IBM Systems and Technology Group University 2005

ON DEMAND BUSINESS™

2/10/2005

IBM Systems and Technology Group University 2005

Implementing AIX 5L on eServer i5: Let's Shed Some Light

Vess Natchev
AIX 5L on eServer i5 Technical Specialist
IBM Rochester

2/10/2005

Agenda

- iSeries and pSeries hardware terminology
- HMC view of AIX 5L direct I/O
- Supported hardware overview
- Supported hardware FAQ
- List of supported hardware for AIX 5L
- Overview of Virtual I/O
- Virtual Ethernet, SCSI and serial adapters
- AIX 5L system management
- Sizing for AIX 5L
- Performance tools

Hardware and Terminology

eServer i5 520 + expansion I/O tower

eServer i5 570 + expansion I/O drawers

- Traditional iSeries terminology
  - I/O towers
  - Main storage, IOP, IOA
  - SLIC, PTFs
- Traditional pSeries terminology
  - I/O drawers
  - HMC
  - Memory, PCI adapter
  - Kernel, patches
- eServer i5
  - Model 520 can be in tower or rack
  - Model 570 is rack only
  - Now has towers and drawers
  - Both towers and drawers appear as "units" in HMC
  - Units contain buses and adapters

HMC View of Direct I/O

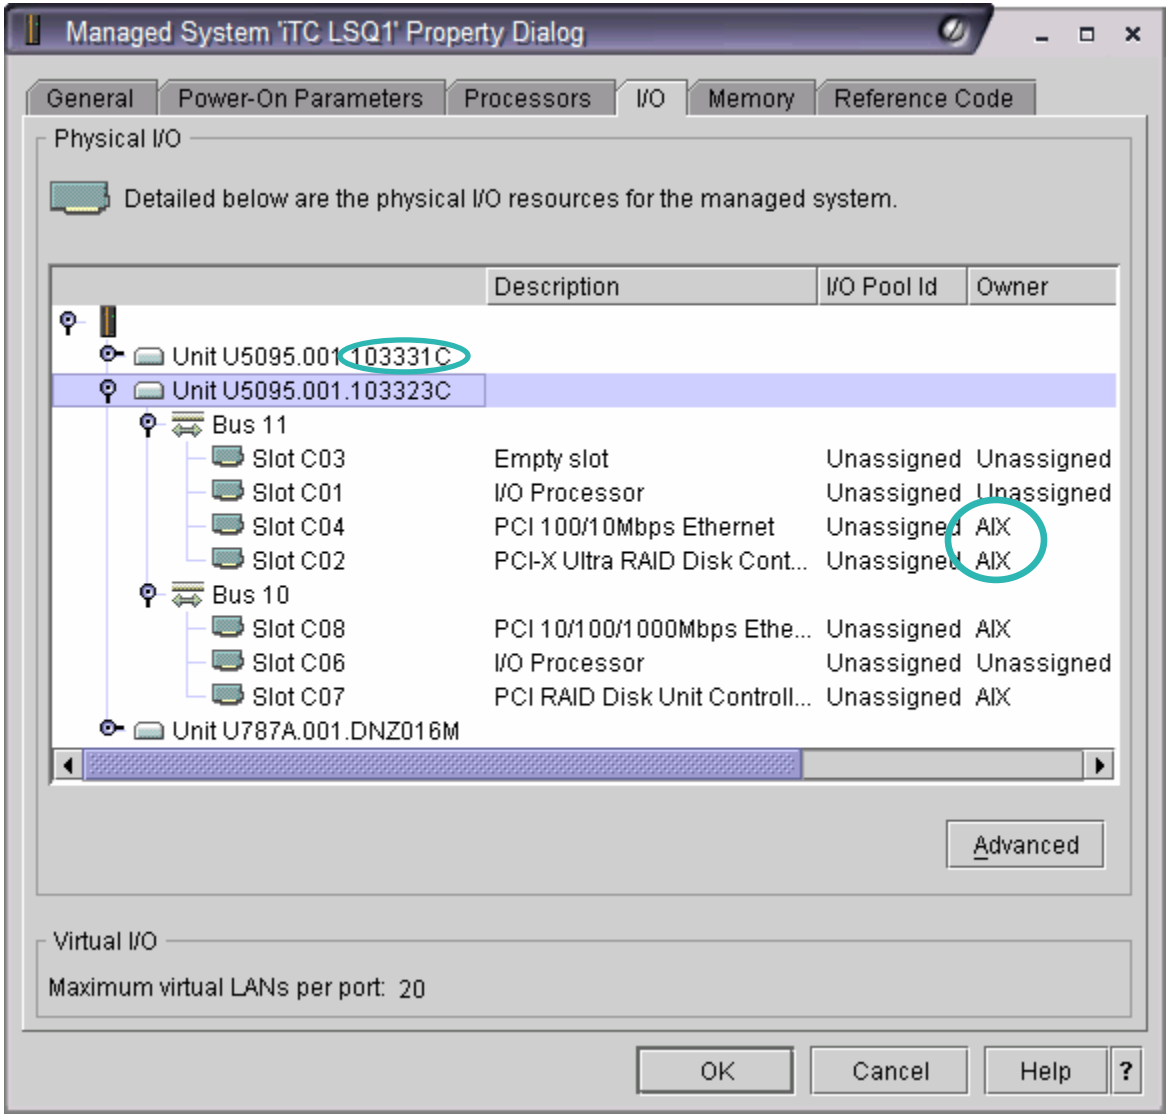

- HMC view of eServer i5's physical I/O resources:
  - AIX 5L and Linux use IOAs, not IOPs
  - Direct I/O assigned in each partition's profile
  - Towers and drawers have serial numbers
- Common set of AIX 5L and Linux adapters for eServer i5 and p5.
- HOWEVER: not all eServer i5 AIX 5L and Linux adapters are supported on eServer p5, and vice versa

Supported Hardware Overview

- Supported I/O for AIX 5L document: "AIX 5L Facts and Features"
- AIX 5L and Linux on eServer i5 support dynamic movement of I/O resources between partitions (**performed in HMC**)
  - "Required" I/O adapters cannot be moved dynamically, only "desired" ones

- TIP: In consolidated environment with multiple AIX 5L 5.3 partitions using Virtual Storage, one DVD or tape drive can be purchased, and later switched between those partitions without IPL (reboot). Controller for DVD/tape drive in this case is 0645/5712

- HMC **will enforce** placement rules for AIX 5L and Linux disk controllers: they cannot be "downstream" of i5/OS IOP
  - Applies only to internal DASD controllers (such as 0627/2780)
  - Does not apply to Fibre Channel controllers (such as 0625/5704)

- Same rules enforced by LVT. **Always use latest LVT from** http://www-1.ibm.com/servers/eserver/iseries/lpar/systemdesign.htm

Supported Hardware FAQ

- **Is there a Fibre Channel adapter for AIX 5L on eServer i5?**

Yes, the 0625/5704. It does show up in the LVT as "Fibre Channel Tape Controller," because it is one for i5/OS. For AIX 5L and Linux, it controls disk (FAStT, ESS) and tape

- **There are 2 embedded Ethernet ports in the system unit. Can they be split between 2 different partitions (of any kind)?**

No, the dual Ethernet ports are controlled by the same adapter on the backplane. They will show up as 2 network resources in a partition, but cannot be split up between any 2 partitions. They are usable by i5/OS or AIX 5L

- **There are 2 embedded USB ports in the system unit. Can they be split between 2 different partitions (of any kind)?**

No, the two USB ports are controlled by the same adapter on the backplane. They will show up as 2 USB resources in a partition, but cannot be split up between any 2 partitions. They are usable by AIX 5L, but not i5/OS

- **Do I have to use a 7311-Dxx for AIX 5L features on eServer i5?**

No, AIX 5L features can be placed in any of the system units, 0595/5095, 5094/5294 or 7311-Dxx.
Use LVT to find out which AIX 5L features are valid where in each system unit or tower

Supported Hardware FAQ

- **What are my options for CD/DVD for AIX on eServer i5?**

There are several:

-- An external SCSI DVD drive, attached to a 0645/5712

-- The 2nd DVD slot in the system unit, controlled by the embedded IDE controller (eServer i5 520 and 550 only)

-- The 1st DVD slot in the system unit, controlled by the 570B/5709

-- A DVD drive in a 5094/5294 tower, controlled by an 0627/0628/0645

There is no recommendation, the choice depends on the full system configuration.

NOTE: these are the DVD options for AIX 5L without regard to any other partitions. For example, if i5/OS controls the 5709 in the system unit, the 1st DVD slot will not be an option

- **If I allocate an 0645/5712 to an AIX 5L partition, can it control both an external DVD/tape, and internal disks?**

Yes. However: the 0645 is a 2-bus SCSI adapter, each of which has an external and an internal port. Therefore, only 1 of the ports for the same bus can be used at the same time

- **If I have an older HVD (high-voltage differential) external SCSI tape drive, is there an adapter for it for AIX 5L on eServer i5?**

Yes, the 6204. All other supported SCSI adapters (such as the 0627 or 0645) are LVD

Supported Hardware List for AIX 5L

IBM @server i5

	9406-520 rack drawer; deskside	9406-550 rack drawer; deskside	9406-570 rack drawer	9406-595 rack drawer
Storage Interfaces				
0628 Direct Attach 5703 PCI-X SCSI RAID Disk Unit Controller	Y	Y	Y	Y
0627 Direct Attach 2780 PCI-X Ultra4 RAID Disk Controller	Y	Y	Y	Y
5709 RAID Enabler Card	Y	Y	Y	N
0625 Direct Attach 5704 2 Gigabit Fibre Channel PCI-X	Y	Y	Y	Y
0645 Direct Attach 5712 PCI-X Tape Controller	Y	Y	Y	Y
6203 PCI-3 SCSI Adapter	N	Y	Y	Y
6204 Differential SCSI Adapter	Y	Y	Y	Y
0611 Direct Attach 2765 Fibre Channel Tape Controller	Y	Y	Y	Y
0642 PCI Ultra 3 Raid Adapter	Y	Y	Y	N
0638 PCI 40MB SSA Adapter	N	N	Y	Y
0639 128MB DIMM SSA Adapter Memory	N	N	Y	Y
0640 32MB Non-volatile Fast Write Cache for 6225 and 6230	N	N	Y	Y
Communications, Connectivity and Encryption				
0637 10/100 Mbps 4-port Ethernet Adapter	N	N	Y	Y
5706 1Gbps Ethernet Adapter	Y	Y	Y	Y
0620 Direct Attach 5700 PCI 1 Gbps Ethernet IOA	Y	Y	Y	Y
0621 Direct Attach 5701 PCI 1 Gbps Ethernet UTP IOA	Y	Y	Y	Y
2943 8 Port Async Adapter	Y	Y	Y	Y
0635 SDLC/X.25-2 port Adapter	Y	Y	Y	Y
2947 PCI Multipprotocol Adapter	Y	Y	Y	Y
4960 Cryptographic Accelerator	N	N	Y	Y
4963 Cryptographic Co-processor	N	N	Y	Y
0634 Direct Attach 2944 128 port Async Adapter	Y	Y	Y	Y
8136 Remote Async Node (Rack)	Y	Y	Y	Y
8137 RS232 Remote Async Node	Y	Y	Y	Y
6312 Quad Digital Trunk Adapter	Y	Y	N	N
4959 PCI 16/4 Mbps Token-Ring IOA	Y	Y	Y	Y
4953 155Mbps UTP ATM Adapter	N	N	Y	Y
4957 155Mbps ATM Fibre Adapter	N	N	Y	Y
2946 PCI 622 Mbps ATM Fibre Adapter	N	N	Y	Y
5718 10 Gbps Ethernet Adapter	Y	Y	Y	Y
4962 PCI 100/10 Mbps Ethernet IOA	Y	Y	Y	Y
2732 PCI Serial HIPPI Adapter	N	N	Y	Y

Supported Hardware List for AIX 5L

IBM @server i5

	9406-520 rack drawer; deskside	9406-550 rack drawer; deskside	9406-570 rack drawer	9406-595 rack drawer
Displays and Display Adapters [a]				
0632 USB 2.0 Adapter	Y	Y	Y	Y
0633 Graphics Adapter	Y	Y	Y	Y
1876 20" Flat Panel Monitor	Y	Y	Y	Y
3637 15" Color Monitor	Y	Y	Y	Y
3639 17" Color Monitor	Y	Y	Y	Y
3628 21" Color Monitor	N	N	Y	Y
3638 21" Color Monitor	Y	Y	Y	Y
2737 4-port USB 1.1 Adapter	N	N	Y	Y
8841 USB Mouse and Cable to Keyboard	Y	Y	Y	Y
88XX/187X Keyboard	Y	Y	Y	Y
8244 PCI WS Audio Adapter	Y	Y	N	N
2848 PCI 2D Entry Graphics Adapter	Y	Y	Y	Y
6546-00E/00N/01E/01S/0AE/0AN/0BE/0BN/0BS/21N/2BN/31N/4AN/6BN Display	Y	Y	Y	Y
Storage and StorageSubsystems				
4326/7508 35.16 GB 15K rpm Disk Unit	Y	Y	Y	Y
4327/7509 70.56 GB 15K rpm Disk Unit	Y	Y	Y	Y
4328/7510 141 GB 15K rpm Disk Unit	Y	Y	Y	Y
1893 36.4 GB 10K rpm Disk Unit	Y	Y	Y	Y
1894 73.4 GB 10K rpm Disk Unit	Y	Y	Y	Y
1895 146.8 GB 10K rpm Disk Unit	Y	Y	Y	Y
1896 36.4 GB 15K rpm Disk Unit	Y	Y	Y	Y
1897 73.4 GB 15K rpm Disk Unit	Y	Y	Y	Y
2591 External 1.44 GB Diskette Drive	Y	Y	Y	Y
Expansion Towers				
5094/5294 PCI-x Expansion Tower	Y	Y	Y	Y
5095/0595 PCI-x Expansion Tower	Y	Y	Y	Y
7311-D10 I/O Drawer	N	N	Y	Y
7311-D20 I/O Drawer	Y	Y	Y	Y

Supported Hardware List for AIX 5L

IBM @server i5

	9406-520 rack drawer; deskside	9406-550 rack drawer; deskside	9406-570 rack drawer	9406-595 rack drawer
Fibre Channel Directors, Switches, Hubs and Routers				
2031-016/032/216/224/232/L00 McDATA Fibre Channel Switch	Y	Y	Y	Y
2032-001/064/140 McDATA Fibre Channel Director	Y	Y	Y	Y
3534-F08 SAN Switch	Y	Y	Y	Y
Optical Drives and Libraries				
2640 DVD-ROM	Y	Y	Y	N
5751 DVD-RAM	Y	Y	Y	N
4625 CD-ROM	Y	Y	Y	Y
4631 DVD-ROM	Y	Y	Y	Y
4430 DVD-RAM	Y	Y	Y	Y
4630 DVD-RAM	Y	Y	Y	Y
4633 DVD-RAM	Y	Y	Y	Y
7210-020 CD-ROM Drive	Y	Y	Y	Y
7210-025 DVD-RAM Drive	Y	Y	Y	Y
7210-030 DVD Drive	Y	Y	Y	Y
3995-X63/C6X Optical Library	Y	Y	Y	Y
Tape Drives				
1889 80/160 GB Tape Drive	Y	Y	N	N
4482 4 GB ¼" Tape Drive	Y	Y	Y	Y
4682 4 GB ¼" Tape Drive	Y	Y	Y	Y
4684 30 GB ¼" Tape Drive	Y	Y	Y	Y
4487 50 GB ¼" Tape Drive	Y	Y	Y	Y
4687 50 GB ¼" Tape Drive	Y	Y	Y	Y
4685 80 GB ¼" Tape Drive	Y	Y	Y	Y
5753 QIC 30 GB Tape	Y	Y	N	N
5754 QIC 50 GB Tape	Y	Y	N	N
6134 60/150 GB 8 mm Tape Unit	Y	Y	N	N
6258 36/72 GB 4 mm Tape Unit	Y	Y	N	N

