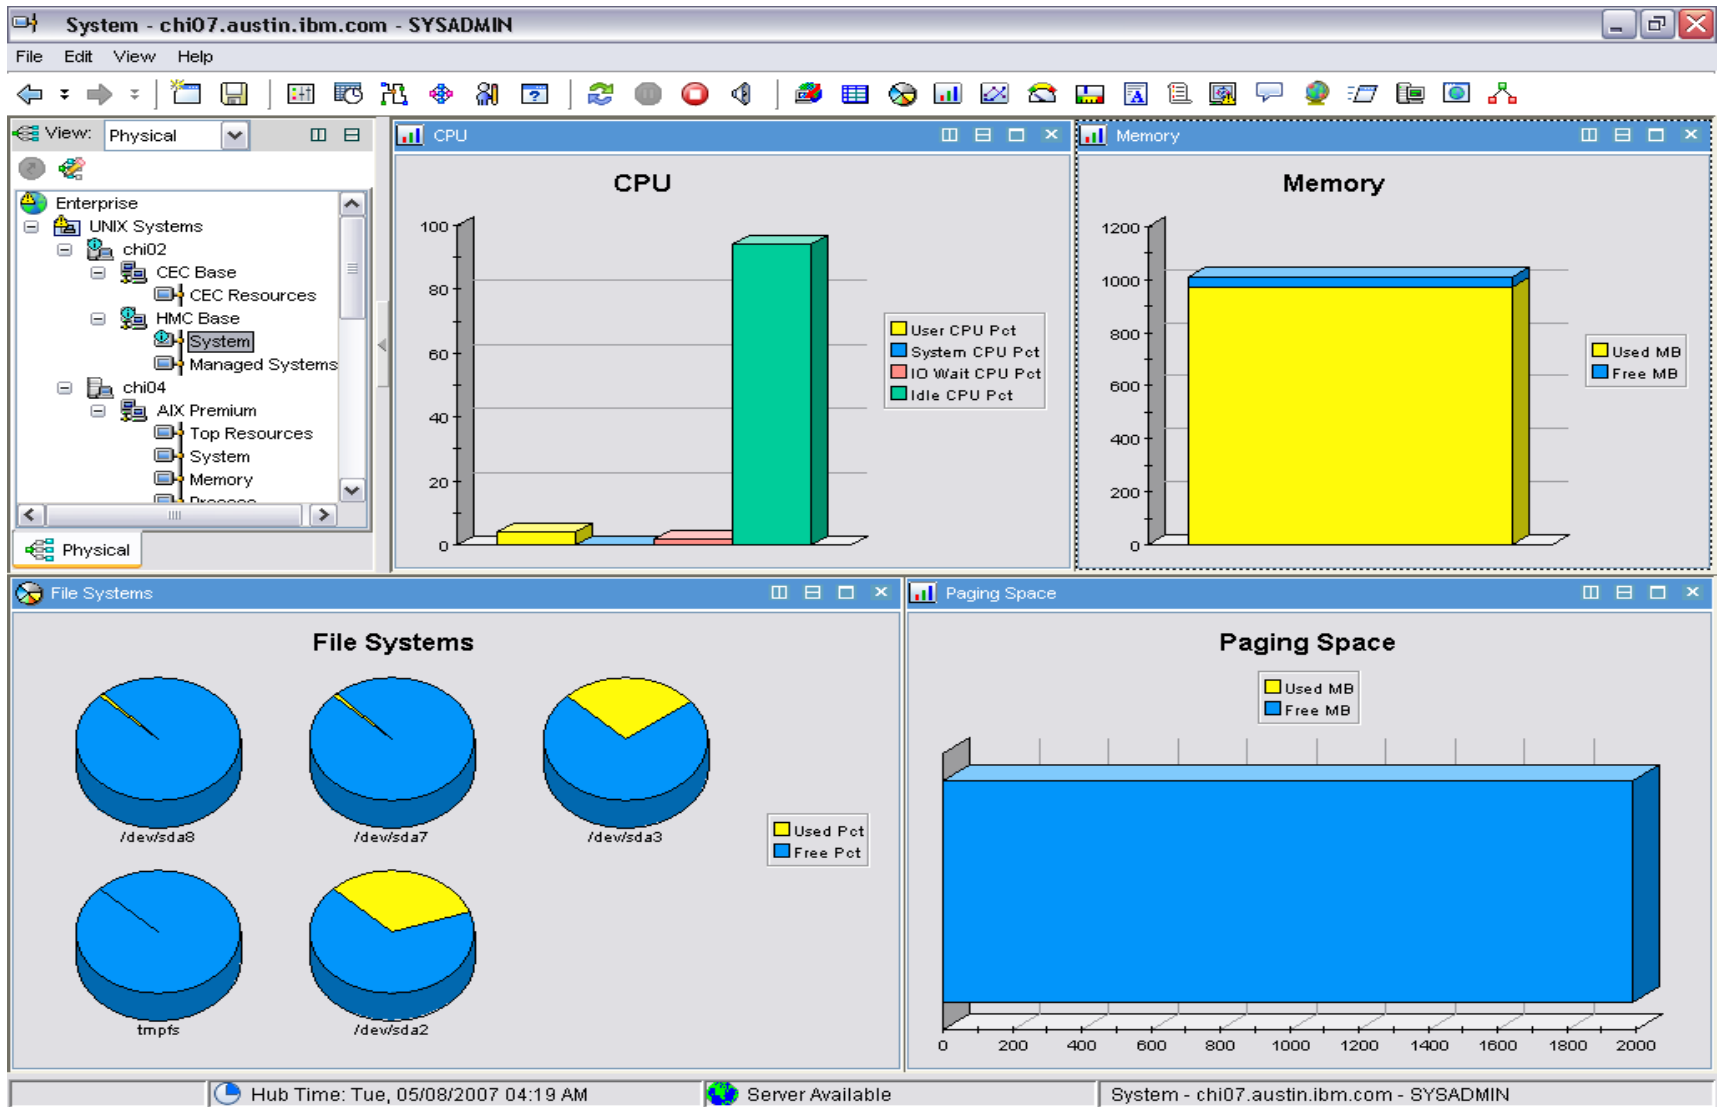
Hub Time: Wed, 10/08/2008 08:10 PM Server Available Status - lep2003e.raleigh.ibm.com - SYSADMIN \*ADMIN MODE\*

