IBM @server POWER5 Processors Virtual SCSI Throughput Analysis

January 24, 2005

Elisabeth Stahl
Industry Solutions and Proof of Concept Centers
IBM Systems and Technology Group
estahl@us.ibm.com

Table of Contents

I. Executive Overview

With the arrival of IBM POWER5™ processor-based systems, advanced virtualization capabilities, including Micro-Partitioning™ and Virtual I/O, have significantly extended IBM's leading-edge virtualization technologies.

This throughput analysis highlights Virtual SCSI performance on the IBM @server® p5 570 system with the AIX 5L™ V5.3 operating system. This study is a follow-on to the previous "IBM @server POWER5 Virtual SCSI Processors Performance Study."

Throughput, as a measure of the amount of work performed over a period of time, is an important performance expectation. This paper describes the effects of varying throughput, including workload used, systems configuration, performance tools, benchmark test script, performance results and conclusions.

This analysis examines the performance of a Virtual SCSI (VSCSI) implementation using the Iozone workload. Local I/O is compared with Virtual I/O using FAStT (DS4500) and SCSI configurations at various throughput levels. Additionally, the effects of simultaneous multi-threading vs. single threading (ST) are considered.

VSCSI implementation is highlighted as a solution for clients looking to consolidate I/O resources with a small amount of processor overhead. The new Virtual SCSI capability on POWER5 systems creates new opportunities for consolidation, while demonstrating strong performance and manageability.

II. Introduction

Virtualization is a critical component in the on demand operating environment, and the system technologies implemented in the POWER5 processor-based systems provide a significant advancement in the enablement of these functions. Virtualization creates virtual server and storage resources and "maps" them to physical resources. IBM Virtualization Engine™ systems technologies integrated into the @server p5 system hardware and operating systems include the POWER™ Hypervisor™, Micro-Partitioning, Virtual LAN, Virtual I/O, Capacity on Demand, simultaneous multi-threading, and multiple operating system support.

Virtual I/O, which is the focus of this analysis, provides the capability for a single physical I/O adapter to be used by multiple logical partitions of the same server, allowing consolidation of I/O resources. Virtual I/O allows POWER5 systems to support more partitions than they have slots for I/O devices by enabling the sharing of the I/O adapters among partitions. Virtual SCSI (VSCSI) enables a partition to access block-level storage that is not a physical resource of that partition.

Virtual SCSI is based on a client/server relationship. The Virtual I/O Server owns the physical resources while the logical client partitions access the virtual SCSI resources provided by the Virtual I/O Server. The server partition has physically attached I/O devices and exports one or more of these devices to other partitions. The client partition is a partition that has a virtual client adapter node defined in its device tree and relies on the server partition to provide access to one or more block interface devices. Figure 1 shows an example of partition access to Virtual SCSI devices.

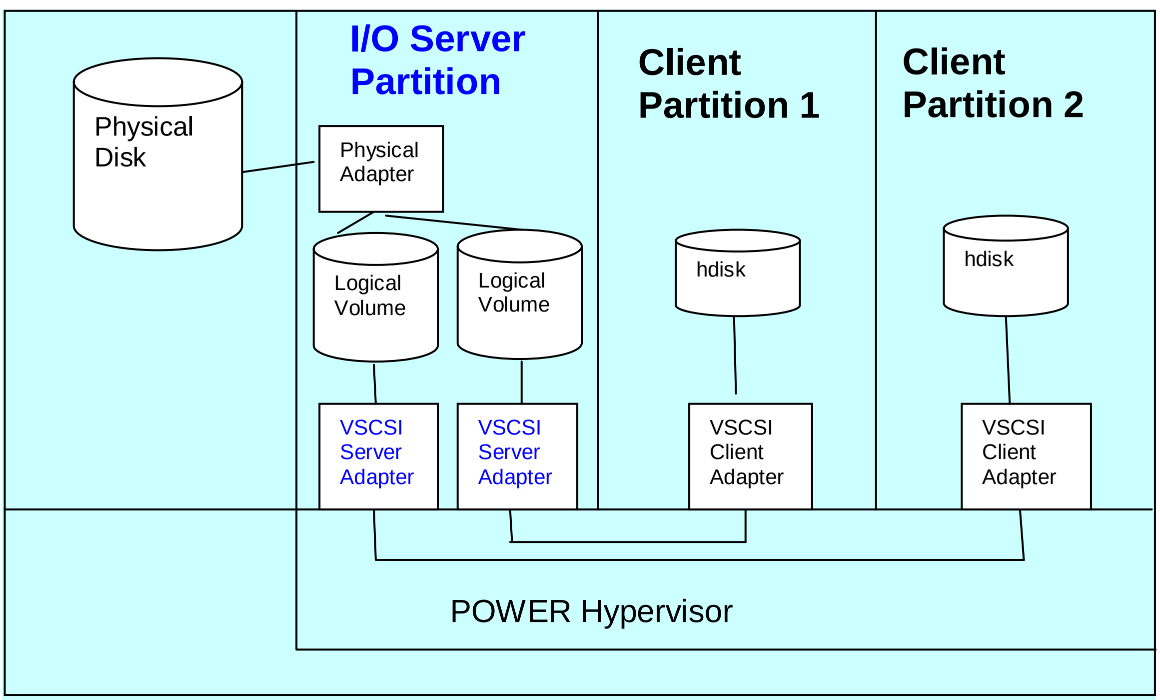

Figure 1. Virtual SCSI Architecture Overview Example

The client and server drivers operate together with the Hypervisor providing the means of communication between the two. The client emulates a physical SCSI adapter to the disk drivers by accepting requests for storage services and using interpartition communication facilities to transmit those requests to the server driver. The server driver completes the requests and converts the results and returns the results back to the client driver. See “Advanced Virtualization on @server p5 Servers Architecture and Performance Considerations” (SG24-5768) for transport and protocol details.

This analysis examines the performance of a VSCSI implementation using the Iozone workload running on an IBM @server p5 570 with AIX 5L V5.3. It is a follow-on to the previous “IBM @server POWER5 Virtual SCSI Processors Performance Study.”

Throughput, as a measure of the amount of work performed over a period of time, is an important performance expectation. Local I/O will be compared with Virtual I/O using FAStT (DS4500) and SCSI configurations at varied throughput rates. Additionally, the effects of simultaneous multi-threading versus single threading will be discussed as they relate to this Virtual I/O environment.