Virtual I/O Overview

Virtual I/O

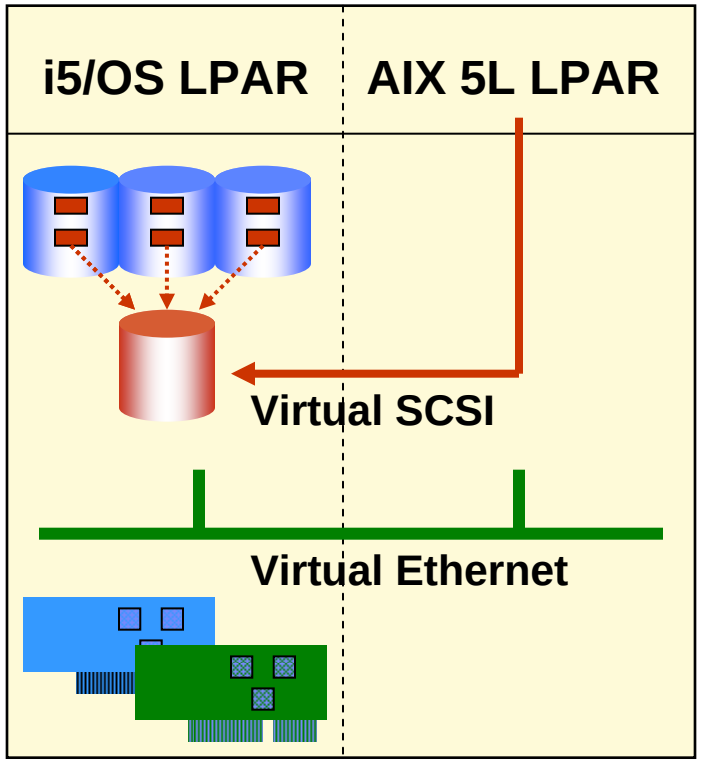

- i5/OS provides virtual disk to AIX 5L
- Improves asset utilization and ROI
- Uses the IBM Virtualization Engine

Direct I/O

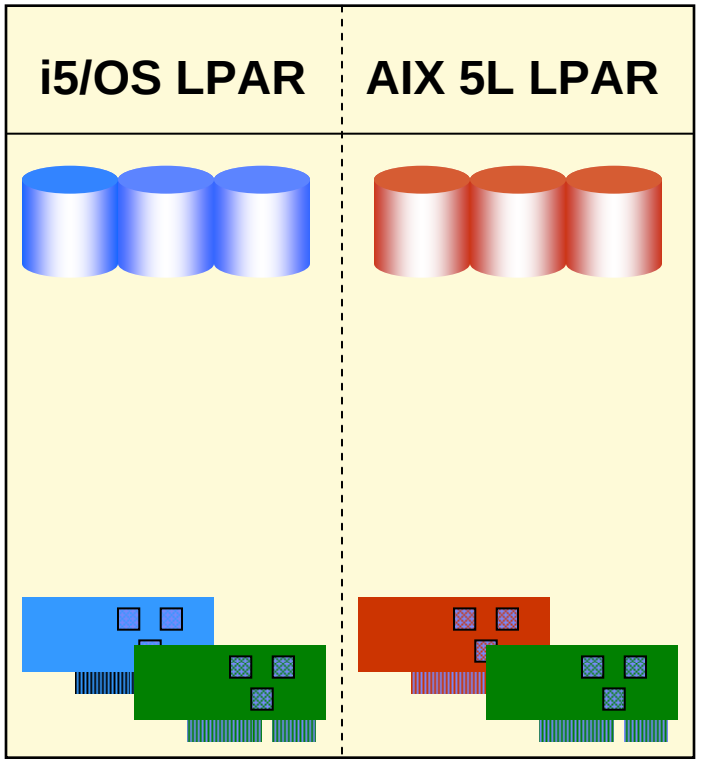

- Resources dedicated to AIX 5L
- AIX 5L management of disk, NICs
- AIX 5L independent of other LPARs

Virtual I/O Overview

- **Each partition has virtual I/O "slots"**
  - Configurable for each partition in HMC
- **Slots can have virtual adapter instance**
  - Ethernet, serial, or SCSI
- **Virtual adapters configured in partition profile**
- **Maximum number of virtual adapters cannot be changed without de/reactivation**
- **Can be dynamically added or removed just like physical I/O slots**
  - Cannot be dynamically moved to another partition
  - Configuration of what is in the slot can be redefined without a restart of the partition
- **For AIX 5L, virtual I/O available at version 5.3 or later; direct I/O only at version 5.2**

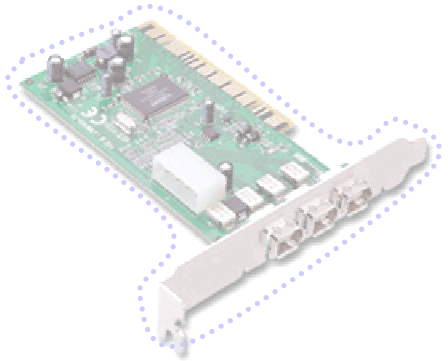

Virtual Ethernet

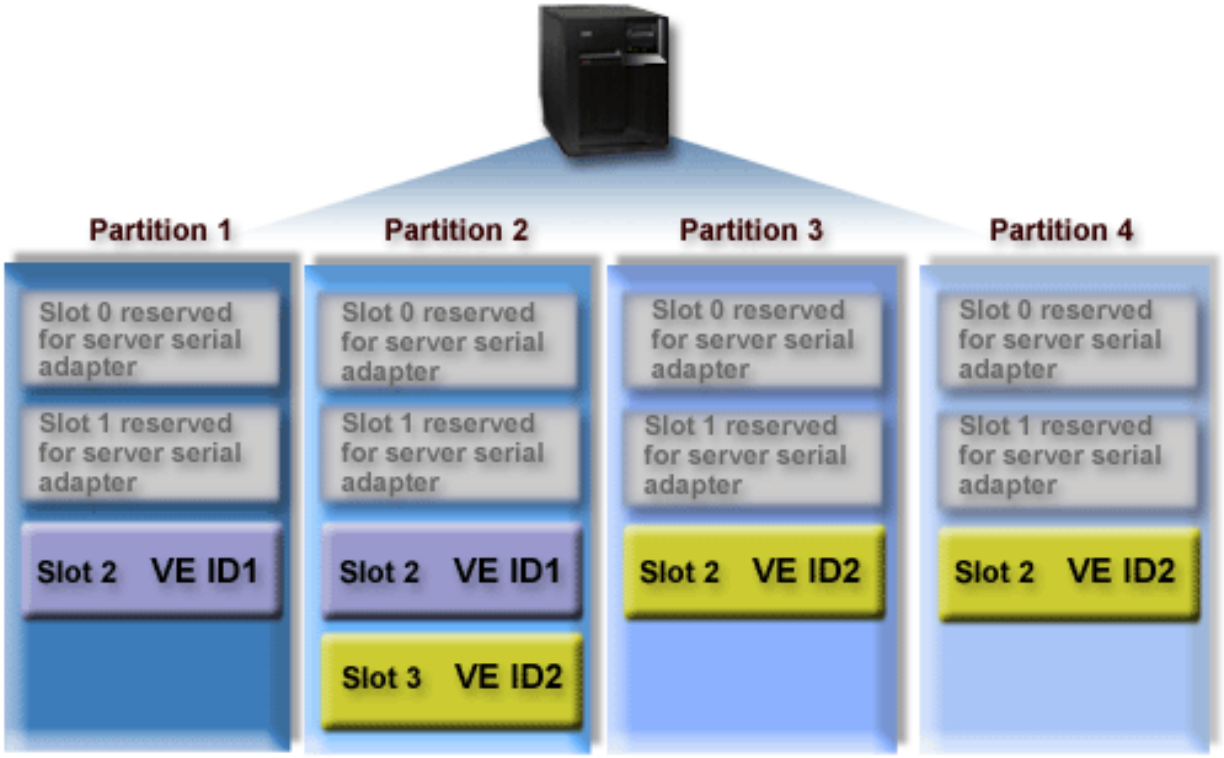

- Total number of VE LANs is 4094
- Two VE adapters that have the same PVID (port virtual ID) are on the same VLAN
- Above, P1 and P2 are on VLAN 1, P2, P3, and P4 are on VLAN 2
- Virtual slot numbers do not matter
- By using IEEE 802.1Q VE adapters, the same adapter can be on several VLANs (can have several VIDs)
- Only AIX and Linux can use IEEE 802.1Q adapters

Virtual Ethernet

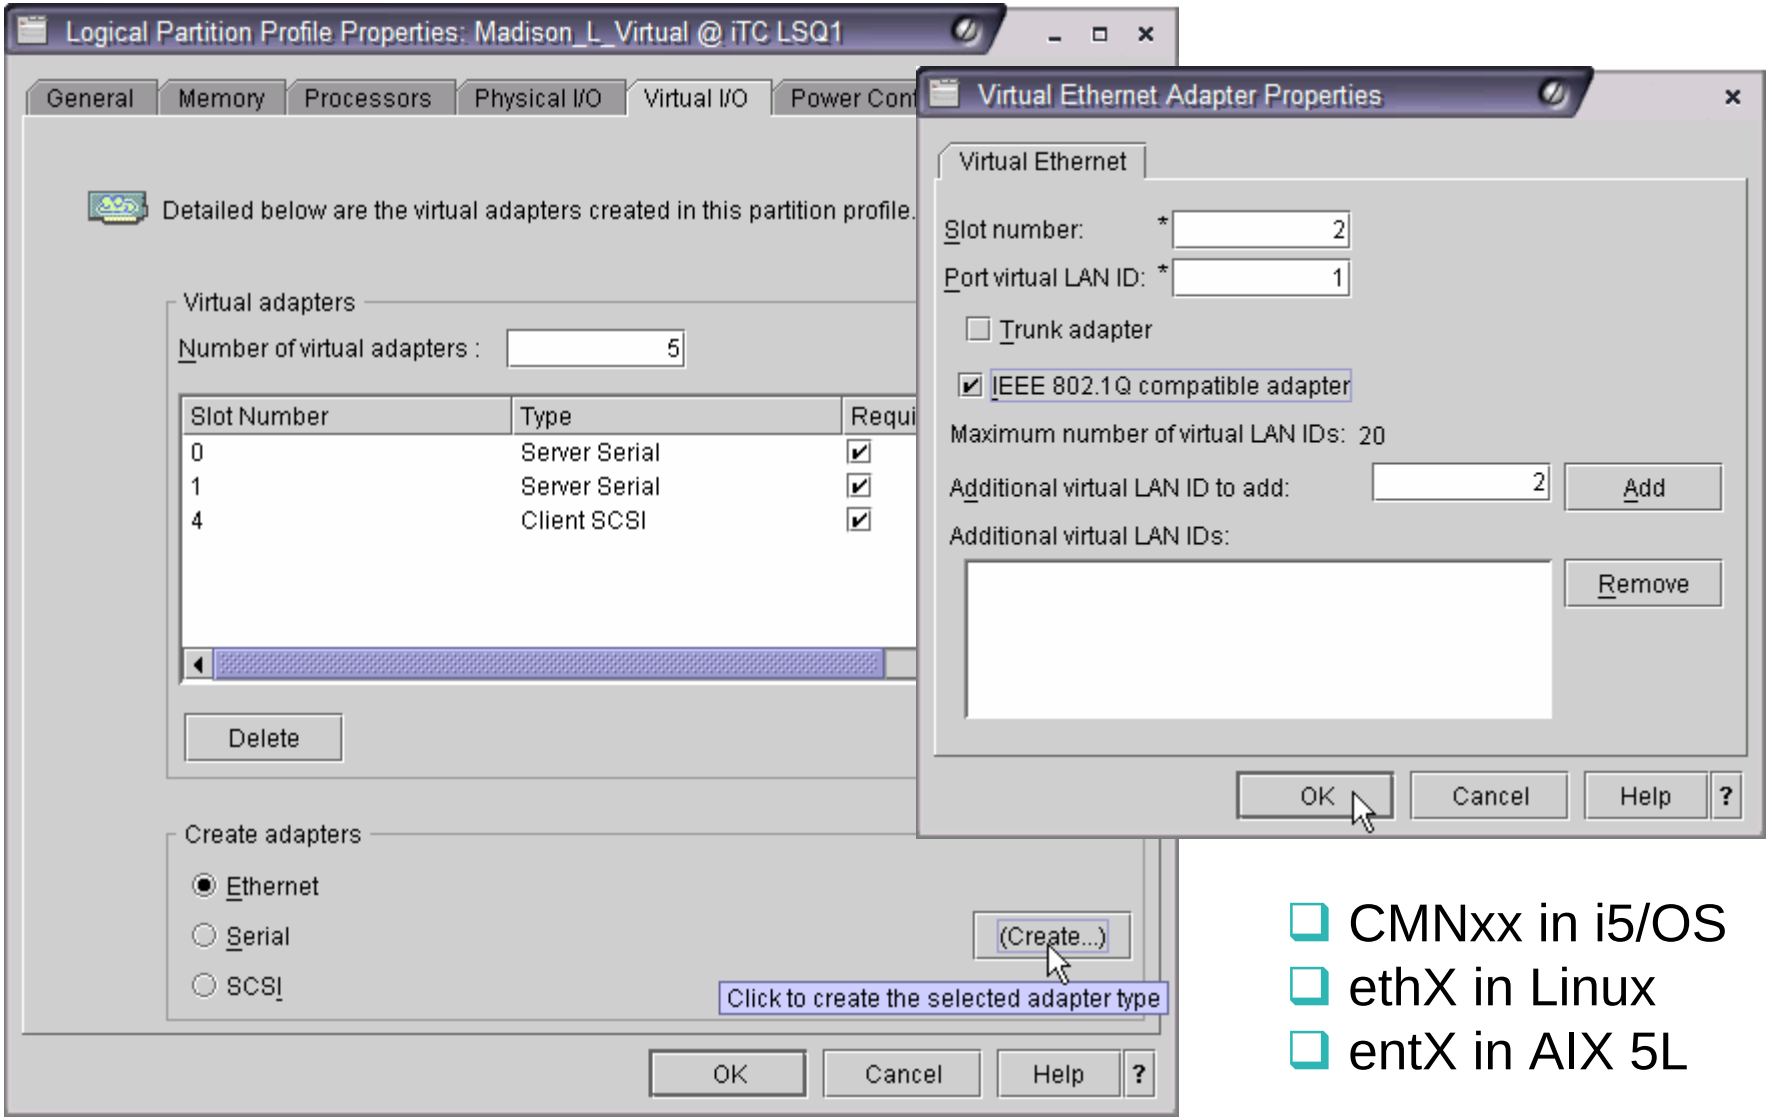

- CMNxx in i5/OS
- ethX in Linux
- entX in AIX 5L

Virtual SCSI

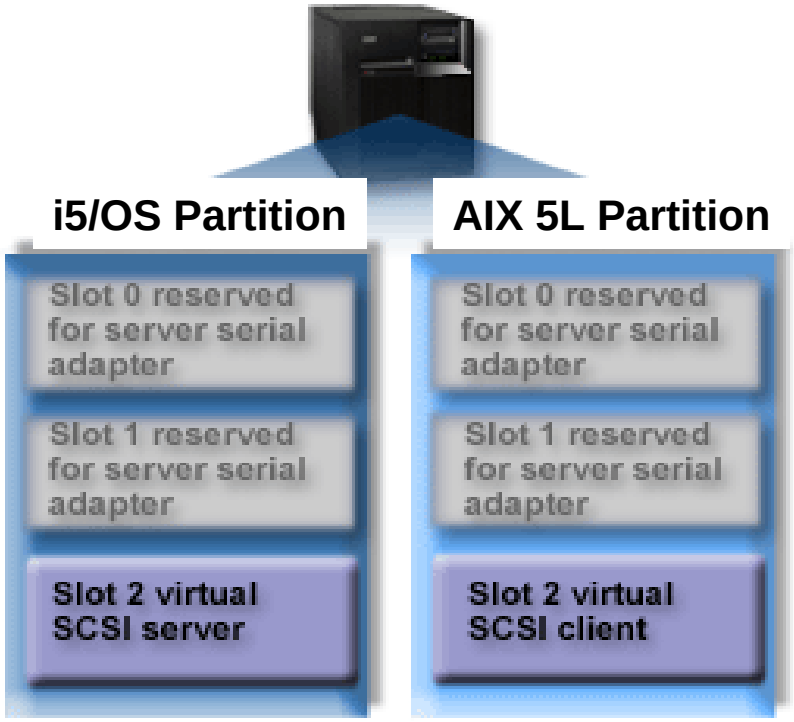

- Virtual Storage for AIX 5L on eServer p5 provided through Virtual I/O Server
- Virtual I/O Server not supported on eServer i5