# HMC AGENT

- Benefits
  - Ensures availability of the hardware management console and reports on the status of systems managed by the HMC.
- Features
  - Provides information on all the managed systems connected to the HMC
  - Monitors the resource (CPU, memory, storage, network) usage on the HMC
- Install on an LPAR – remotely monitors the HMC



# HMC System Resources



# HMC AGENT: SITUATION EVENTS MIGRATION AWARENESS (V 6.2.2 new)

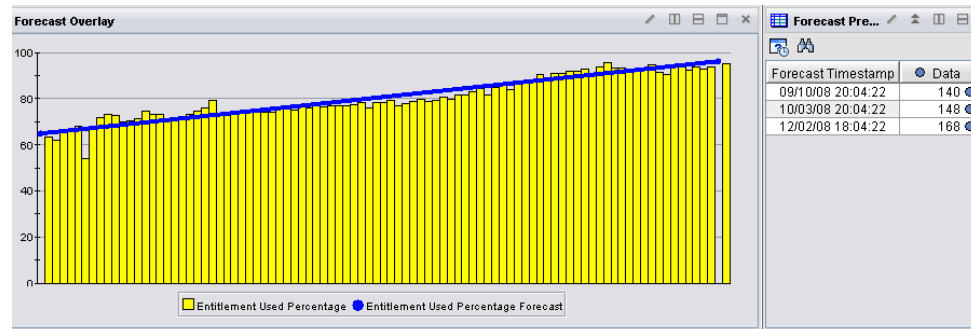

The screenshot displays the HMC Agent interface with the following components:

- Navigator:** Shows a tree view of managed CECs. The selected node is 'KPH\_LPAR\_Deleted\_Info' under 'p7-hmc::PH'.
- Initial Situation Values:** A table showing event details.
 

Event	Node	Timestamp	Time	LPAR Name	LPAR Num	CEC Name
LPAR deleted	p7-hmc:robot16:PH	02/24/10 15:19:32	Wed Feb 24 15:16:13 2010	p7b09	10	p7b-SN1000CB
- Current Situation Values:** An empty section for current values.
- Command View:** A 'Take Action' panel with fields for 'Name' (set to '<Select Action>') and 'Command'. A 'Run' button is at the bottom.
- Expert Advice:** A panel displaying the following information:
  - Expert Advice** (IBM logo)
  - KPH\_LPAR\_Deleted\_Info**
  - Situation Description**: A Logical Partition has been deleted from this CEC.
  - Suggested Actions**: When an LPAR is deleted from a CEC and the resources are freed, it is usually a good time to see if redistribution of resources to the remaining LPARs is warranted.
  - Copyright IBM Corp. 2005, 2010 All Rights Reserved US Government Users
- Status Bar:** Shows 'Hub Time: Wed, 02/24/2010 03:04 ...', 'Server Available', and the window title 'KPH\_LPAR\_Deleted\_Info - p7b09 - Wed Feb 24 11:24:19 CST 2010 - 9.42.23.40 - SYSADMIN'.