III. Test Workload Description

For the VSCSI throughput analysis, the Iozone workload was used. Iozone is a filesystem benchmark tool which generates and measures a variety of file operations. The benchmark tests file I/O performance for the following operations:

> *Read, write, re-read, re-write, read backwards, read strided, fread, fwrite, random read, pread, mmap, aio_read, aio_write*

In this study, filesystems were created on the target devices outlined in Table 1.

Filesystem Name	Device	Description
/local_fs	Local – hdisk0	2GB jfs2 filesystem on local SCSI disk
/fast_loc_fs1	Local – hdisk5	200GB jfs2 filesystem on local FAStT (DS4500)
/fast_loc_fs3	Local – hdisk6	200GB jfs2 filesystem on local FAStT (DS4500)
/fast_loc_fs4	Local – hdisk7	200GB jfs2 filesystem on local FAStT (DS4500)
/scsi_virt_fs	VSCSI – hdisk1	2GB jfs2 filesystem using a virtual scsi disk which is a scsi disk served from the I/O Server
/fast_virt_fs1	VSCSI – hdisk2	200GB jfs2 filesystem using a virtual scsi disk which is a FAStT (DS4500) disk served from the I/O Server
/fast_virt_fs3	VSCSI – hdisk3	200GB jfs2 filesystem using a virtual scsi disk which is a FAStT (DS4500) disk served from the I/O Server
/fast_virt_fs4	VSCSI – hdisk4	200GB jfs2 filesystem using a virtual scsi disk which is a FAStT (DS4500) disk served from the I/O Server

Table 1. Filesystems used for VSCSI Throughput Analysis

The Iozone workload was run for each test using commands as highlighted in Figure 2. The Iozone tool provides output as shown in Figure 2.

```
File size set to 1024000 KB
        Command line used: iozone -s1000m -i0 -t3 -F /fast_virt_fs1/test /fast_virt_fs3/test3 /fast_virt_fs4/test4
        Output is in Kbytes/sec
        Time Resolution = 0.000001 seconds.
        Processor cache size set to 1024 Kbytes.
        Processor cache line size set to 32 bytes.
        File stride size set to 17 * record size.
        Throughput test with 3 processes
        Each process writes a 1024000 Kbyte file in 4 Kbyte records

        Children see throughput for  3 initial writers  =  631442.91 KB/sec
        Parent sees throughput for  3 initial writers   =  104765.00 KB/sec
        Min throughput per process                      =  172194.53 KB/sec
        Max throughput per process                      =  267237.88 KB/sec
        Avg throughput per process                      =  210480.97 KB/sec
        Min xfer                                        =  676480.00 KB

        Children see throughput for  3 rewriters        =  568090.59 KB/sec
        Parent sees throughput for  3 rewriters         =  132455.63 KB/sec
        Min throughput per process                      =  165386.73 KB/sec
        Max throughput per process                      =  203534.47 KB/sec
        Avg throughput per process                      =  189363.53 KB/sec
        Min xfer                                        =  832228.00 KB
```

Figure 2. Iozone Example

Iozone was run for this analysis in throughput mode. This option allows the user to specify how many processes to have active during the measurement :

Iozone –s# –iy –tx –F /fs1/file1 /fs2/file2 . . . /fsx/filex

where

- # = the file size of the file to test followed by k(Kbytes), m(Mbytes), or g(Gbytes)
- y = the parameter to specify which tests to run (0=write/rewrite, 1=read/reread, . . .)
- x = the number of concurrent processes.

Option -F is followed by the temporary file names (/filesystem/file name) to be used in the throughput testing. The number of names should be equal to the number of processes that are specified.

In the example in Figure 2, three concurrent processes were run. The filesystems (option -F) used were virtual SCSI FAStT systems with filenames *test1, test3, test4* created. The file size (option -s) in all cases was set to 1000MB or 1GB. Write tests (option -i0) were performed in all cases.

IV. System Architecture and Performance Tools

Configuration

The @server p5 570 system used in this benchmark was configured into logical partitions as shown in Figure 3. The Virtual I/O Server partition is using one dedicated processor and acts as the AIX 5L hosting partition. The VSCSI client partition is using two processors and acts as the AIX 5L hosted partition.

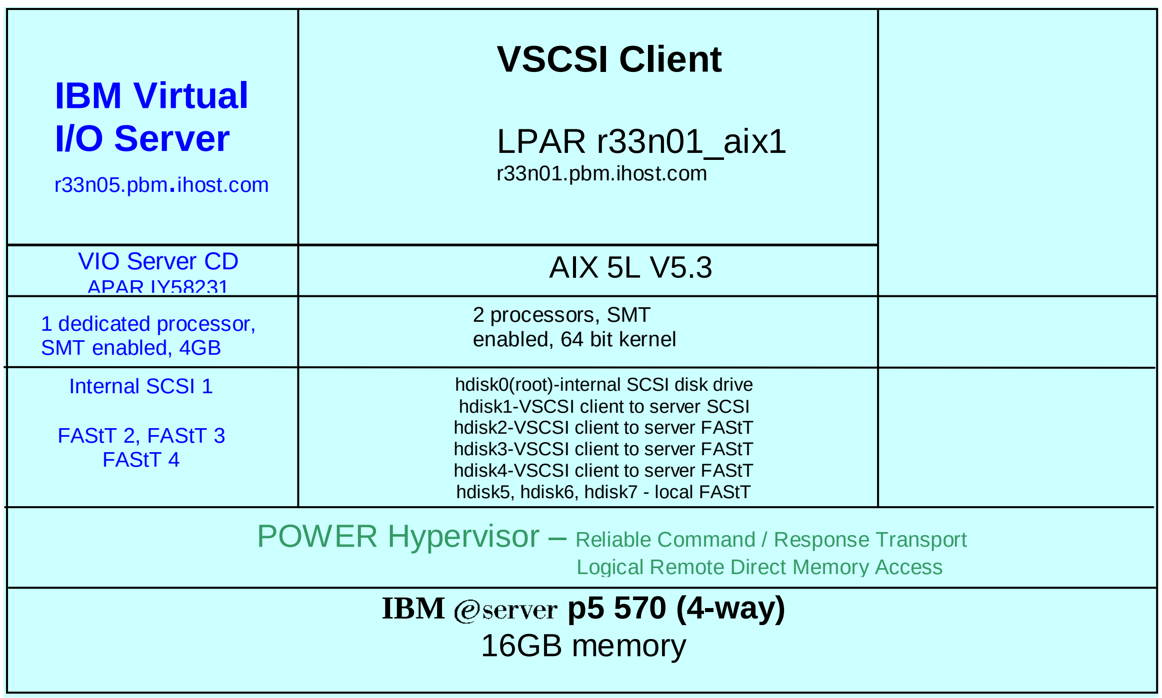

Figure 3. VSCSI Throughput Analysis System Configuration

System Performance Tools

The iostat AIX 5L Performance Tool and the AIX® Performance Toolbox were used for this study.

iostat

The iostat command is used for monitoring system input/output device load by observing the time the physical disks are active in relation to their average transfer rates. The iostat command generates reports that can be used to determine an imbalanced system configuration to tune the I/O load between physical disks and adapters. iostat is traditionally used to detect I/O bottlenecks by monitoring the disk utilization and can also be used to identify CPU problems, assist in capacity planning, and provide insight into solving I/O problems.