- Virtual SCSI server and client adapters
- i5/OS is server, Linux, AIX 5L are clients
- Required for accessing Virtual Disk, CD/DVD, tape from i5/OS
- Virtual disk = NWSSTG object
- NWSSTG created in IFS
- NWSD object connects server-client SCSI adapter pair with NWSSTG
- One NWSD/multiple NWSSTG possible per server-client adapter pair
- AIX 5L, Linux see virtual disk as physical drive
- /dev/sdX in Linux
- hdiskX in AIX 5L
- Leverage RAID-5, multiple disk arms, scatter-loading, single-level storage
- No support for Virtual CD/DVD or tape for AIX 5L

Virtual SCSI

AIX 5L

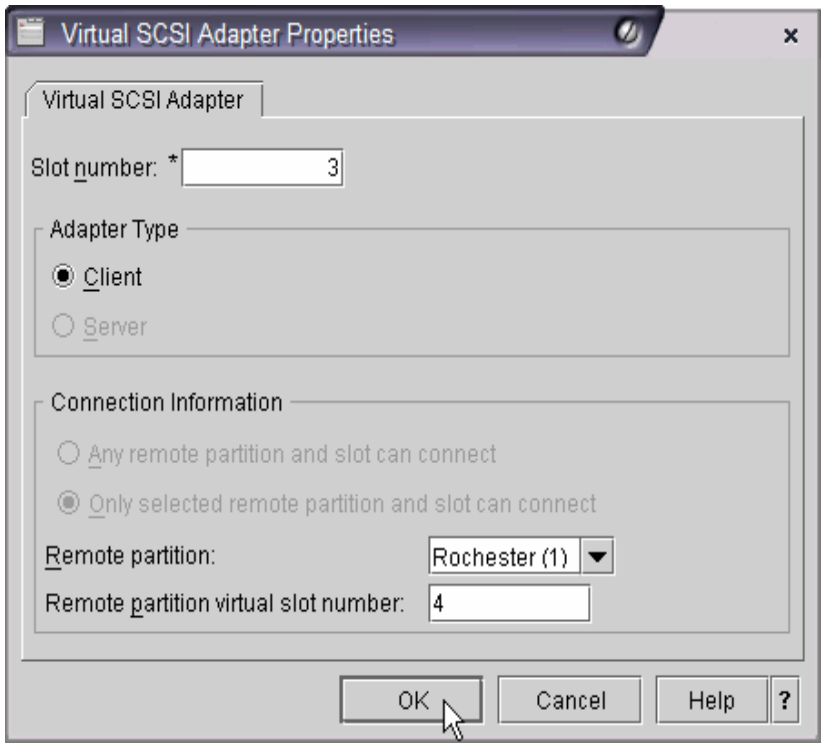

i5/OS

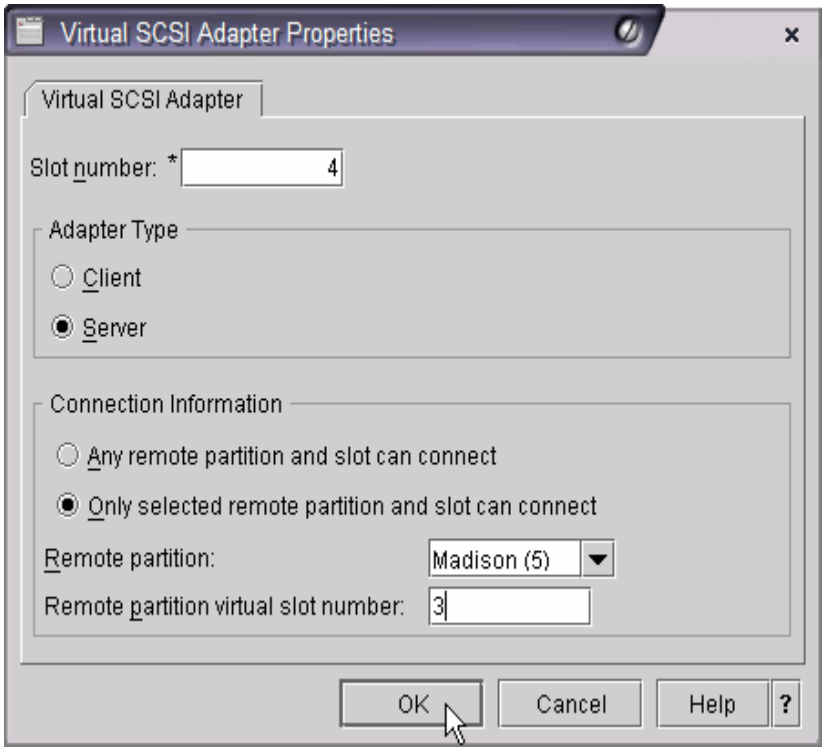

Individual slot numbers do no matter, as long as they are configured in pairs

Virtual SCSI: i5/OS View

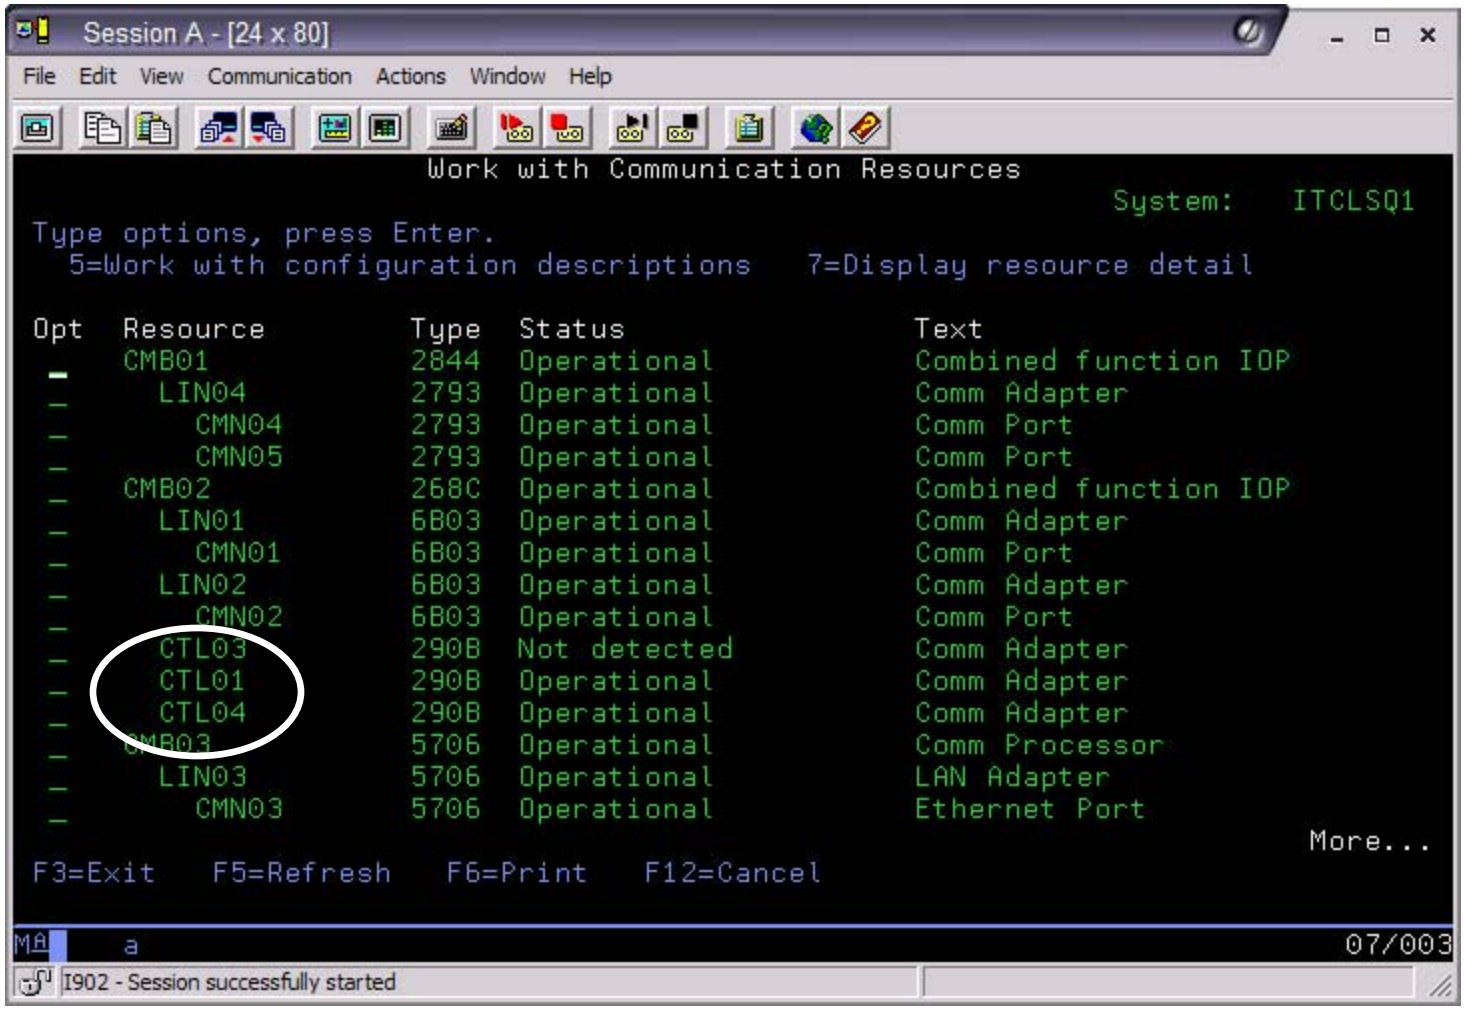

"Resource name" for AIX 5L NWSD

HA with Virtual SCSI

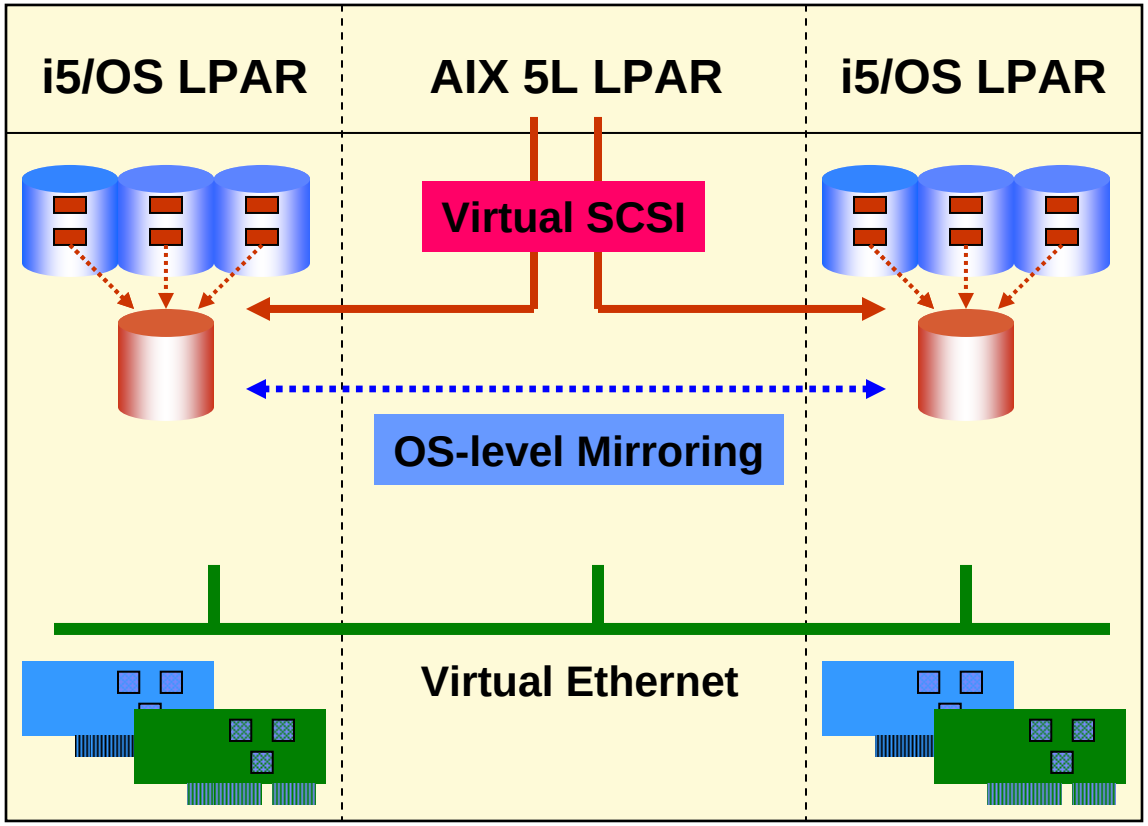

- AIX 5L partitions can access virtual disk from two or more i5/OS partitions
- Two virtual disks of equal size from separate i5/OS partitions allow mirroring of AIX 5L system disk
- Two Virtual SCSI server/client adapter pairs required, one for each i5/OS partition providing storage
- Mirroring accomplished with OS tools AIX 5L (**LVM**)
- AIX 5L partition becomes highly available, able to withstand failure of either host i5/OS partition