# Performance Analyzer - Predict resource constraints

- Future predictions based off historical warehoused data.
- Predict when CPU, Memory, Storage and Network will meet limitations
- Proactively respond to near-term & long term performance problems
- Warehouse future trends for capacity analysis reports



# Performance Reporting & Monitoring



- Overall capacity/utilization; capacity by CEC, Pool, LPAR – forecast trend
- Top and low utilized LPARs
- Workload utilization trends – peak v/off peak
- Workload placement

# Deployment Tips

- Operating System
  - AIX 5.3 TL 5 and above (32 bit / 64 bit)
  - AIX 6.1 (64 bit)
- OpenSSH to communicate with HMC for CEC, VIOS and HMC agents
  - key.pl script provided to configure OpenSSH communication
  - On Vios, use the cfgsvc command to configure the agents. More details:  
<http://www-01.ibm.com/support/docview.wss?uid=swg21424888>  
<http://www-01.ibm.com/support/docview.wss?uid=swg21432003>
- CEC Agent now supports remote monitoring
  - For support of live partition mobility, it is recommended to run the CEC agent on the VIOS partition.

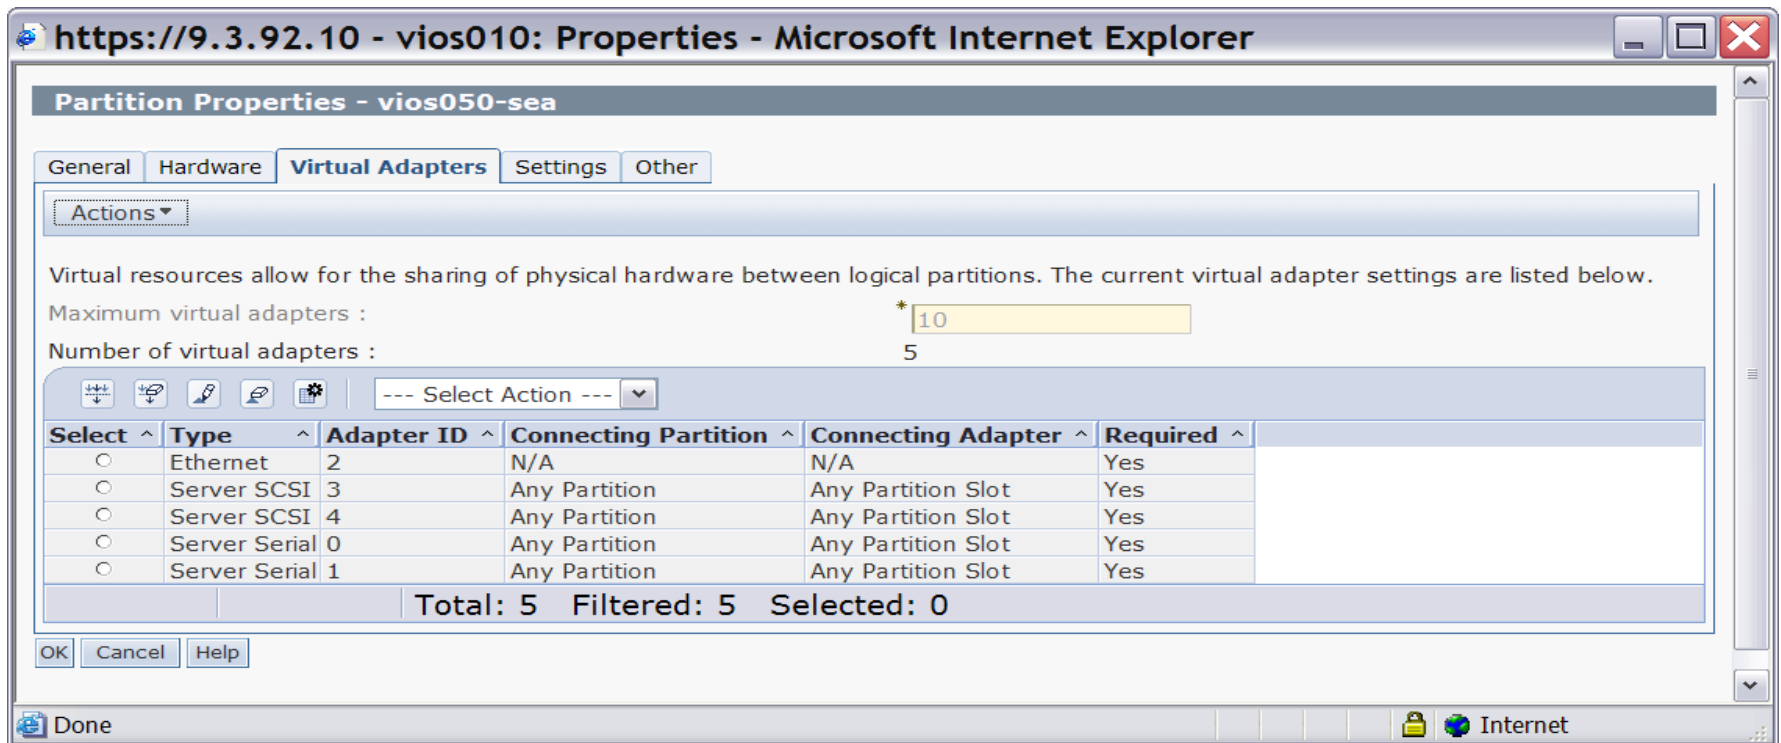
## Deployment Tips cont...