The iostat parameters used in this study were *interval* and *count*. Interval specifies the update period or amount of time between each report in seconds. Count specifies the number of iterations. The value of count determines the number of reports generated at interval seconds apart. In this study, the interval used was 20 seconds with a count of four as in Figure 4.

```
r33n01:/u/b4p527a/src/current $ iostat 20 4
```

Figure 4. Sample Performance Tools Command

AIX Performance Toolbox

The AIX Performance Toolbox consists of two parts: the manager and the agent. The agent must be loaded on all nodes that are to be monitored by the manager. In the configuration for this VSCSI analysis, the manager was configured on the client partition (r33n01) and the agents were configured on the client partition (r33n01) and the I/O Server partition (r33n05).

The AIX Performance Toolbox (PTX) can be used for load monitoring, analysis and control, and capacity planning. Load monitoring, which was used in this study, assists in monitoring system resources to detect performance problems. Analysis and control determines the correct tool for analyzing a problem and the root cause of the problem so that the necessary corrective action can be taken. Capacity planning performs long-term monitoring to determine in advance the correct quantity of additional resources required.

Many tools are available in PTX. The components used in this analysis were xmperf and 3dmon.

The xmperf program is used to monitor the statistics of a system. The xmperf program has the ability to monitor statistics on the system where it is running (in this case, the client partition) as well as on remote systems (the I/O partition). Xmperf can monitor performance in real time as well as record performance and play back a recording.

The windows or panes that are displayed by the xmperf program are referred to as consoles. Monitoring instruments occupy an area within a console. Only values from the same system can be displayed in an instrument. For this analysis, a recording instrument was created which monitored information over a period of time. The displayed values are pushed to the left as they are replaced by new values.

Figure 5 shows a user-defined console with three monitoring instruments created. The top instrument contains measurements for %CPU busy from the central processor statistics on the client partition. The middle instrument contains measurements for %CPU busy on the I/O server partition. The bottom instrument contains %busy from the disk statistics.

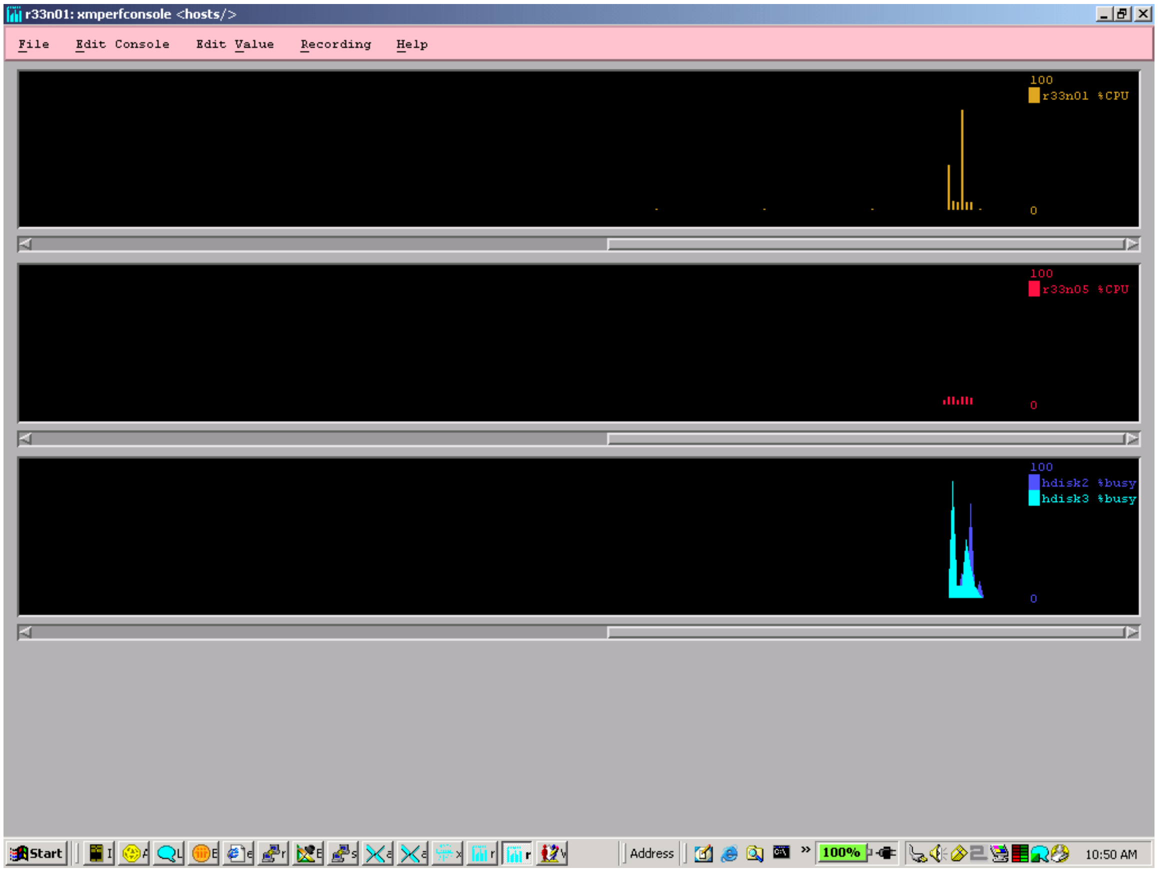

Figure 5. Example of xmperf User-Defined Console with Monitoring Instruments

The 3dmon program is useful for monitoring the same statistics on numerous hosts across a network. The results of the 3dmon program are three-dimensional and graphical. In this analysis, both host names for the client and server partitions are selected.