Virtual Serial

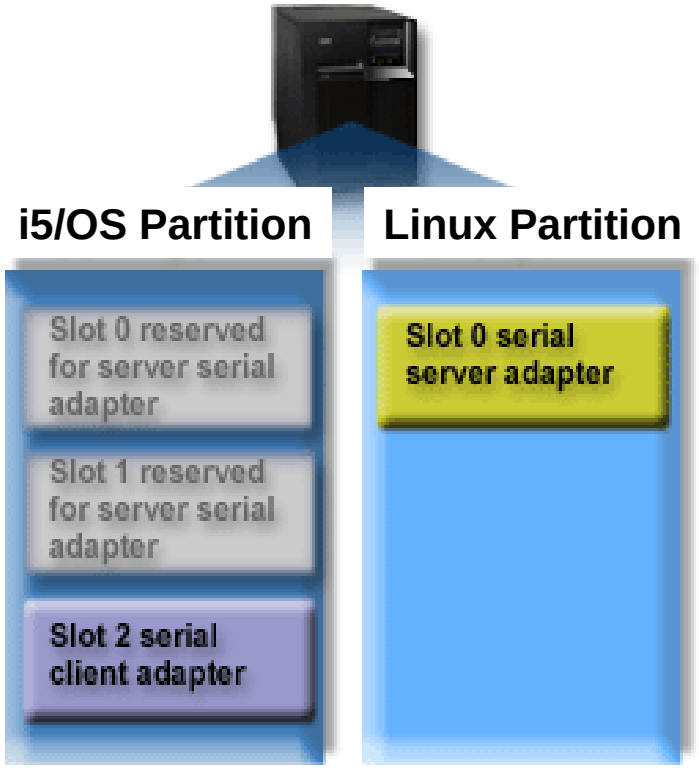

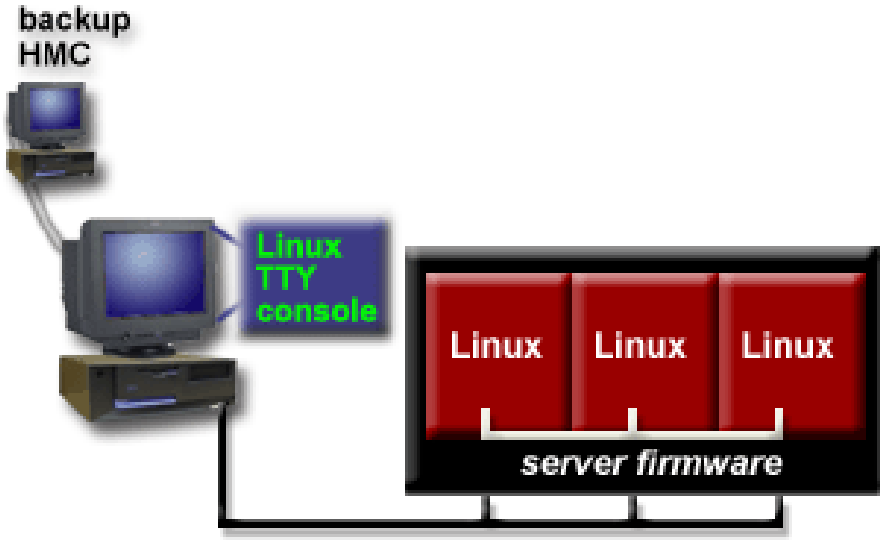

- First 2 virtual slots in every partition reserved for virtual serial server adapters for system console in HMC
- For i5/OS, virtual serial adapters provide 5250 console
- For AIX 5L, they provide character console

Virtual Serial

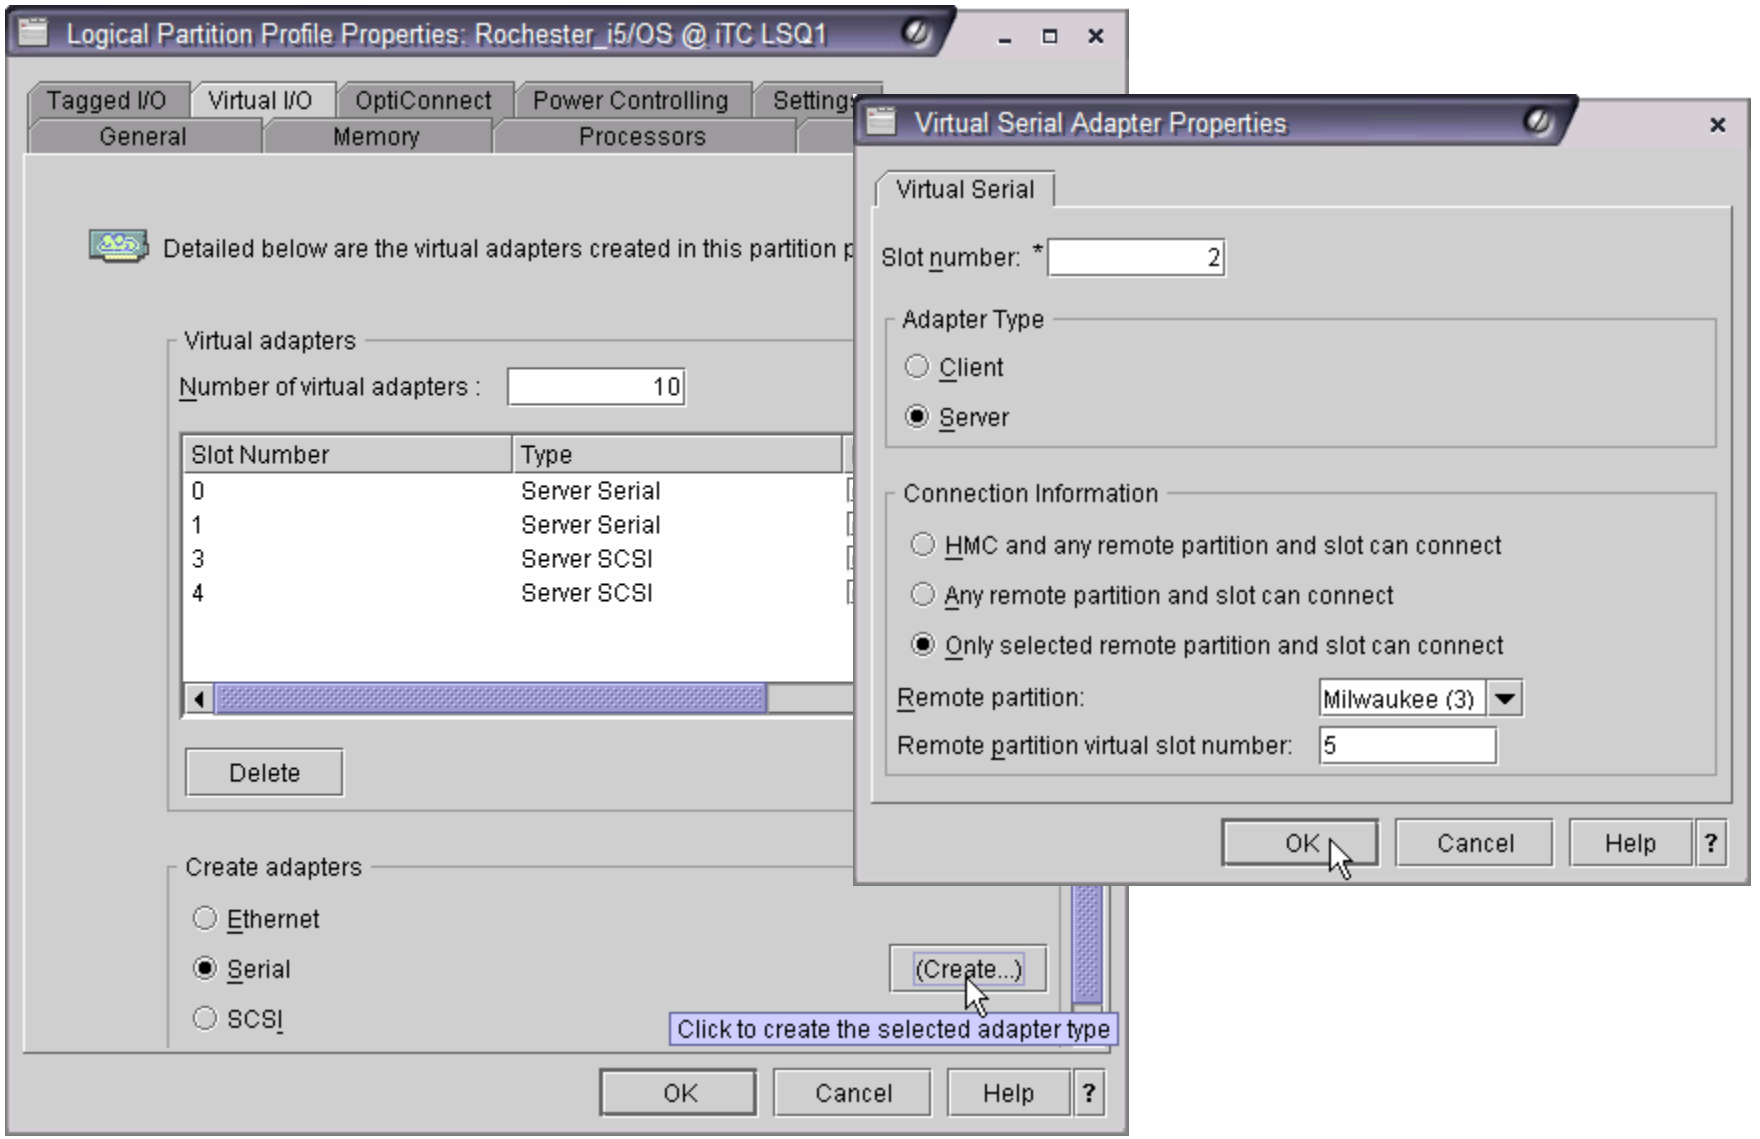

AIX 5L System Management: SMIT

- **System Management Interface Tool**
- **GUI or command-line**
- **Similar to iSeries Navigator**
- **Local or remote access**
- **Menu-driven**
- **Comprehensive management:**
  - Software installation
  - Hardware devices
  - Disk storage
  - Security and users
  - Networking
  - Jobs and subsystems
- **SMIT fast paths on command line:**
  - smit devices; smit tcpip
- **Similar to green-screen system management**
  - F4 to prompt fields, F3 to exit

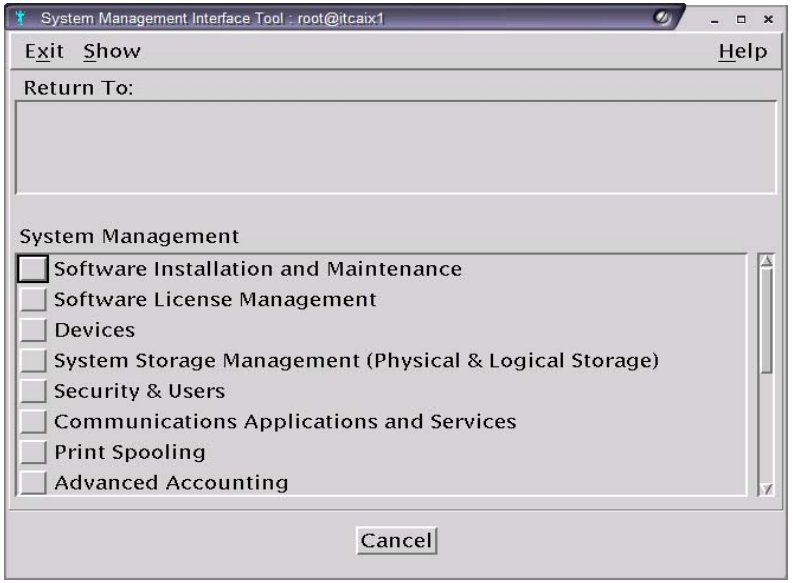

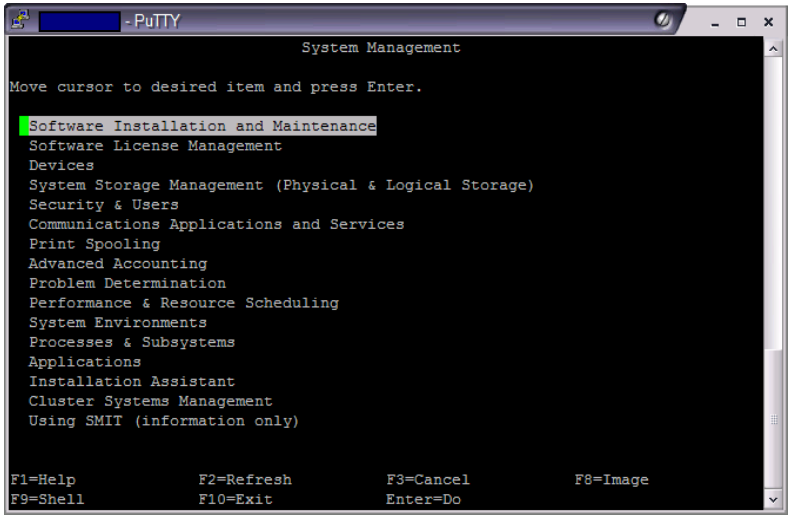

AIX 5L System Management: WebSM

- **Web-based System Manager**
- **Installs with AIX 5L**
- **Remote client connects to server running in AIX 5L**
- **Secure connection**
- **Similar to iSeries Navigator**
- **All aspects of system management**
- **Download WebSM client for Windows, Linux from your HMC!**
  http://*your_hmc*/remote_client.html
- **Common PC look and feel; interface similar to HMC**

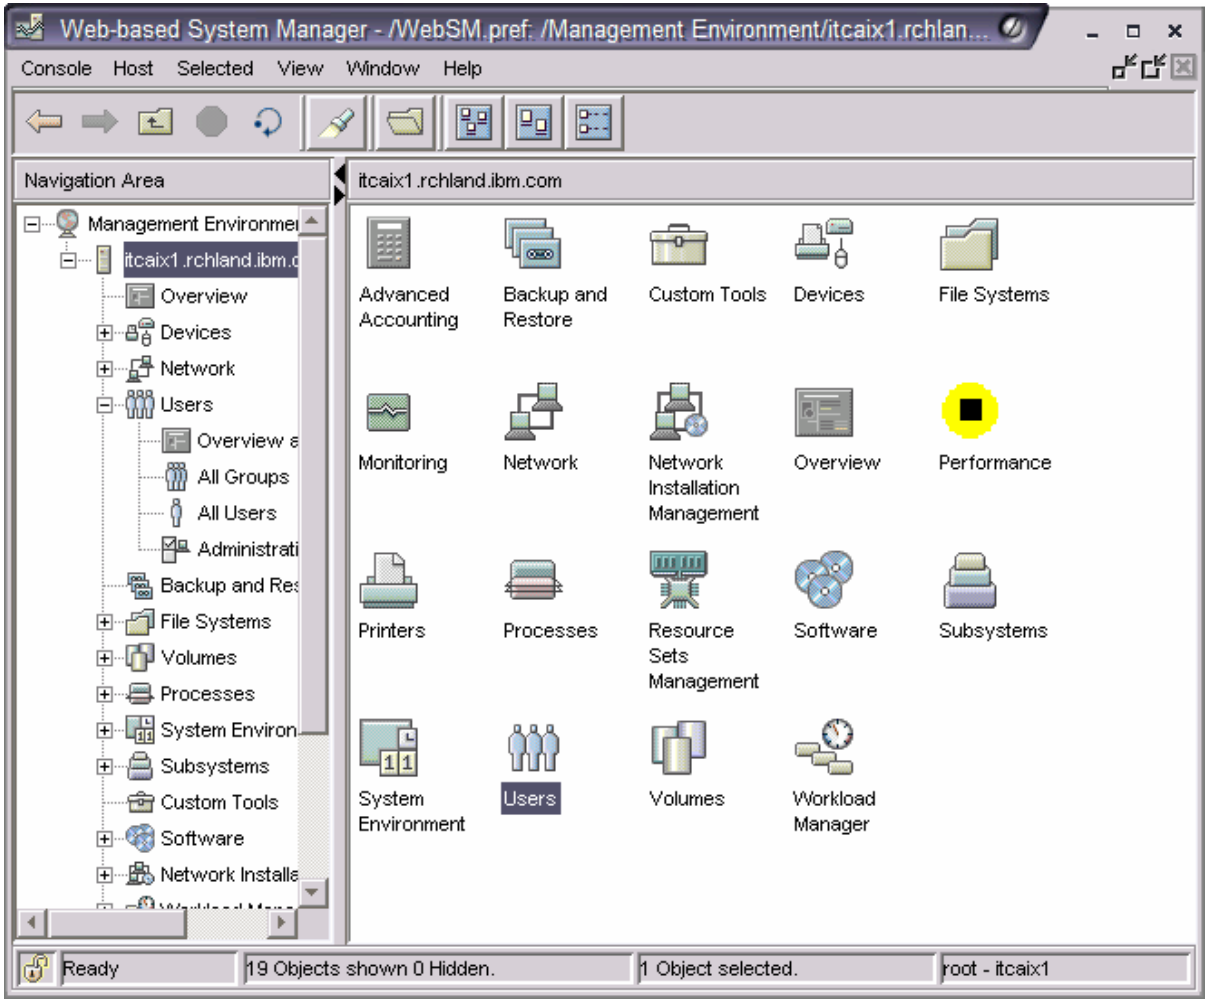

Sizing for AIX 5L on eServer i5

- **AIX 5L performance on eServer i5 is equivalent to eServer p5 with the same configuration**

Model	# of CPUs	GHz	AIX 5L 5.3 rPerf
p5-520	**2**	**1.65**	**9.86**
p5-570	**2**	**1.65**	**9.86**
p5-570	**4**	**1.65**	**19.66**
p5-570	**8**	**1.65**	**37.22**
p5-570	**12**	**1.65**	**53.43**
p5-570	**16**	**1.65**	**68.40**

With the announcement of eServer p5 systems and the availability of AIX 5L v5.2ML4 and AIX 5L v5.3, customers, business partners and independent software vendors can deploy applications running on AIX 5L v5.2 or AIX 5L v5.3 on either eServer i5 or eServer p5 systems.