- CEC agent makes RSiOpen( ) call to all LPARs, hence:
  - Port 2279 needs to be open on LPARS
  - xmservd Or xmtopas needs to be running on LPARS
  - Multi-homed LPARS require an underlying AIX fix addressed by APAR IZ63567
  - Dependence on perfagent.tools. Refer to:  
<https://www-304.ibm.com/support/docview.wss?uid=swg21447016>
- Shared memory corruption
  - Affects all the Power Systems agents
  - Problem is with underlying SPMI libraries
  - AIX 6.1 – APARS IZ56426, IZ64808
  - AIX 5.3 – APARS IZ56425
- HMC Agent requires DNS to be configured



# Deployment Tips cont...

- VIOS Agent Storage Mapping Client data will be missing if the HMC settings for VIOS LPAR Virtual Adapter “Connecting Partition” and “Connecting Adapter” are set to “Any Partition” or “Any Partition Slot”, respectively. They need to be set to the specific Partition and Connecting Adapter of the VIOS client in order for the ITM VIOS agent to map them properly.



https://9.3.92.10 - vios010: Properties - Microsoft Internet Explorer

Partition Properties - vios050-sea

General Hardware **Virtual Adapters** Settings Other

Actions

Virtual resources allow for the sharing of physical hardware between logical partitions. The current virtual adapter settings are listed below.

Maximum virtual adapters : \* 10

Number of virtual adapters : 5

--- Select Action ---

Select ^	Type ^	Adapter ID ^	Connecting Partition ^	Connecting Adapter ^	Required ^
<input type="radio"/>	Ethernet	2	N/A	N/A	Yes
<input type="radio"/>	Server SCSI	3	Any Partition	Any Partition Slot	Yes
<input type="radio"/>	Server SCSI	4	Any Partition	Any Partition Slot	Yes
<input type="radio"/>	Server Serial	0	Any Partition	Any Partition Slot	Yes
<input type="radio"/>	Server Serial	1	Any Partition	Any Partition Slot	Yes

Total: 5 Filtered: 5 Selected: 0

OK Cancel Help

Done Internet

# Deployment Tips cont...

- Upgrading VIOS
  - Make sure the agents are stopped before VIOS upgrade:  
<https://www-304.ibm.com/support/docview.wss?uid=swg21448092>
- Mapping of Agents and VIOS releases

<b>VIOS Release</b>	<b>Based on AIX</b>	<b>ITM Fileset Version</b>	<b>ITM Version</b>
1.5.2.0	53N	6.1.2.1	06.10.01.00
2.1.0.0	61D	6.2.0.1	06.20.00.00 (pk) 06.20.00.01 (va)
2.1.1.0	61F	6.2.1.0	06.20.01.00
2.1.2.0	61H	6.2.2.0	06.21.00.03
2.1.3.0	61J	6.2.3.0	06.22.00.00
2.2.0.0	61L	6.2.4.0	06.22.01.00

# Monitoring the Power AIX environment

<b>Capabilities</b>	<b>Key Benefits</b>
Proactive monitoring of HMC & VIOS performance and availability	Ensures availability of HMC/VIOS for making configuration changes/creating LPARs, etc. and collecting performance & availability data for AIX environment
Real time, historical & predictive usage monitoring of CPU, Memory, Storage & Network resources for the Frame, pool and by LPAR & WPAR	Understand whether utilization goals are being met, if additional LPARs can be added to the environment, predict when resources will meet/exceed capacity
Monitor AIX processes	Understand which processes are consuming the most resource
Monitor performance of resources and availability of critical resources (Servers, Disks, adapters, MPIO paths, etc.)	Reconfigure resources for optimal performance, fix failed adapters, etc. to keep redundancy.

# Documentation

- User Guides
  - <http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc/using.htm>
- Wiki
  - <http://www.ibm.com/developerworks/wikis/display/tivolimonitoring/System+P+Agents>
- Forum – ITM for Power Systems
  - <http://www.ibm.com/developerworks/forums/forum.jspa?forumID=1031&start=15>



QUESTIONS?