Figure 6 shows the typical chessboard window output of the 3dmon command. The names of the hosts can be seen on the right. In front of each bar is the name of the statistic that is being monitored and displayed. On top of each bar is the actual value of the statistic being measured. Since two hosts are being monitored, the three-dimensional display is staggered.

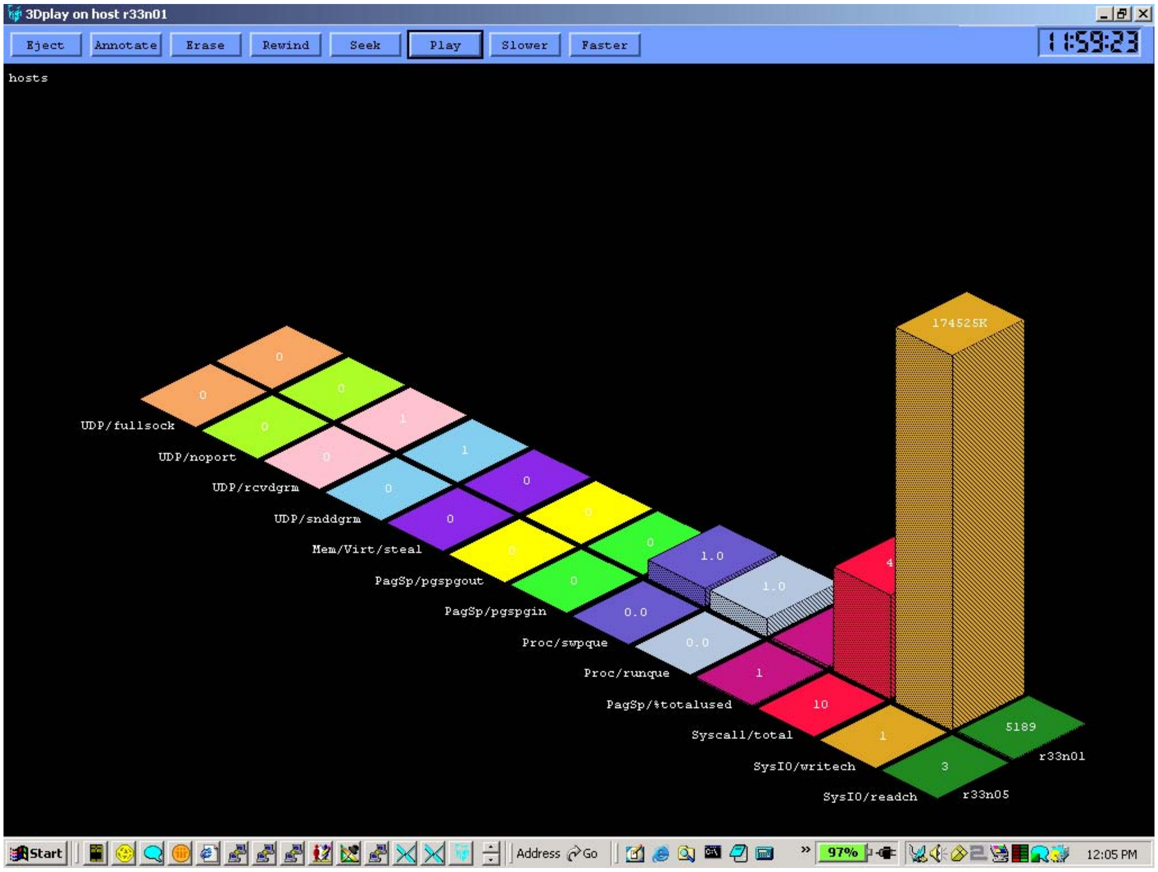

Figure 6. 3dmon Example for VSCSI Client and Server Partitions

V. Benchmark Test Script

The tests in the script performed in the VSCSI Throughput Analysis are listed in Table 2.

Virtual I/O Benchmark VSCSI Script	Test Details
Test 1	Baseline – Local (Native) I/O with Internal SCSI and with Local FAStT (DS4500) at varied throughput rates
Test 2	VSCSI Overhead - VSCSI Client to VIO Server SCSI and VIO Server FAStT (DS4500) at varied throughput rates
Test 3	Simultaneous Multi-threading to Single Threaded Effect on local I/O
Test 4	Simultaneous Multi-threading to Single Threaded Effect on VSCSI

Table 2. Virtual I/O Throughput Analysis VSCSI Script

Test 1: Baseline – Local I/O

The baseline tests for this study use a local internal SCSI device and local FAStT (DS4500). See the Virtual SCSI Configuration in Figure 1 for more details.

Test 2: Overhead of VSCSI using SCSI and FAStT (DS4500)

This test uses the VSCSI feature of the POWER5 processor-based system backed by an internal SCSI device and backed by FAStT900 (DS4500) on the hosting partition (I/O Server).

Test 3: Simultaneous Multi-threading Effect on Local I/O

A new feature of the POWER5 processor is simultaneous multi-threading which presents the kernel with two logical CPUs per virtual or dedicated CPU. Simultaneous multi-threading is a hardware design enhancement in POWER5 processors that allows two separate instruction threads to execute simultaneously on the processor. More details on simultaneous multi-threading can be found in Advanced Virtualization on @server p5 Servers Introduction and Basic Configuration (SG24-7940). This test examines the effect of setting SMT=OFF, thereby utilizing a single threaded (ST) environment, with local I/O.

Test 4: Simultaneous Multi-threading Effect on VSCSI

This test examines the effect of setting SMT=OFF, thereby utilizing a single threaded (ST) environment, with VSCSI.

VI. Performance Results

For each of the benchmark tests, performance data was obtained for throughput and CPU utilizations on both the client partitions and the I/O Server partition. A CPU busy to throughput ratio was then computed in order to analyze linear scalability of overhead. The lower the ratio, the more throughput achieved per CPU cycle.

Baseline and VSCSI Implementations – Tests 1 and 2

Table 3 highlights the performance results comparing the workload using local I/O vs. VSCSI at varied throughput rates. Note that all tests were performed with local FAStT or VSCSI backed by FAStT on the Virtual I/O Server. The highest throughput test also included access to local SCSI and VSCSI backed by SCSI on the I/O Server.

