



Performance Testing of AIX 5L Virtual Storage on IBM System i5™

Vess Natchev
AIX 5L and System i5 Technical Specialist
IBM Rochester, MN

March 14th, 2006

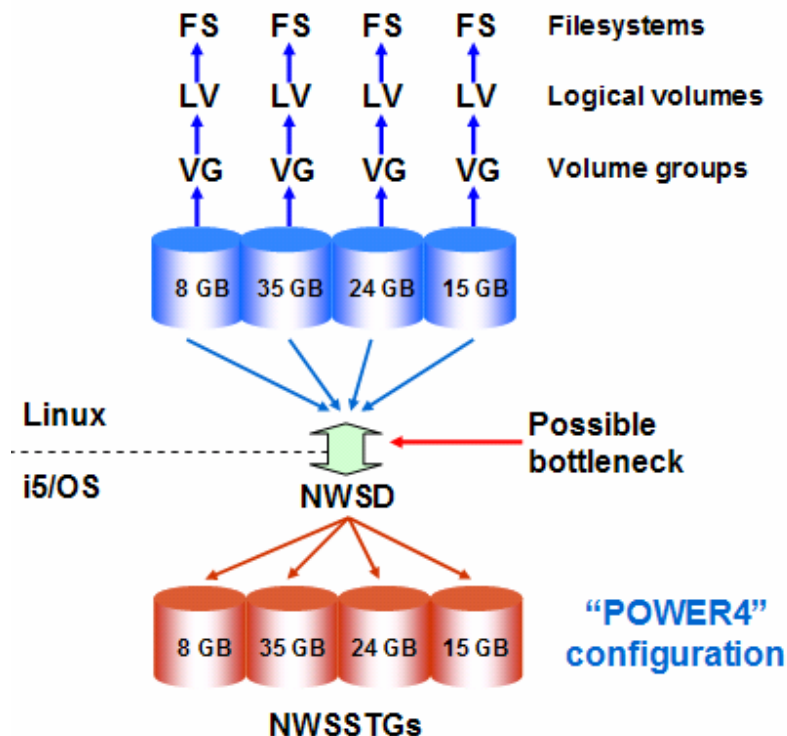


1.0 Overview: Virtual Storage Hosted by i5/OS

Virtual storage for AIX 5L hosted by i5/OS adds to the value of the IBM System i5 by simplifying the IT infrastructure, enabling rapid test and deployment of new applications, and allowing AIX 5L to leverage the i5/OS single-level store and object-oriented architecture. Hosting virtual storage for AIX 5L in i5/OS simplifies the IT environment because all management of physical adapters, disk units or SAN volumes is performed within i5/OS, using an established set of management skills. An administrator can add new virtual disks in i5/OS and assign them to a running AIX 5L logical partition (LPAR) dynamically, which can significantly decrease the time required to begin testing a new application. While virtual disks appear in AIX 5L as physical disk units (*hdisks*), they are really objects in i5/OS and take full advantage of scatter-loading over multiple physical disk arms and of any i5/OS disk protection (RAID or mirroring).

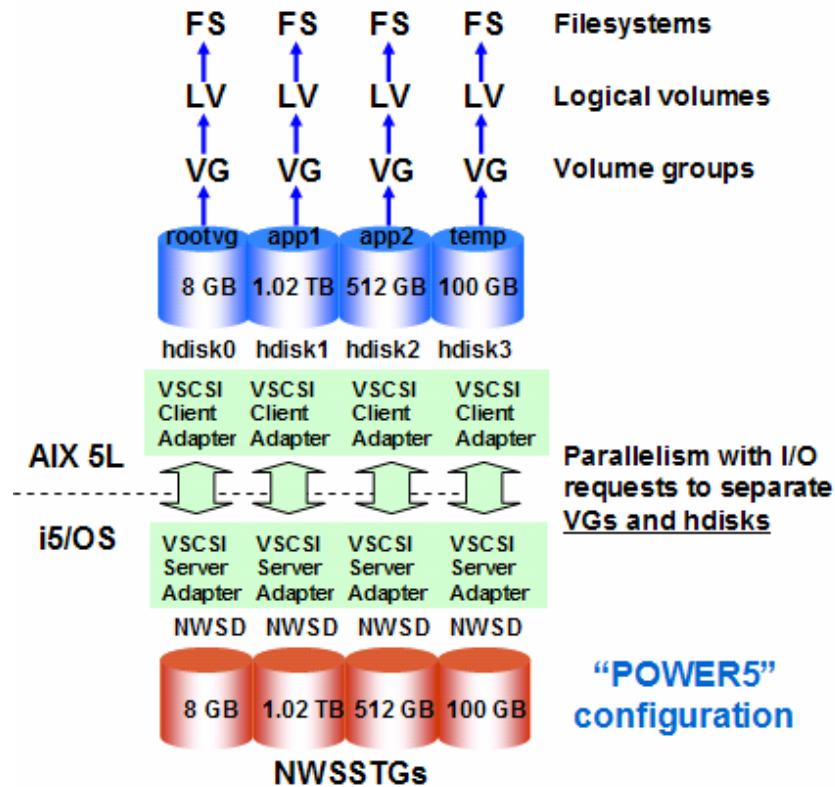
1.1 “POWER4” vs. “POWER5” Virtual Storage Configurations