Since eServer i5 and eServer p5 are built with the same POWER5 processors and server technology, AIX 5L applications that run on eServer p5 can run on eServer i5 unchanged.

While most of the I/O options are the same between the i5 & p5, support is dependent upon specific I/O availability. You should verify that adapters running in AIX partitions on eServer i5 servers are utilizing I/O Adapters supported by AIX 5L.

Some selected AIX licensed program products are not available on eServer i5.

Notes:

1) The rPerf (Relative Performance) is an estimate of commercial processing performance. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and maximum memory available. However, the model does not simulate disk or network I/O operations. IBM pSeries and IBM RS/6000 Performance Report provides rPerf values for servers other than POWER5 servers -- http://www-1.ibm.com/servers/eserver/pseries/hardware/system_perf.html

2) The rPerf results of IBM eServer p5 systems are based on AIX 5L V5.3 which supports simultaneous multi-threading. Simultaneous multi-threading results in a 30% boost in rPerf. For purposes of estimating rPerf for AIX 5L V5.2 on eServer p5 systems divide the published rPerf number by 1.3 and round down to the nearest 1/100th

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

Sizing for AIX 5L on eServer i5

- Sizing recommendations for AIX 5L applications can come from ISVs or IBM Software Group
  - Guidelines for several ISV applications available at IBM PartnerWorld for Developers site: http://www.developer.ibm.com/welcome/eserver/e3/CSFServlet?mvcid=main&packageid=3002
  - Redbook on pSeries sizing available: http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg247071.html?Open
- Determine requirements for CPU (rPerf), memory, disk and network I/O on eServer p5
  - Memory, disk and network requirements will be the same on eServer i5
- Factor in logical partitioning (LPAR) considerations:
  - rPerf numbers produced on non-LPAR machine. AIX 5L on eServer i5 will always run on LPAR machine
  - If AIX 5L partition will use shared processors, factor in between 1 and 10% overhead, depending on total number of partitions using shared processors
  - rPerf numbers not available for single-CPU or sub-processor AIX 5L partitions. Use linear interpolation of existing rPerf numbers to estimate rPerf number for AIX 5L partition using less than full CPU.
- Proceed by identifying server model and CPU capacity
  - Use rPerf numbers from p5 system, factor in LPAR overhead
  - Derive estimate using Workload Estimator – use “Generic (AIX) Workload”

Sizing for AIX 5L on eServer i5

- Use eServer Workload Estimator (WLE): http://www-912.ibm.com/wle/EstimatorServlet
- Configure “Generic (AIX) Workload”
- Enter adjusted rPerf required for AIX 5L workload. Round up.
- Enter amount of memory
- derived from eServer p5
- Enter disk configuration derived from eServer p5
- Choose “none” for “Additional Characteristics
- Choose DBCS support as appropriate
- Create other i5/OS workloads to arrive at consolidated system*

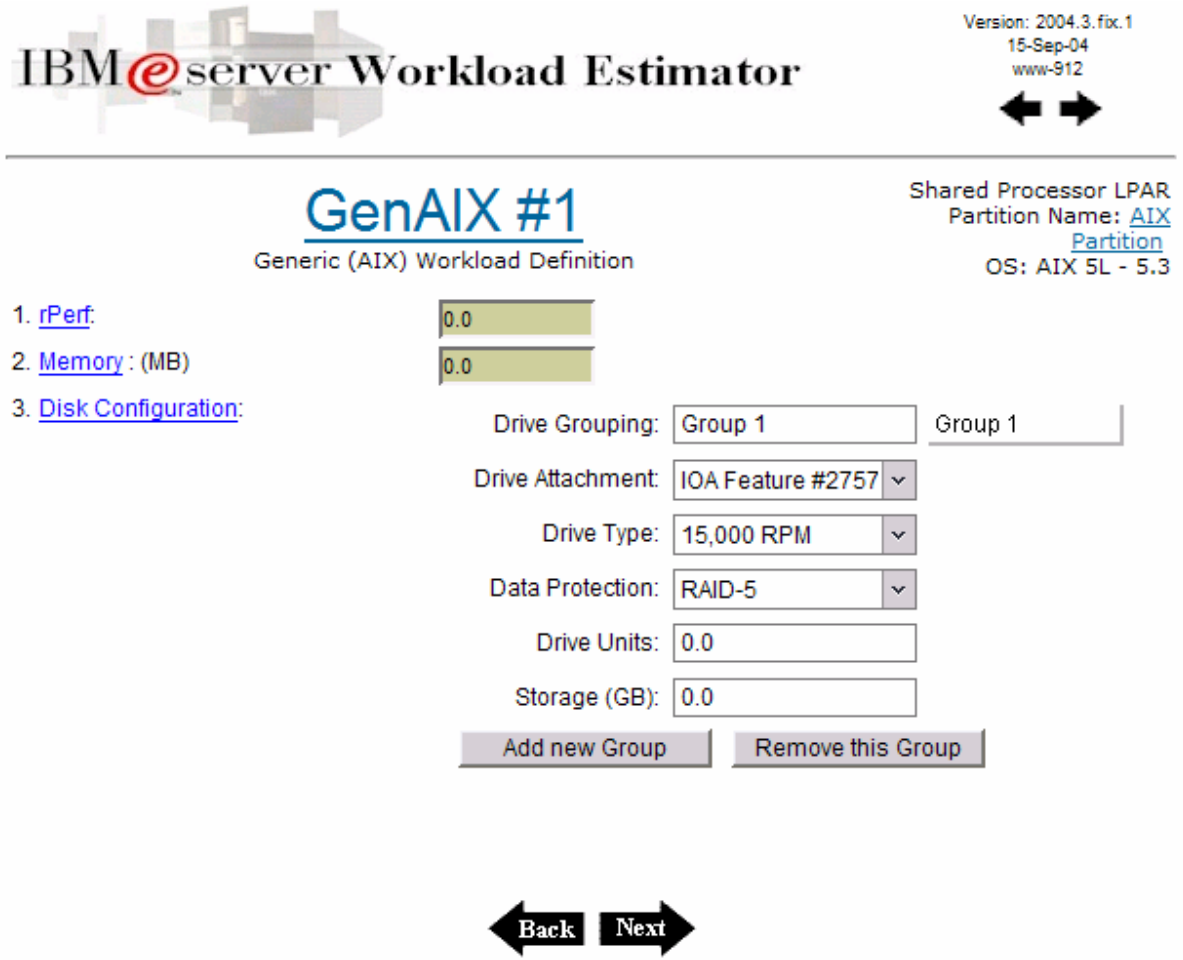

Performance Management: topas

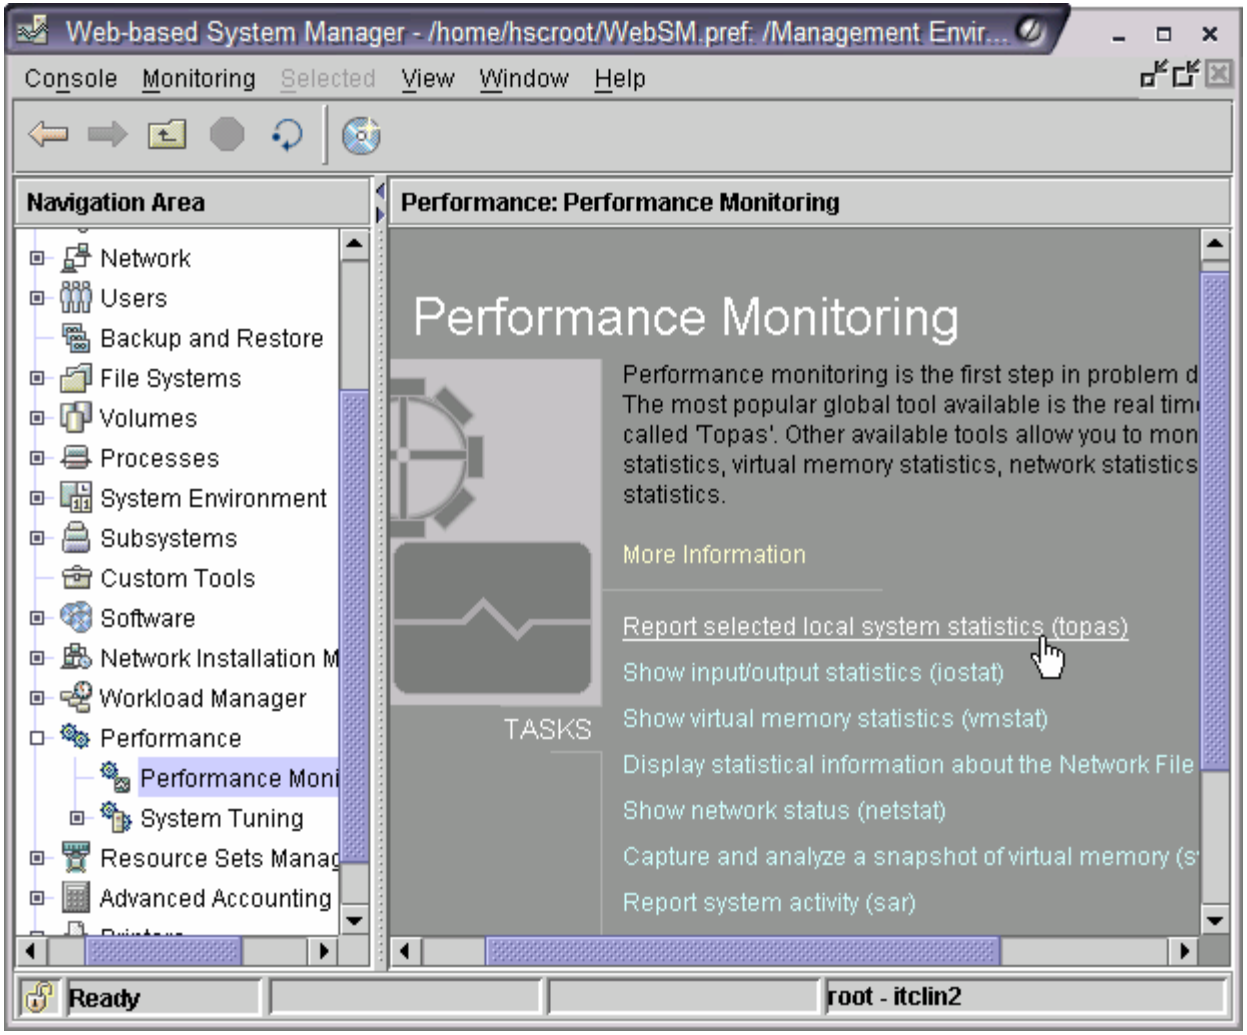

- topas provides snapshot of system activity:
  - CPU utilization
  - Memory
  - Processes
  - Disk I/O
  - Network I/O
  - Paging
- Similar to WRKSYSSTS
- In WebSM: click on first "Performance Monitoring" option to invoke topas