A small amount of overhead in using the VSCSI implementation can be seen in the CPU utilization. Note the linear scalability of the overhead as seen in the CPU busy to Throughput ratios. The overhead increases linearly and the rate does not change as throughput is increased, both in the local and in the VSCSI scenarios.

	Throughput (Mbps)	CPU-Client (% utilization)	CPU-Server (% utilization)	Total CPU (% utilization)	CPU busy to Throughput ratio
Local	92	16.8	0	16.8	.18
Local	153	28.8	0	28.8	.19
Local	168	31.2	0	31.2	.19
Local	199	37.4	0	37.4	.19
VSCSI	76	15.8	2.9	18.7	.25
VSCSI	110	22.2	4.8	27.0	.25
VSCSI	120	27.7	5.0	32.7	.27
VSCSI	172	35.9	7.0	42.9	.25

Table 3. Representative Performance Results for Local vs. Virtual SCSI

Effect of Setting SMT=OFF – Tests 3 and 4

Simultaneous multi-threading was set to off in both the client and server partitions to look at the effect of single threading. Table 4 highlights the performance results comparing the workload using local I/O vs. VSCSI. Note that all tests were performed with local FAStT or VSCSI backed by FAStT on the Virtual I/O Server. The highest throughput test also included access to local SCSI and VSCSI backed by SCSI on the I/O Server.

A small amount of overhead in using the VSCSI implementation can be seen in the CPU utilization. Note the linear scalability of the overhead as seen in the CPU busy to Throughput ratios. The overhead increases approximately linearly and the rate does not significantly change as throughput is increased, both in the local and in the VSCSI scenarios.

	Throughput (Mbps)	CPU-Client (% utilization)	CPU-Server (% utilization)	Total CPU (% utilization)	CPU busy to Throughput ratio
Local	94	18.4	0	18.4	.20
Local	154	34.2	0	34.2	.22
Local	164	36.4	0	36.4	.22
Local	205	47.2	0	47.2	.23
VSCSI	78	17.3	3.5	20.8	.27
VSCSI	110	24.8	4.4	29.2	.27
VSCSI	121	31.4	4.9	36.3	.30
VSCSI	175	45.5	7.3	52.8	.30

Table 4. Representative Performance Results for Local vs. Virtual SCSI with SMT=OFF

Note that at the data points where the throughput is approximately the same as the SMT=ON scenario, the CPU % utilization increases and the CPU busy to Throughput ratio increases, especially as the throughput levels increase.

The performance gain with simultaneous multi-threading demonstrates that local I/O as well as VSCSI benefits from simultaneous multi-threading. Without this feature, at the same throughput more CPU cycles must be used.

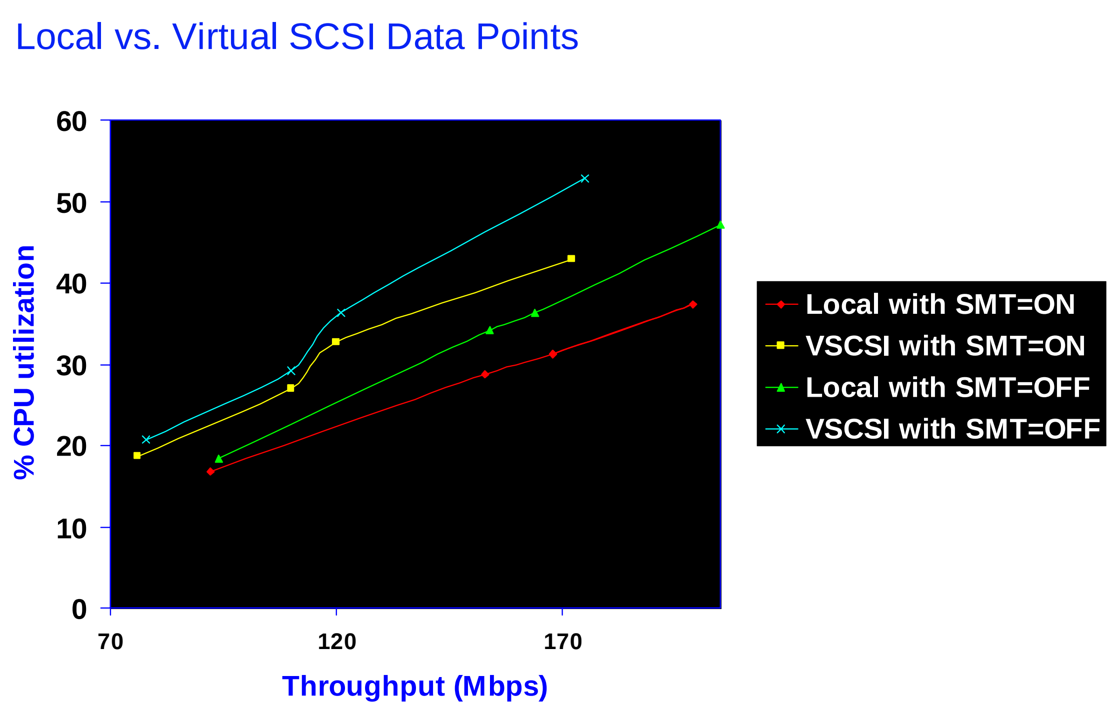

Figure 7. Graph of Local vs. Virtual SCSI Data

Figure 7 highlights the comparison between local and virtual I/O with simultaneous multi-threading implemented and then not implemented. Figure 8 shows the same data points but with linear trend lines added to highlight the rate of overhead increase. Note that as expected the line representing local I/O with SMT=ON has the least slope (rise / run) and the line representing VSCSI with SMT=OFF has the greatest slope. The lines representing local and VSCSI configurations with the same simultaneous multi-threading setting have very similar slopes.