Virtual storage hosted by i5/OS has been available on iSeries systems since early 2001, when i5/OS was known as “OS/400.” At the time, hosted storage was available to any Linux LPAR on the iSeries, and the configuration object that linked a virtual disk to a Linux LPAR was a Network Server Description (NWSA). Each NWSA corresponded to a single Linux LPAR, but multiple virtual disks, or Network Server Storage Spaces (NWSSTGs) could be attached to an NWSA. Prior to the announce of the POWER5-based eServer i5 system in 2004, this configuration – 1 LPAR/1 NWSA/1 or more NWSSTGs – was the only possibility. The paper will refer to this virtual storage configuration (see diagram below) as “POWER4.”



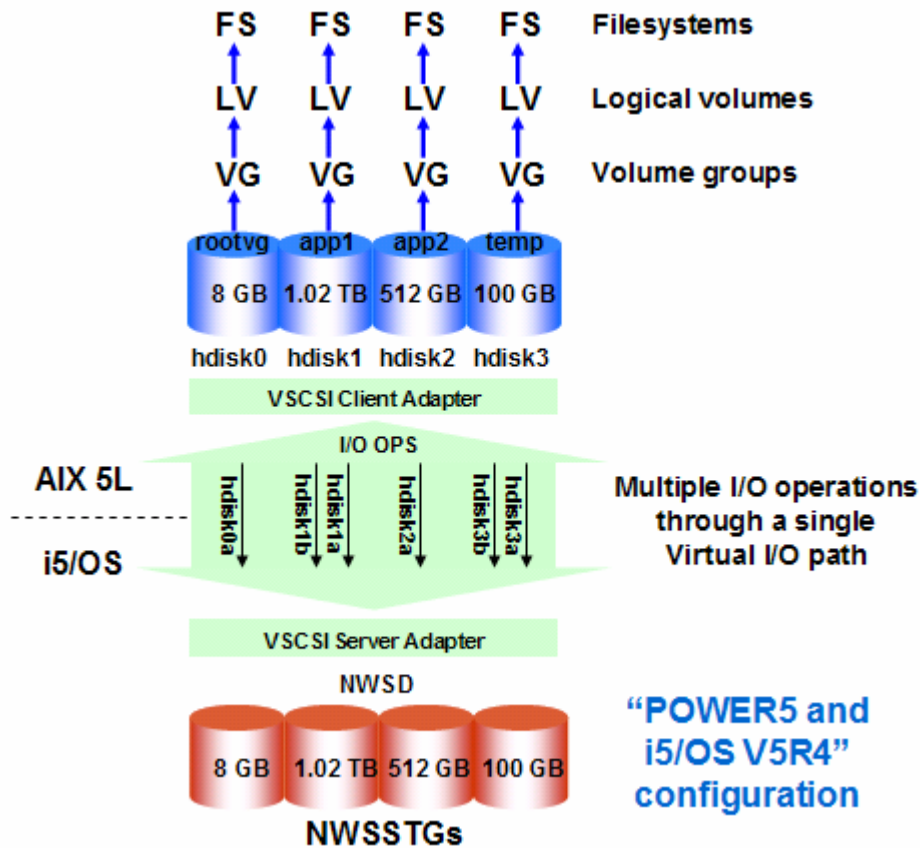
Since 2001, the “POWER4” configuration has performed well and is the most popular configuration choice for Linux partitions on iSeries systems. On the POWER5-based System i5 the “1 LPAR/1 NWSSTG/1 or more NWSSTGs” configuration is still fully supported. In addition to the single NWSSTG to define the virtual configuration, there is a pair of Virtual SCSI client (in the AIX 5L or Linux LPAR) and server (in i5/OS) adapters configured in the HMC. The POWER5 virtual storage architecture is based on the industry-standard SCSI protocol (please see <http://www.t10.org> for more information).