Performance Management: topas

CPU usage

Number of network packets, inbound and outbound

% disk busy

"High offenders" by CPU and paging space usage

Paging statistics

Memory usage for different types of tasks

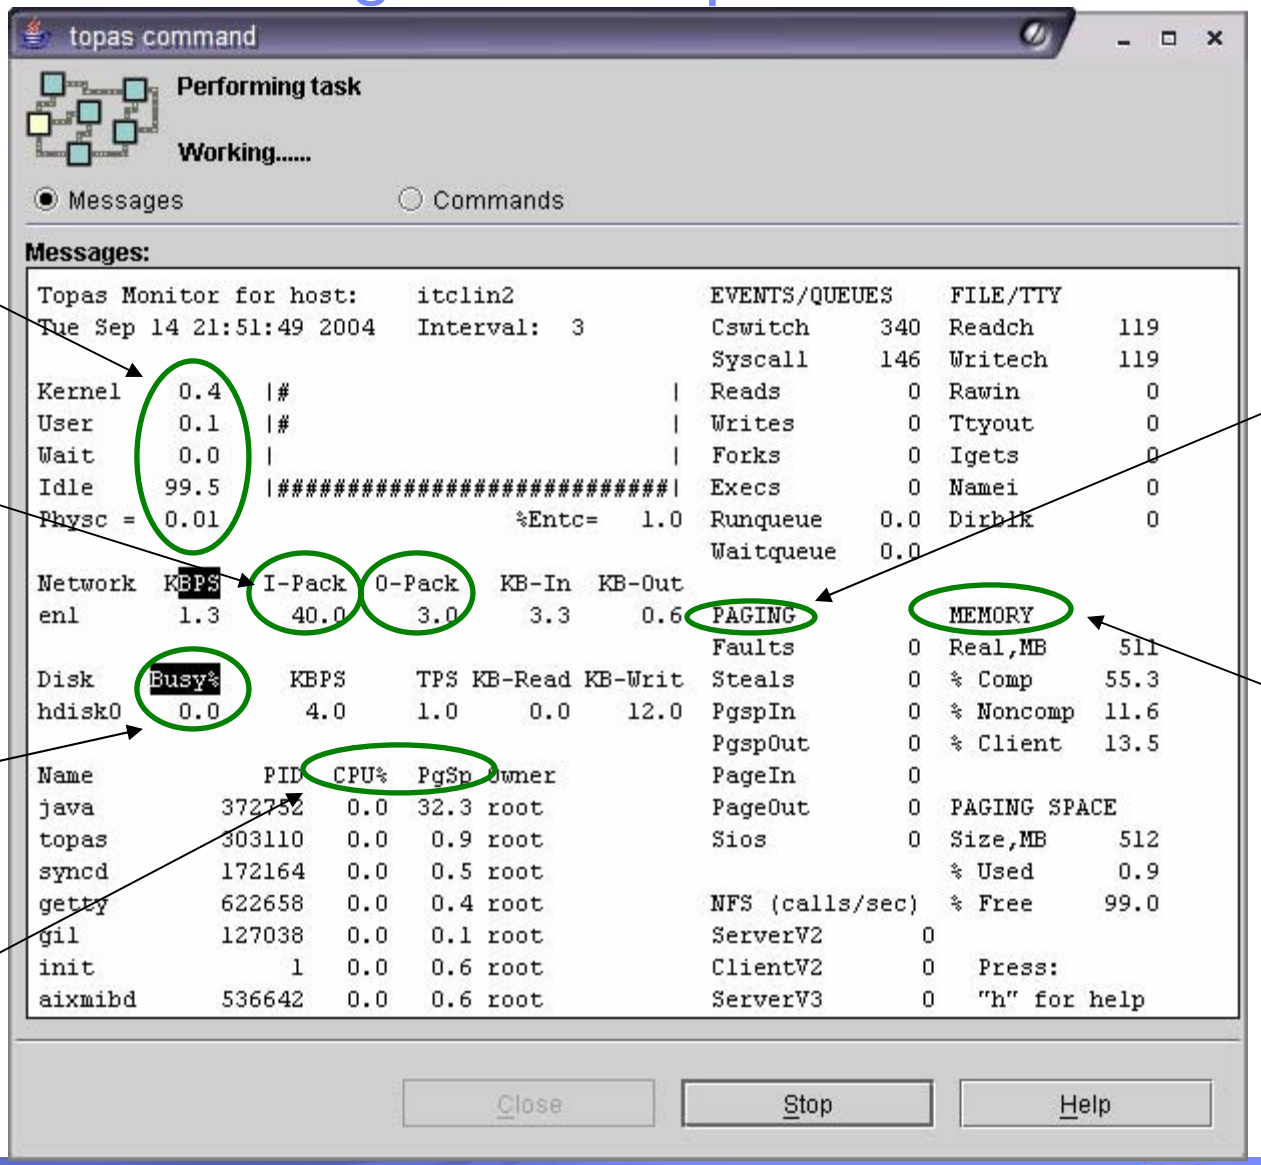

Performance Management: iostat

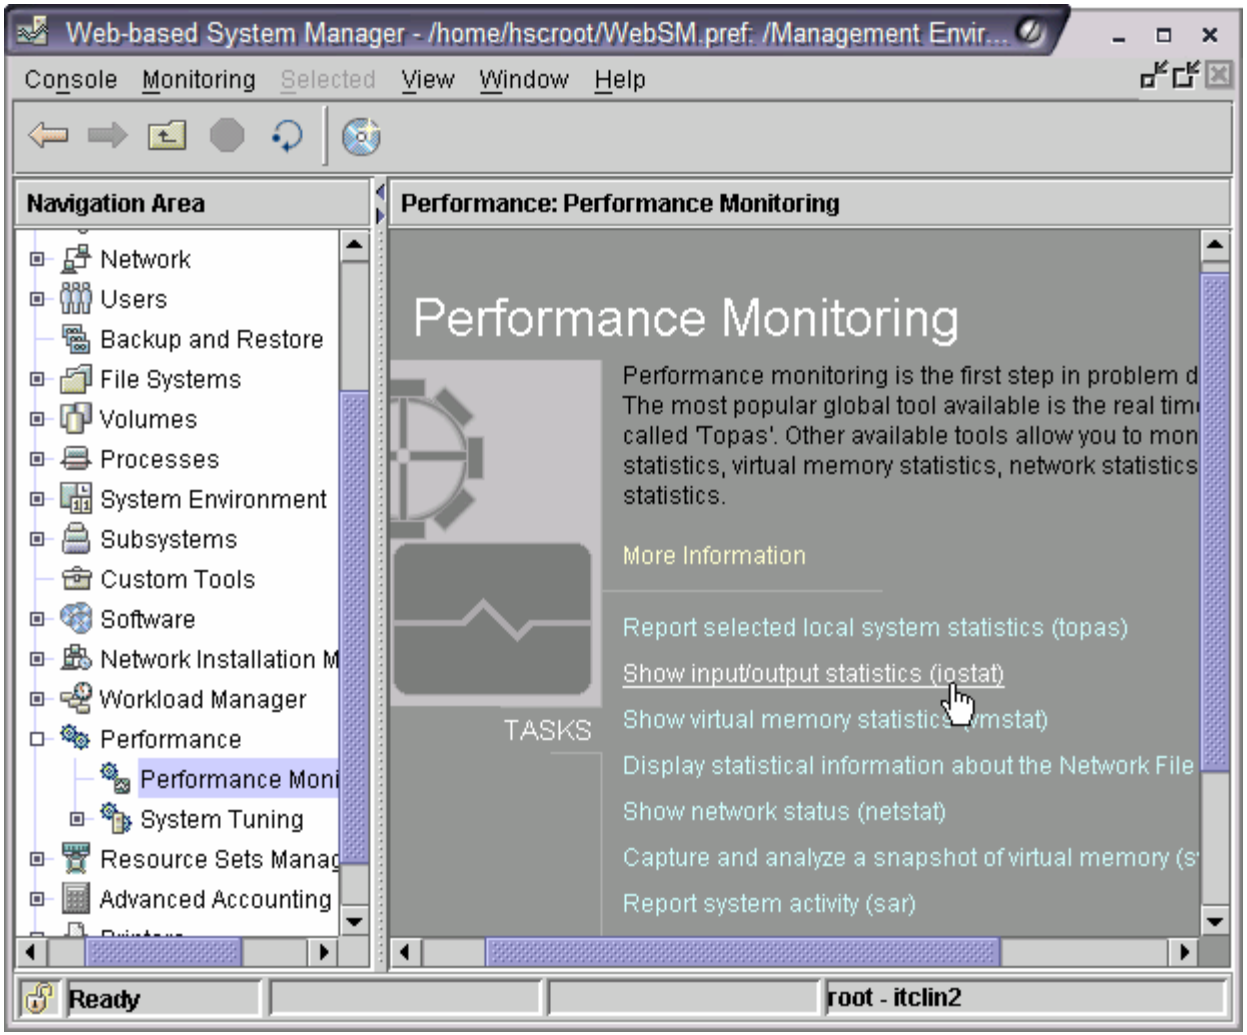

- iostat displays disk and CPU activity
- Most often used for disk utilization monitoring
- Click on second menu option to start it