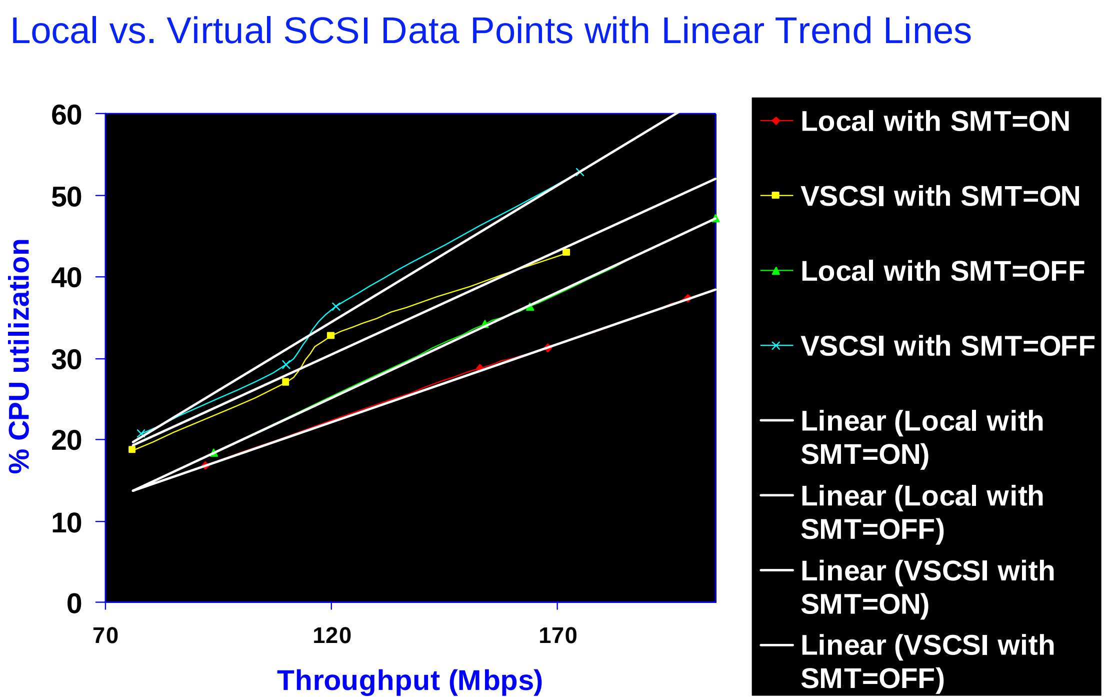

Figure 8. Graph of Local vs. Virtual SCSI Data Points with Linear Trend Lines

Figure 9 shows an AIX Performance Toolbox xmperf user-defined console with four monitoring instruments created for this analysis. The top instrument contains measurements for %CPU busy from the central processor statistics on the client partition. The second instrument contains measurements for %CPU busy on the I/O server partition. The third instrument contains %busy from the disk statistics. The bottom instrument contains the number of writes per second for each interval.

The first four tests in the console from left to right are using local I/O with throughput increasing. The next four are using Virtual SCSI with throughput increasing. Note that for the first four tests, the second instrument does not contain any graphs for these tests since there is no %CPU busy on the I/O server. Also note that as the throughput increases in each case, the CPU and disk busy spikes are higher and more frequent. The writes show the pattern of the iozone tests writing and rewriting the data. To better understand the data, these instruments are most useful in analyzing patterns as the tests are run.

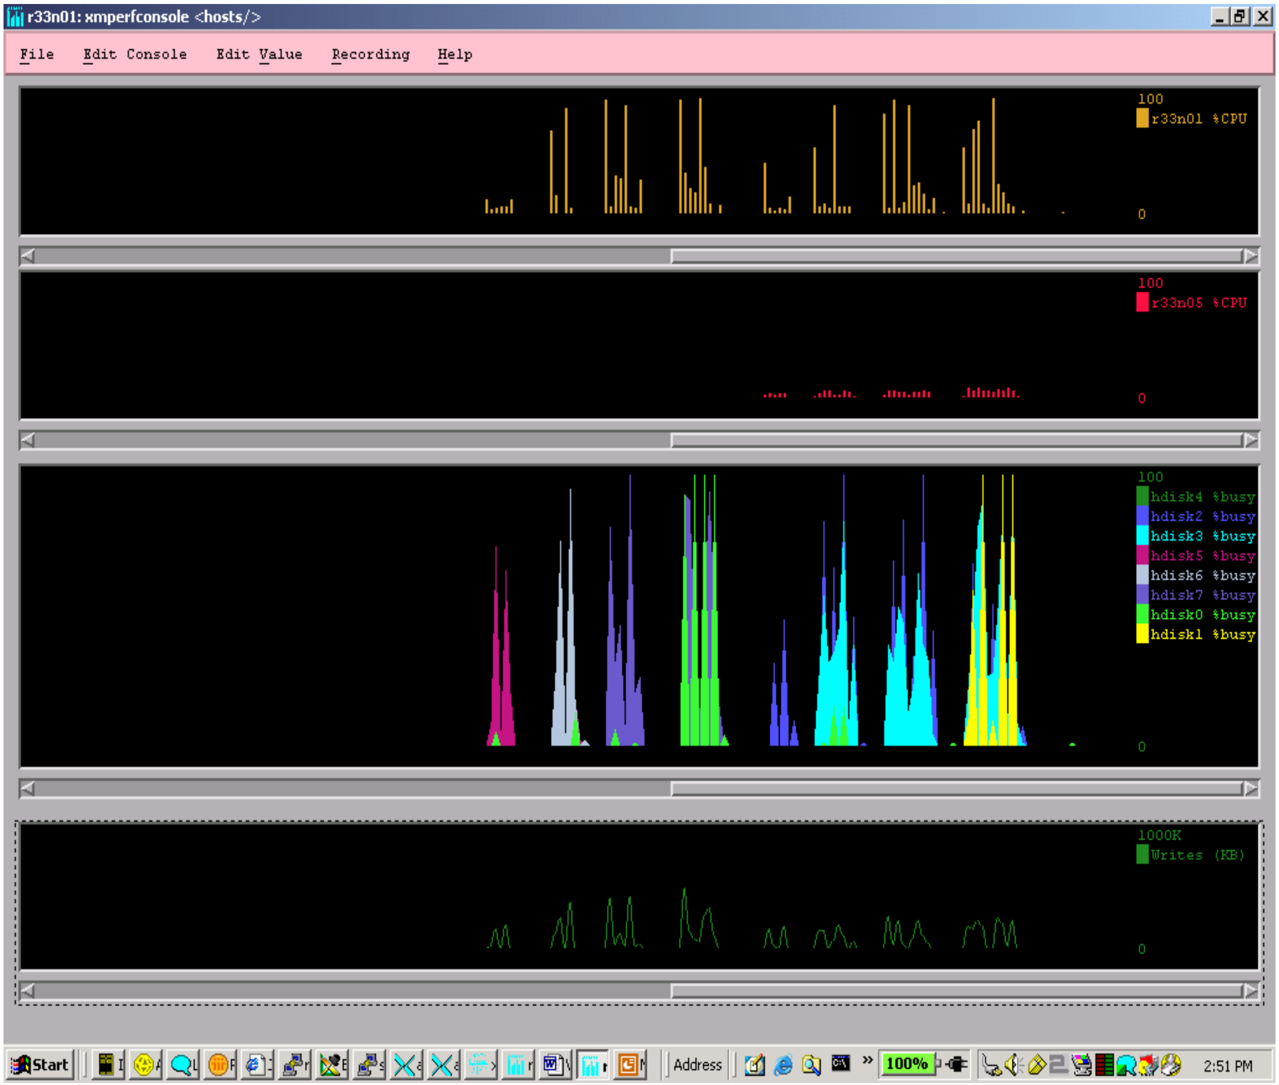

Figure 9. xmperf User-Defined Console with Monitoring Instruments for Local versus VSCSI Tests

* Note that these synthetic tests use simulated workloads and clients should validate their own workload prior to using this technique in a production environment. The results apply to the specific examples described in this paper.

VII. Further Study

Recommendations on further performance testing in the VSCSI environment include:

- Measuring performance of I/O server and client partitions using industry standard and ISV benchmarking workloads.

- Configuring additional micro-partitions in the Virtual I/O Server partition to analyze effect on overhead.

- Performing LVM mirroring analysis with two Virtual I/O Server partitions for complete redundancy. This system hardware did not have the resources to test this configuration.

VIII. Conclusions

Virtualization is an innovative technology that redefines the utilization and economics of managing an on demand operating environment. The POWER5 architecture provides new opportunities for clients to take advantage of virtualization capabilities. Virtual I/O provides the capability for a single physical I/O adapter to be used by multiple logical partitions of the same server, allowing consolidation of I/O resources.

This performance analysis has demonstrated how a client can implement a Virtual SCSI environment and what performance may be in different implementations using both FAStT (DS4500) and internal SCSI devices.

Overhead of VSCSI is small and relatively linear as throughput is increased. Clients should assess the benefits of the VSCSI implementation for their environment. When enabled, simultaneous multi-threading achieves increased performance in a VSCSI environment.

VSCSI implementation is an excellent solution for clients looking to consolidate I/O resources with a modest amount of processor overhead. The new Virtual SCSI capability on the POWER5 architecture creates new opportunities for consolidation, while demonstrating strong performance and manageability.

IX. References

IBM @server POWER5 Virtual SCSI Processors Performance Study
http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100504
http://w3-1.ibm.com/sales/systems/portal/_s.155/254?navID=f320s240&geoID=AM&prodID=IBM%20eServer%20Products&docID=pstivirtualscsiperf

Advanced Virtualization on @server p5 Servers: Introduction and Basic Configuration
http://www.redbooks.ibm.com/redpieces/abstracts/sg247940.html

Advanced Virtualization on @server p5 Servers: Architecture and Performance Considerations
http://www.redbooks.ibm.com/redpieces/abstracts/sg245768.html

FAStT700 and AIX 5L V5.2: Performance Tests using iozone
http://w3-1.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD101189

iozone Filesystem Benchmark
http://www.iozone.org/

IBM @server p5 570 Server Consolidation Using POWER5 Virtualization White Paper
http://www-1.ibm.com/servers/eserver/pseries/hardware/whitepapers/570_serverconsol.html

IBM @server p5 570 Workload Balancing using POWER5 Virtualization White Paper
http://www-1.ibm.com/servers/eserver/pseries/hardware/whitepapers/570_workload.html

AIX 5L Performance Tools Handbook
http://www.redbooks.ibm.com/redbooks/pdfs/sg246039.pdf

Thanks to the pSeries® Benchmark Center team of Julie Peet, Tina Tarquinio and Sharon Selzo; the AIX Performance Team of Jim Chen and Sergio Reyes; Bruno Blanchard in IGS; and the marketing efforts of Terry Brennan and Nha D Vu.

X. Appendix

Devices on Client Partition

L2cache0	Available		L2 Cache
aio0	Available		Asynchronous I/O (Legacy)
dac0	Available	06-08-02	1742-900 (900) Disk Array Controller
dac1	Available	07-08-02	1742-900 (900) Disk Array Controller
dar0	Available		1742-900 (900) Disk Array Router
en0	Available	04-08	Standard Ethernet Network Interface
ent0	Available	04-08	10/100/1000 Base-TX PCI-X Adapter (14106902)
et0	Defined	04-08	IEEE 802.3 Ethernet Network Interface
fast_loc_lv1	Defined		Logical volume
fast_loc_lv2	Defined		Logical volume
fast_loc_lv3	Defined		Logical volume
fast_loc_lv4	Defined		Logical volume
fast_loc_lv5	Defined		Logical volume
fast_loc_vg	Defined		Volume group
fast_virt_lv1	Defined		Logical volume
fast_virt_lv2	Defined		Logical volume
fast_virt_lv3	Defined		Logical volume
fast_virt_lv4	Defined		Logical volume
fast_virt_lv5	Defined		Logical volume
fast_virt_vg	Defined		Volume group
fcnet0	Defined	06-08-01	Fibre Channel Network Protocol Device
fcnet1	Defined	07-08-01	Fibre Channel Network Protocol Device
fcs0	Available	06-08	FC Adapter
fcs1	Available	07-08	FC Adapter
fscsi0	Available	06-08-02	FC SCSI I/O Controller Protocol Device
fscsi1	Available	07-08-02	FC SCSI I/O Controller Protocol Device
fslv00	Defined		Logical volume
fslv01	Defined		Logical volume
fslv02	Defined		Logical volume
fslv03	Defined		Logical volume
gxme0	Defined		Graphics Data Transfer Assist Subsystem
hd1	Defined		Logical volume
hd2	Defined		Logical volume
hd3	Defined		Logical volume
hd4	Defined		Logical volume
hd5	Defined		Logical volume
hd6	Defined		Logical volume
hd8	Defined		Logical volume
hd10opt	Defined		Logical volume
hd9var	Defined		Logical volume
hdisk0	Available	02-08-00-5,0	16 Bit LVD SCSI Disk Drive
hdisk1	Available		Virtual SCSI Disk Drive
hdisk2	Available		Virtual SCSI Disk Drive
hdisk3	Available		Virtual SCSI Disk Drive
hdisk4	Available		Virtual SCSI Disk Drive
hdisk5	Available	07-08-02	1742-900 (900) Disk Array Device
hdisk6	Available	06-08-02	1742-900 (900) Disk Array Device
hdisk7	Available	07-08-02	1742-900 (900) Disk Array Device
inet0	Available		Internet Network Extension
iscsi0	Available		iSCSI Protocol Device
lg_dumplv	Defined		Logical volume
lo0	Available		Loopback Network Interface
loglv00	Defined		Logical volume
loglv01	Defined		Logical volume
loglv02	Defined		Logical volume
loglv03	Defined		Logical volume
lv00	Defined		Logical volume
lvdd	Available		LVM Device Driver
mem0	Available		Memory
nampd0	Available		N/A
pci0	Available		PCI Bus
pci1	Available		PCI Bus
pci2	Available	00-10	PCI Bus

```
pci3        Available 00-12       PCI Bus
pci4        Available 01-10       PCI Bus
pci5        Available             PCI Bus
pci6        Available 01-14       PCI Bus
pci7        Available 05-16       PCI Bus
posix_aio0  Defined               Posix Asynchronous I/O
proc0       Available 00-00       Processor
proc2       Available 00-02       Processor
proc4       Defined   00-04       Processor
proc6       Defined   00-06       Processor
proc8       Defined   00-08       Processor
pty0        Available             Asynchronous Pseudo-Terminal
rcm0        Defined               Rendering Context Manager Subsystem
rootvg      Defined               Volume group
scsi0       Available 02-08-00    PCI-X Ultra320 SCSI Adapter bus
scsi_virt_vg Defined              Volume group
ses0        Available 02-08-00-15,0 SCSI Enclosure Services Device
sisscsia0   Available 02-08       PCI-X Ultra320 SCSI Adapter
sys0        Available             System Object
sysplanar0  Available             System Planar
usb0        Available             USB System Software
usbhc0      Available 03-08       USB Host Controller (33103500)
usbhc1      Available 03-09       USB Host Controller (33103500)
vio0        Available             Virtual I/O Bus
vsa0        Available             LPAR Virtual Serial Adapter
vscsi0      Available             Virtual SCSI Client Adapter
vscsi1      Available             Virtual SCSI Client Adapter
vty0        Available             Asynchronous Terminal
```