Performance Management: iostat

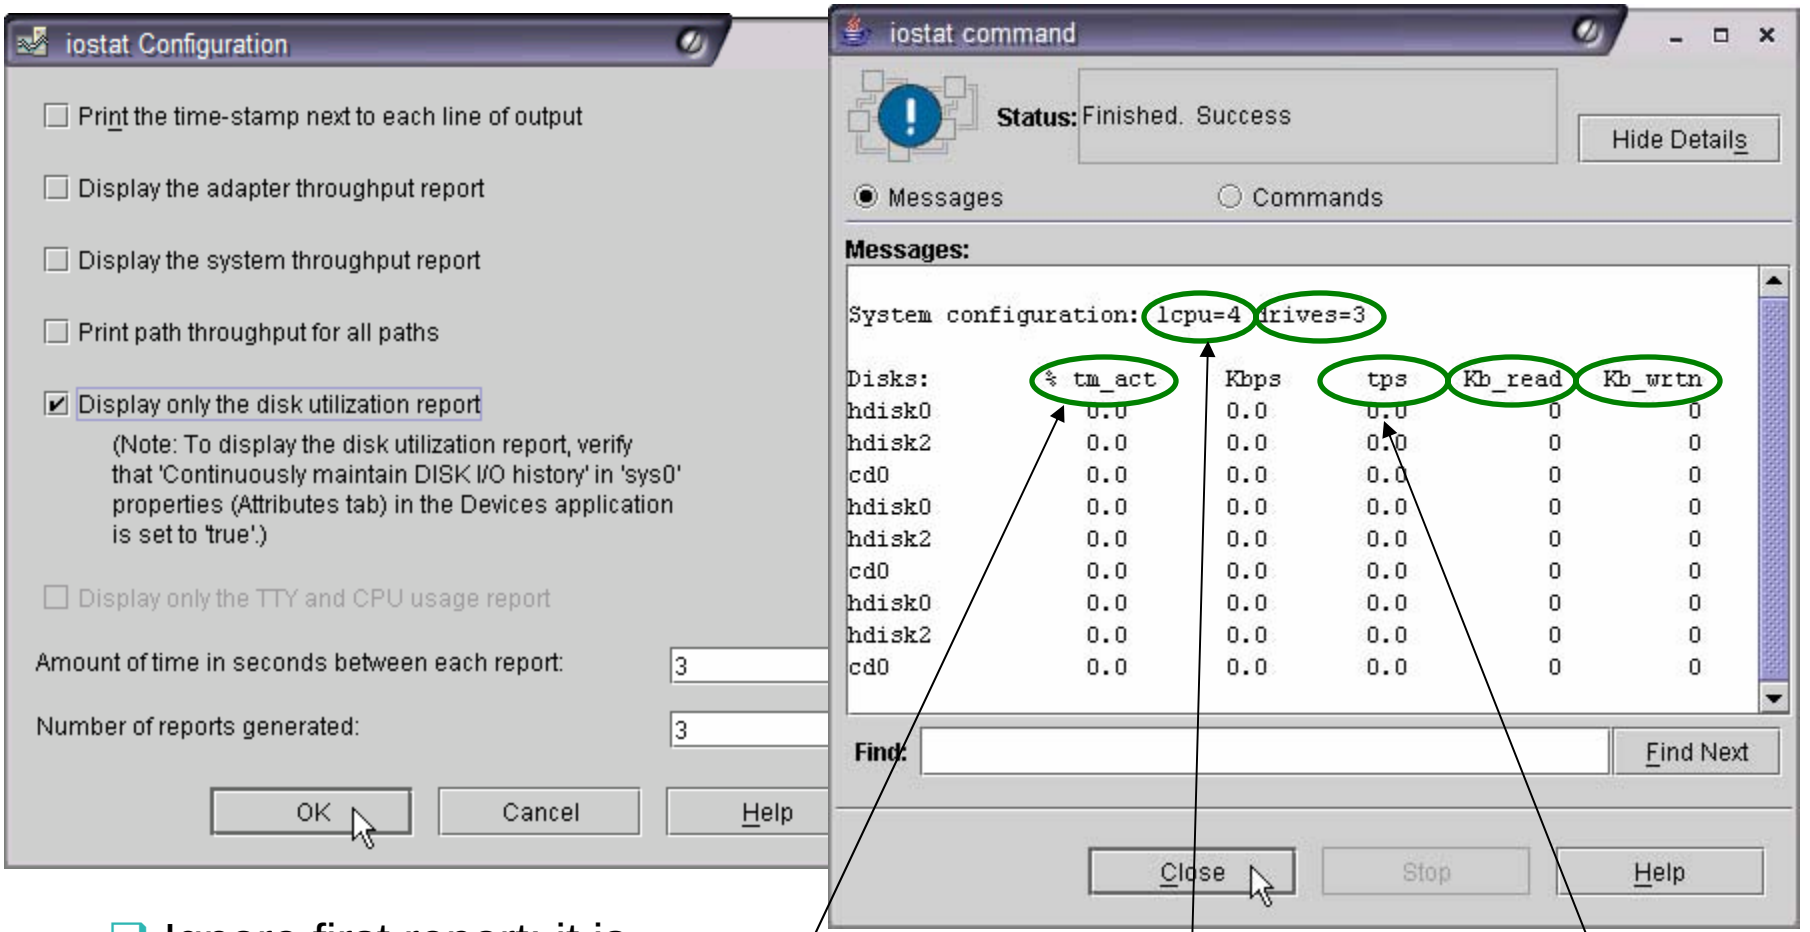

- Ignore first report; it is average of disk activity since boot

% busy

Number of logical CPUs

Transfers (r/w) per second

Performance Management: vmstat

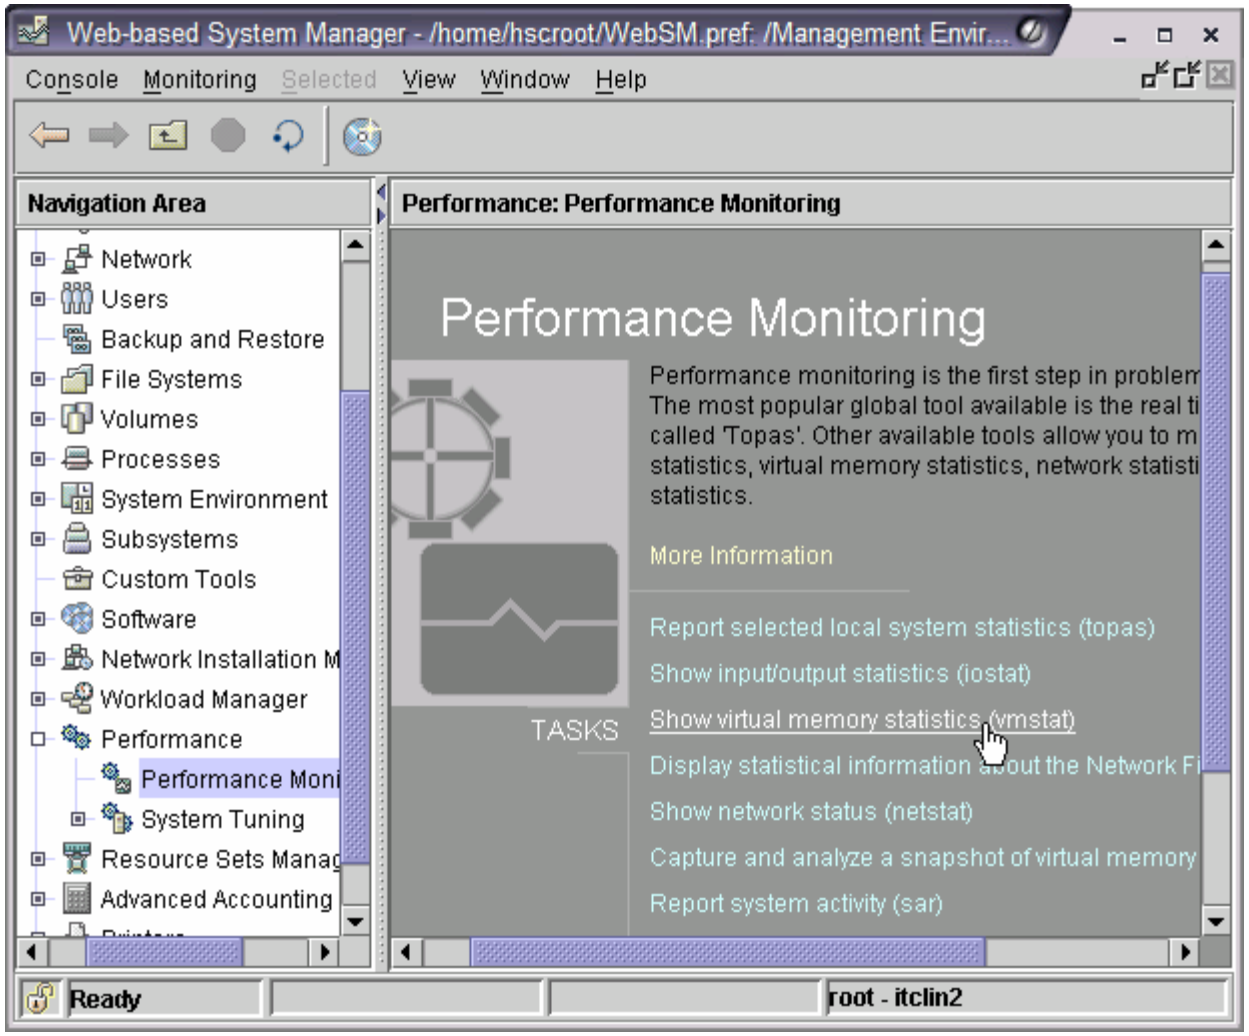

- To display memory utilization, use vmstat

Performance Management: vmstat

Number of active (used) pages

Number of free pages

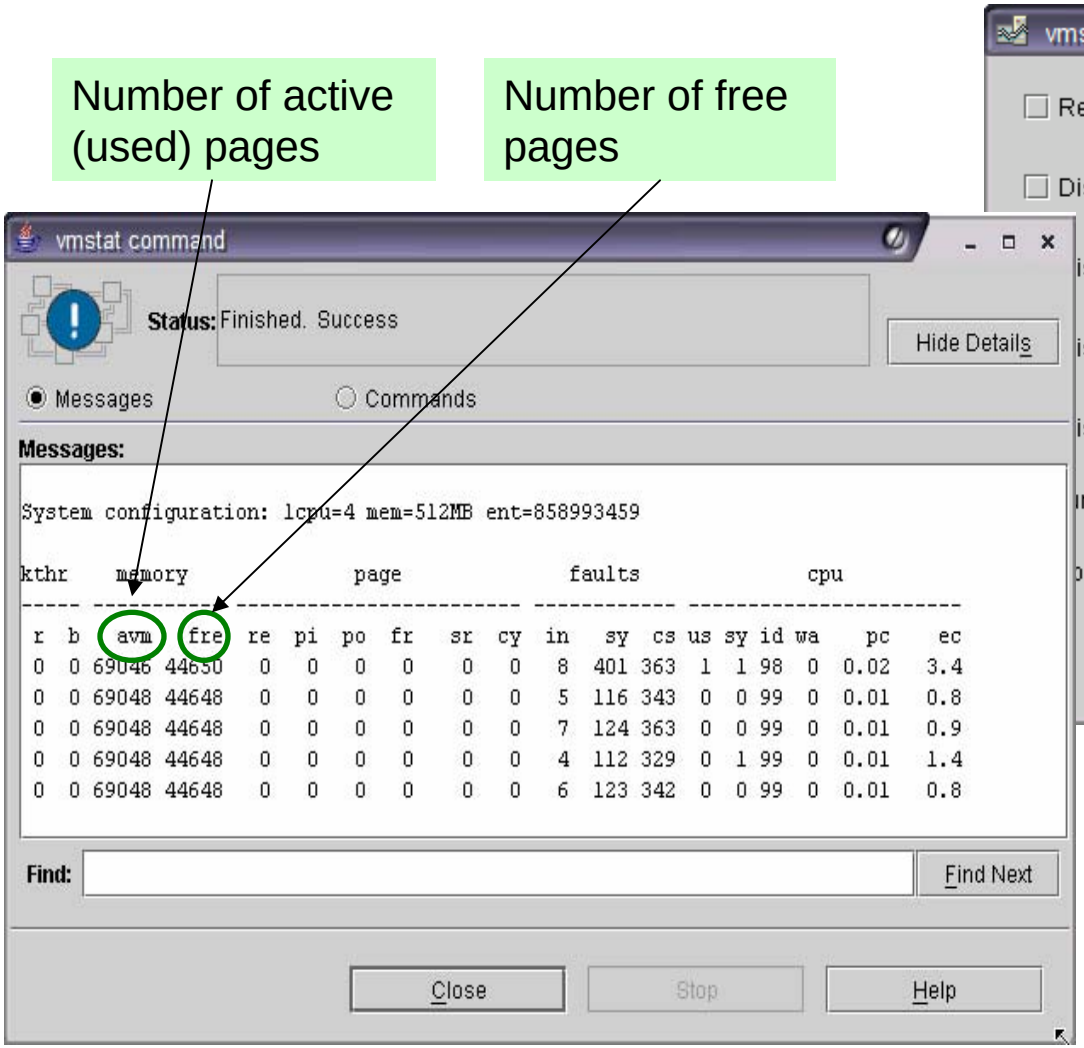

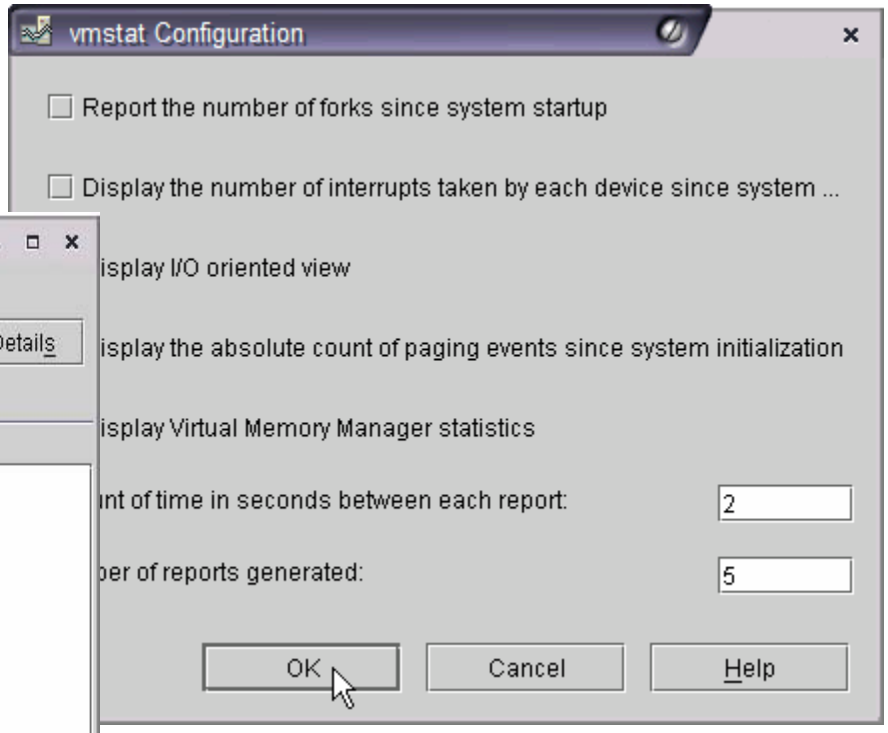

- Portion of real memory utilized for file system caching
- Therefore, low number of free pages is usually not cause for alarm