Devices on I/O Server Partition

```
L2cache0    Available             L2 Cache
aio0        Defined               Asynchronous I/O (Legacy)
cd0         Available 08-08-00    IDE DVD-ROM Drive
dac0        Available 04-08-02    1742-900 (900) Disk Array Controller
dac1        Available 06-08-02    1742-900 (900) Disk Array Controller
dar0        Available             1742-900 (900) Disk Array Router
en0         Defined   03-08       Standard Ethernet Network Interface
en1         Defined   03-09       Standard Ethernet Network Interface
en2         Available 09-08       Standard Ethernet Network Interface
ent0        Available 03-08       2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent1        Available 03-09       2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent2        Available 09-08       10/100/1000 Base-TX PCI-X Adapter (14106902)
et0         Defined   03-08       IEEE 802.3 Ethernet Network Interface
et1         Defined   03-09       IEEE 802.3 Ethernet Network Interface
et2         Defined   09-08       IEEE 802.3 Ethernet Network Interface
fcnet0      Defined   04-08-01    Fibre Channel Network Protocol Device
fcnet1      Defined   06-08-01    Fibre Channel Network Protocol Device
fcs0        Available 04-08       FC Adapter
fcs1        Available 06-08       FC Adapter
fscsi0      Available 04-08-02    FC SCSI I/O Controller Protocol Device
fscsi1      Available 06-08-02    FC SCSI I/O Controller Protocol Device
gxme0       Defined               Graphics Data Transfer Assist Subsystem
hd1         Defined               Logical volume
hd2         Defined               Logical volume
hd3         Defined               Logical volume
hd4         Defined               Logical volume
hd5         Defined               Logical volume
hd6         Defined               Logical volume
hd8         Defined               Logical volume
hd10opt     Defined               Logical volume
hd9var      Defined               Logical volume
hdisk0      Available 07-08-00-3,0  16 Bit LVD SCSI Disk Drive
hdisk1      Available 07-08-00-4,0  16 Bit LVD SCSI Disk Drive
hdisk2      Available 07-08-00-5,0  16 Bit LVD SCSI Disk Drive
hdisk3      Available 06-08-02    1742-900 (900) Disk Array Device
hdisk4      Available 04-08-02    1742-900 (900) Disk Array Device
hdisk5      Available 04-08-02    1742-900 (900) Disk Array Device
ide0        Available 08-08       ATA/IDE Controller Device
inet0       Available             Internet Network Extension
iscsi0      Available             iSCSI Protocol Device
```

```
lg_dumplv  Defined              Logical volume
lo0        Available            Loopback Network Interface
lvdd       Available            LVM Device Driver
mem0       Available            Memory
paging00   Defined              Logical volume
pci0       Available            PCI Bus
pci1       Available            PCI Bus
pci2       Available            PCI Bus
pci3       Available 02-14      PCI Bus
pci4       Available 00-12      PCI Bus
pci5       Defined   00-14      PCI Bus
pci6       Available 00-16      PCI Bus
pci7       Available 01-12      PCI Bus
pci8       Available 01-14      PCI Bus
pci9       Available 01-10      PCI Bus
pci10      Defined   01-16      PCI Bus
posix_aio0 Defined              Posix Asynchronous I/O
proc0      Available 00-00      Processor
pty0       Available            Asynchronous Pseudo-Terminal
rcm0       Defined              Rendering Context Manager Subsystem
rootvg     Defined              Volume group
scsi0      Available 07-08-00   PCI-X Ultra320 SCSI Adapter bus
ses0       Available 07-08-00-15,0 SCSI Enclosure Services Device
sisscsia0  Available 07-08      PCI-X Ultra320 SCSI Adapter
sys0       Available            System Object
sysplanar0 Available            System Planar
vhost0     Available            Virtual SCSI Server Adapter
vhost1     Available            Virtual SCSI Server Adapter
vio0       Available            Virtual I/O Bus
virtfastt1 Available            Virtual Target Device - Disk
virtfastt2 Available            Virtual Target Device - Disk
virtfastt3 Available            Virtual Target Device - Disk
virtscsi   Available            Virtual Target Device - Disk
vsa0       Available            LPAR Virtual Serial Adapter
vty0       Available            Asynchronous Terminal
```

IBM Corporation
Marketing Communications
Systems Group
Route 100
Somers, New York 10589

Produced in the United States of America
January 2005