Performance Management: netstat

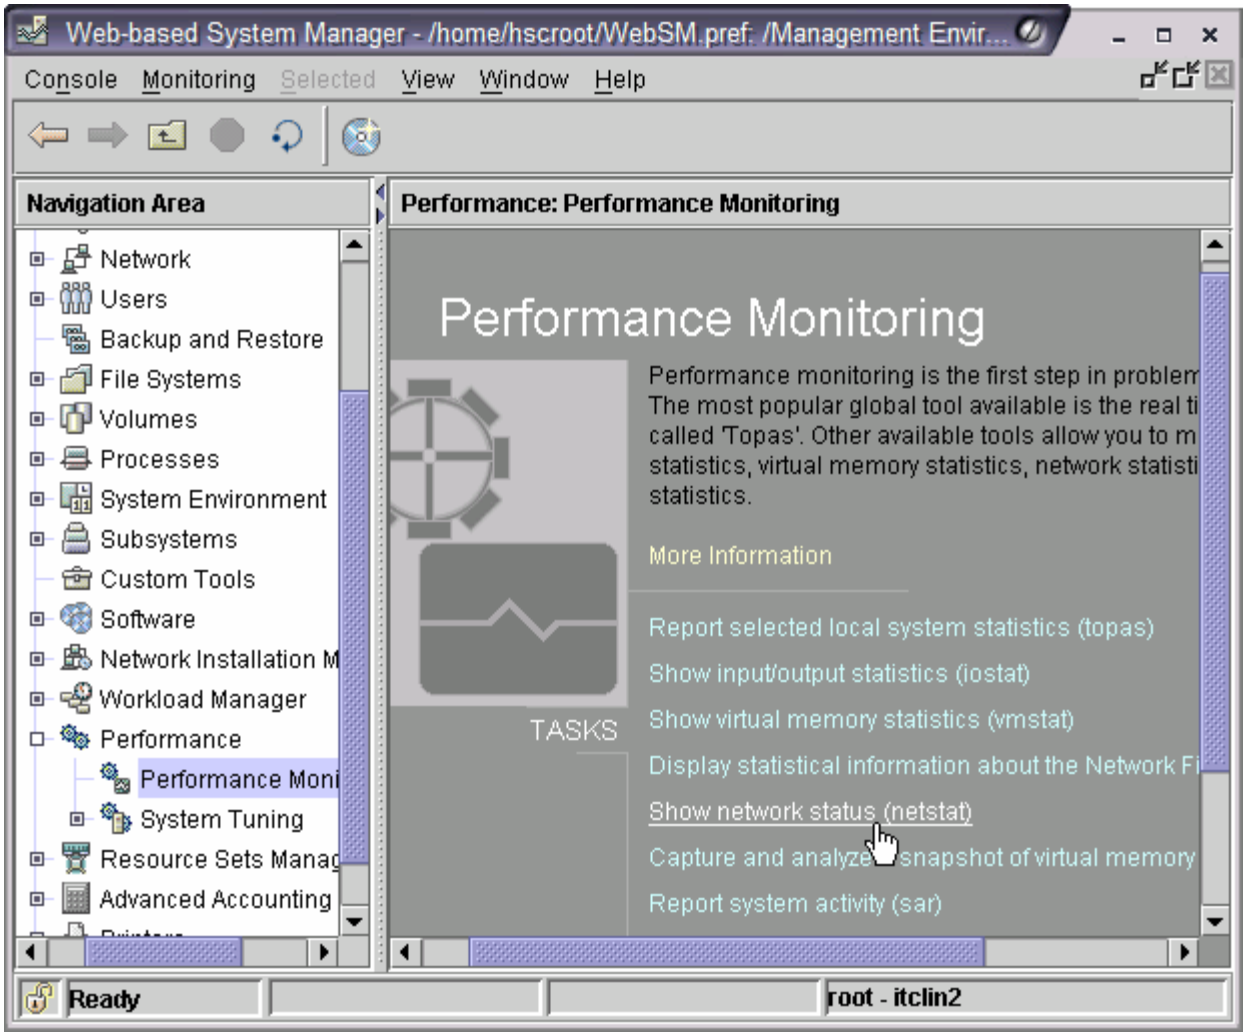

- To network information, use netstat

Performance Management: netstat

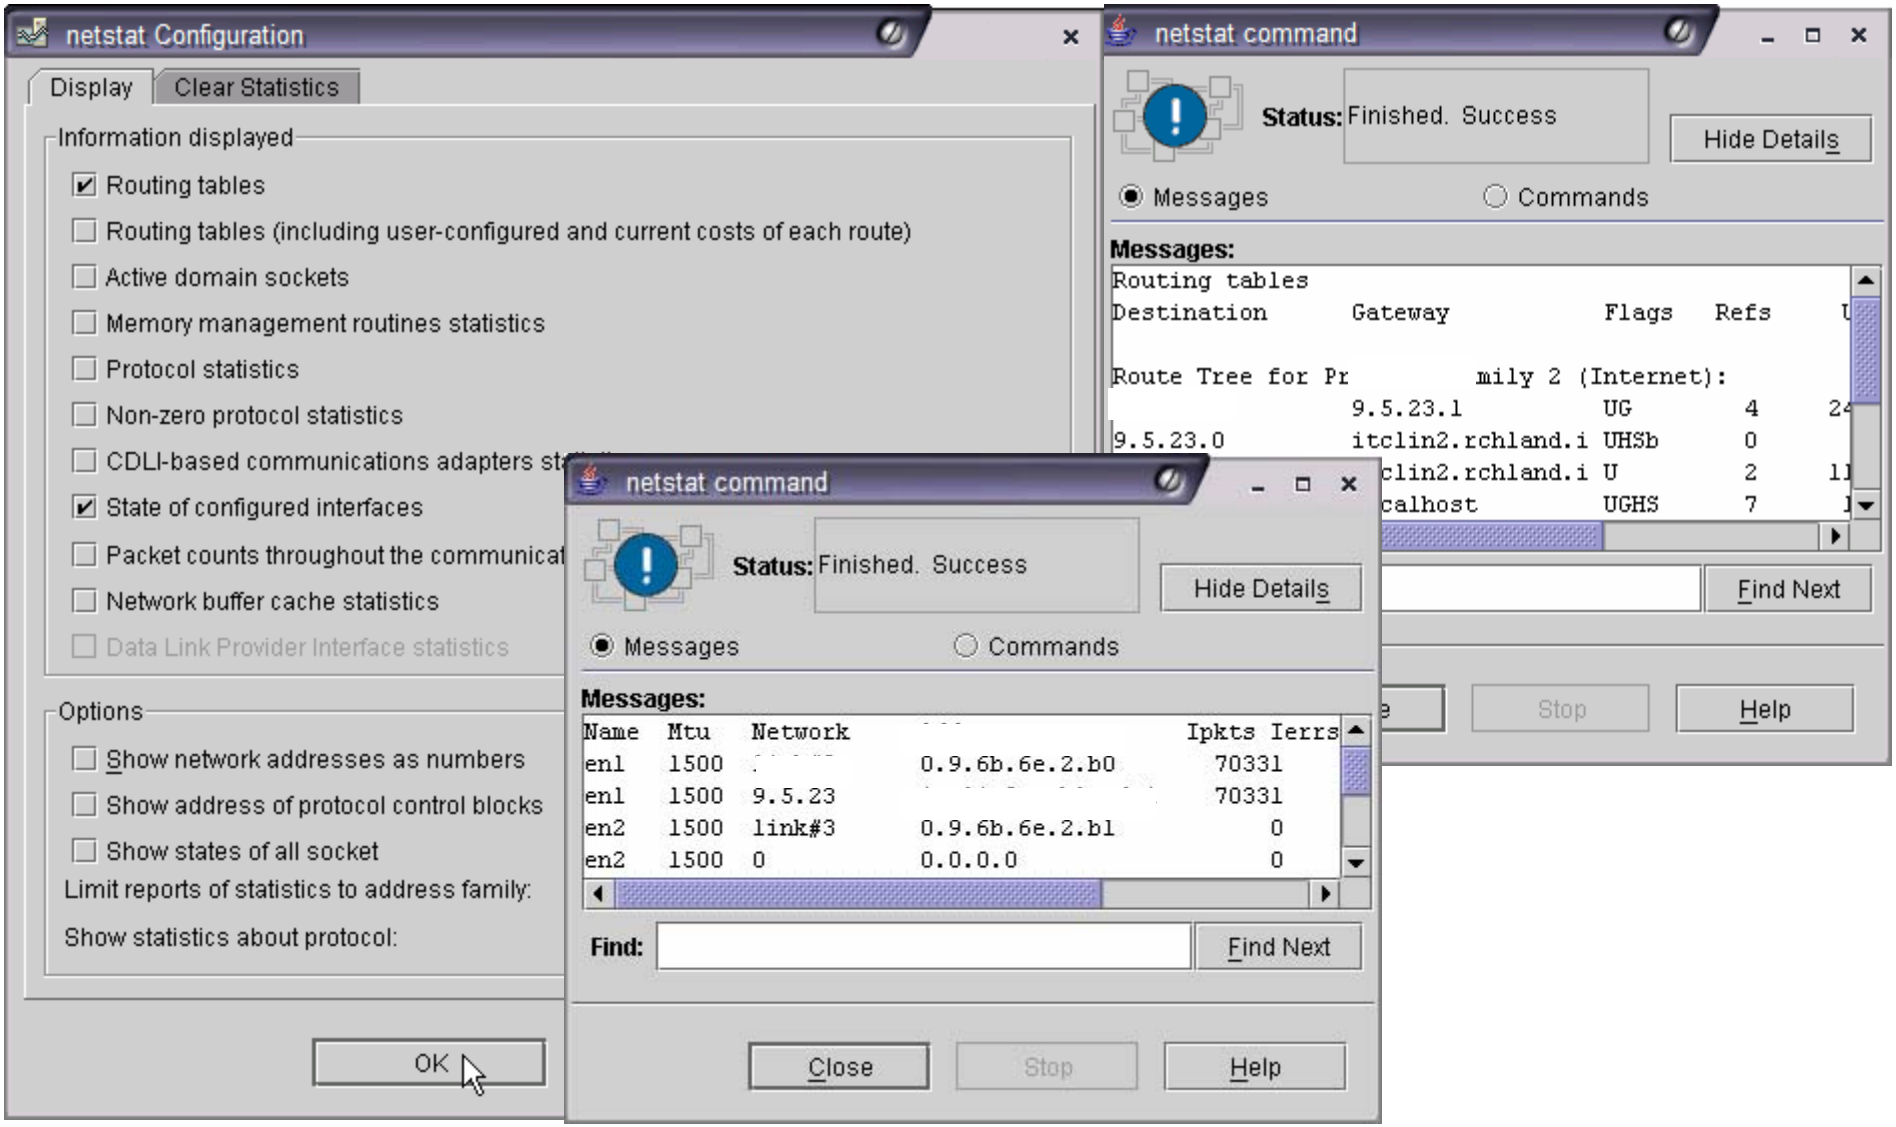

Performance Management: lparstat

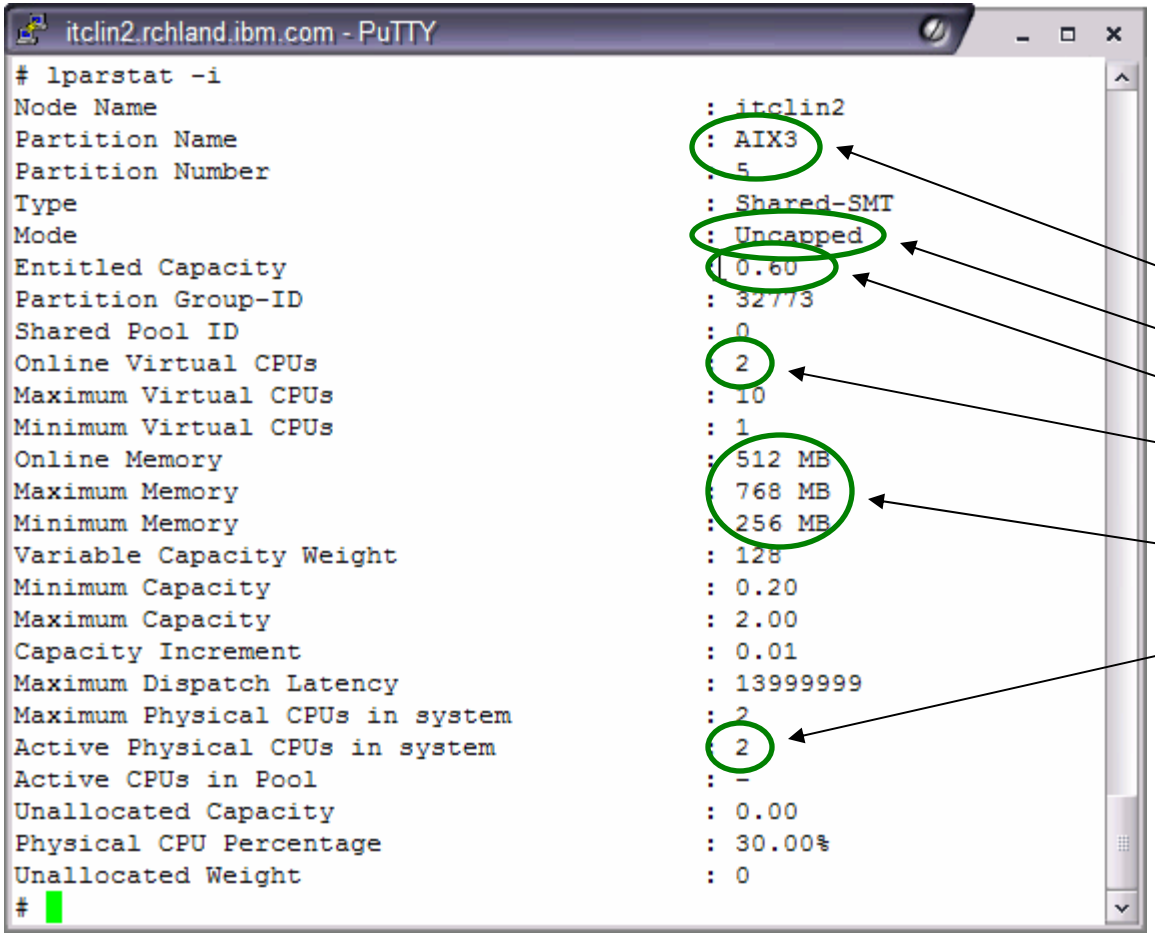

- lparstat displays logical partitioning and system resource allocation information:
  - Partition name and ID
  - Capped vs. uncapped
  - Shared processor units
  - Number of virtual CPUs
  - Memory settings
  - Number of physical CPUs in system

AIX 5L Performance Toolbox (PTX)

- What is it?
  - Separate LPP
  - Comprehensive GUI system monitoring tool
  - More information at http://www-1.ibm.com/servers/aix/products/ibmsw/system_man/perftoolbox.html
- What can it do?
  - Provides a quick and simple solution for obtaining and analyzing detailed system information
  - Generates reports on system activity over hours, days or weeks
  - 24 x 7 analysis of large performance datasets
  - Supports distributed performance monitoring
  - Provides the ability to customize views for monitoring
  - Provides access to thousands of system metrics

AIX 5L Performance Toolbox (PTX)

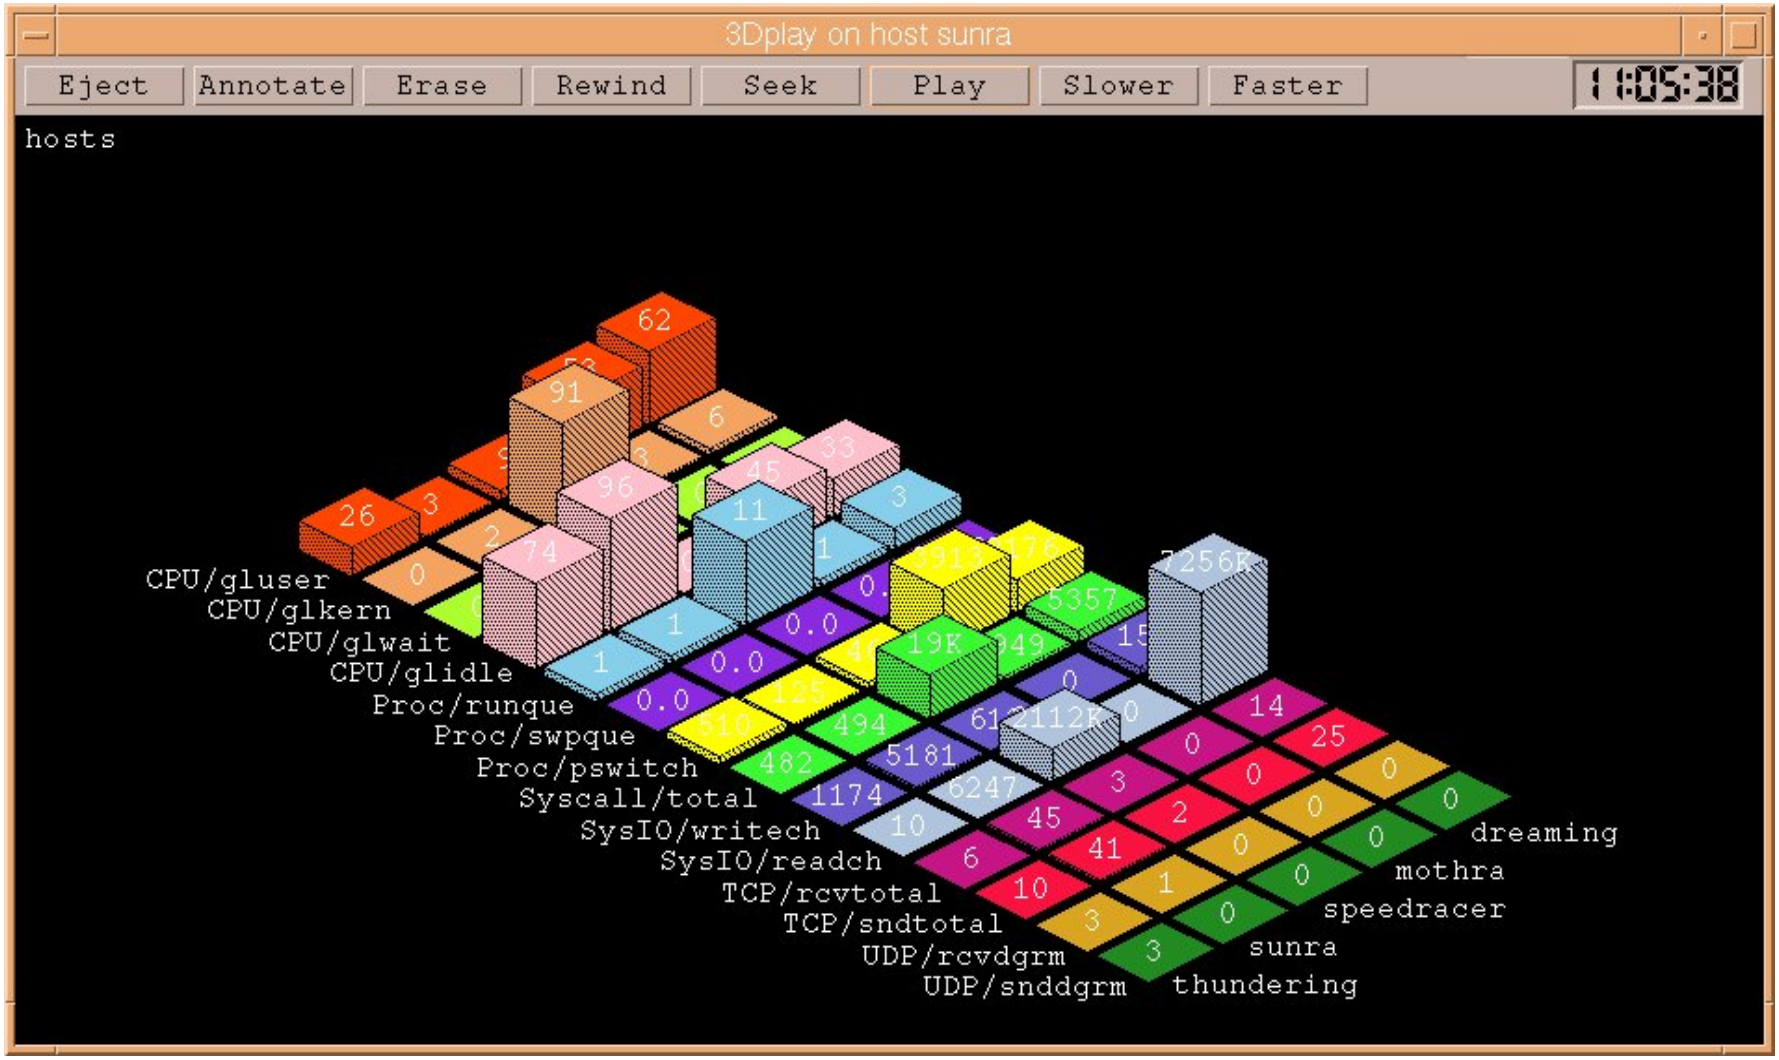

PTX example: 3dmon

AIX 5L Performance Toolbox (PTX)

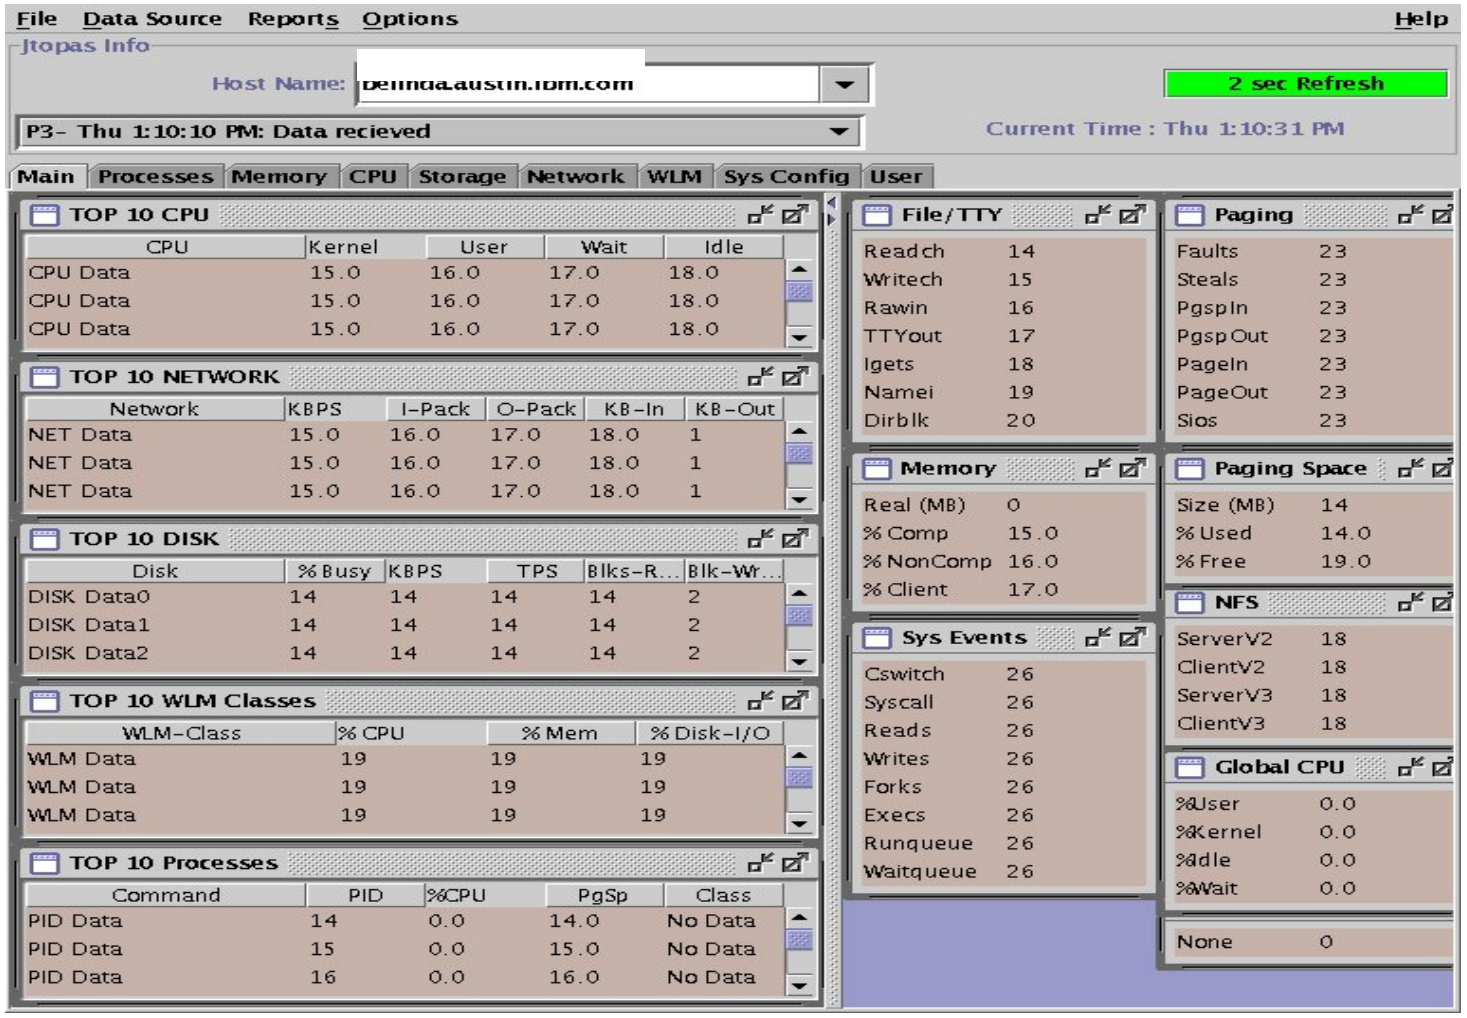

PTX example: jtopas

Additional Resources

- **The Campus for more education:**
  - Technical Course: "Implementing AIX 5L on eServer i5"
    http://www-306.ibm.com/services/learning/ites.wss/us/en?pageType=course_search&sortBy=5&searchType=1&sortDirection=9&includeNotScheduled=15&rowStart=0&rowsToReturn=20&maxSearchResults=200&searchString=as570
  - BPs: http://www.ibm.com/partnerworld/sales/systems/education
    - Includes link to IBM PartnerWorld University (Web lectures for key topics)
    - http://www.ibmweblectureservices.ihost.com/pwu
  - IBMers: w3.ibm.com/sales/systems/education
    - Includes links to the Online Universities for Cross-Brand and each Brand (Web lectures for key topics)
  - Customers: www-1.ibm.com/servers/eserver/education

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DBE, e-business logo, ESCO, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/30, VM/ESA, VSE/ESA, Websphere, xSeries, z/OS, zSeries, z/VM

The following are trademarks or registered trademarks of other companies

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

Intel is a registered trademark of Intel Corporation

* All other products may be trademarks or registered trademarks of their respective companies.

NOTES:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.