The “POWER4” configuration features a possible performance bottleneck because of the single Virtual SCSI (VSCSI) connection between AIX 5L and i5/OS. Even if multiple NWSSTGs are assigned to an AIX 5L LPAR, *there is still a single I/O path from the VSCSI client adapter in AIX 5L to the VSCSI server adapter in i5/OS*. To maximize performance with i5/OS V5R3, a System i5 administrator can create multiple Virtual SCSI client and server pairs of adapters between AIX 5L and i5/OS. Each VSCSI pipeline corresponds to an NWSSTG in i5/OS, with up to 64 NWSSTGs attached to each NWSSTG. This “POWER5” setup has multiple I/O paths, as shown on the diagram below:



1.2 Enhanced I/O Parallelism with i5/OS V5R4

With the introduction of System i5 and i5/OS V5R4 on January 31st, 2006, the virtual storage architecture has been further improved by implementing *Tagged Command Queuing*. In addition to enabling an AIX 5L LPAR to perform multiple simultaneous I/O operations through multiple VSCSI connections, AIX 5L can now *send multiple I/O requests to multiple NWSSTGs in i5/OS through a single Virtual I/O path*:



2.0 Statement of Purpose, Applicability, and Audience

IBM conducted a limited number of performance tests to determine the effect of the i5/OS V5R4 virtual storage enhancements and other factors on AIX 5L I/O throughput. This paper will describe the results of those tests and provide recommendations for configuring virtual storage for AIX 5L at i5/OS V5R3 and V5R4 for best performance. The results and recommendations in the paper apply to POWER5-based System i5 systems running i5/OS V5R3 and V5R4.

Please note: another supported method of providing virtual storage to AIX 5L on System i5 is to use a special LPAR called a Virtual I/O Server (VIO Server). VIO Server scenarios were not tested as part of this effort. The focus of this white paper is on the performance of AIX 5L virtual storage provided by i5/OS.

The intended audience for this paper is IBM Field Technical Sales Specialists, Systems Architects, Customer Engineers, Techline support specialists, IBM Business Partners, and other IT professionals designing or implementing AIX 5L virtual storage solutions on System i5.

The performance data contained in this document were determined in a controlled environment. Actual results may vary significantly and are dependent on many factors

including system hardware configuration, and software design and configuration. Users of this document should verify the applicable data for their specific environment.

3.0 Test Environment

The tests involved an i5/OS LPAR hosting virtual storage for up to 3 AIX 5L LPARs. Some of the tests were performed on a System i5 570 and some on a System i5 595, but the same type of DDR-1 memory, 1.65-GHz POWER5 CPUs, SCSI disk adapters and disk units were used in both cases.

i5/OS LPAR configuration:

- Operating system: i5/OS V5R3 or V5R4 (there were no test cases with V5R3M5 Licensed Internal Code)
- 8 GB of memory, 1 GB of memory in the machine pool (QMCHPOOL)
- 1 dedicated 1.65-GHz POWER5 processor
- Virtual SCSI configuration: between 1 and 10 VSCSI server adapters per hosted AIX 5L LPAR
- Physical SCSI configuration: 1 Ultra4 SCSI adapter (FC 2780) and 10 73-GB 15K-RPM drives per hosted AIX 5L LPAR in a RAID 5 set

AIX 5L LPAR configuration:

- Operating system AIX 5L v5.3D (or, ML 03)
- 1 dedicated 1.65-GHz POWER5 processor
- 2 GB of memory
- Virtual SCSI configuration: between 1 and 10 VSCSI client adapters
- Test filesystem = /perf2, which is created using a single logical volume in a volume group. The volume group contains between 1 and 10 virtual disks (hdisks or physical volumes in AIX 5L). The logical volume uses all physical volumes in the volume group (1-10), but it is *not* a striped logical volume. Instead, data in the logical volume are allocated to the physical volumes in 16-MB increments, corresponding to the size of the logical partition in the volume group. This logical volume configuration was chosen, because it is the more commonly used one in AIX 5L environments
- Filesystem mount options: rw, rbrw, dio. In addition to the typical read/write option (“rw”), the “release behind on reads on writes” option (“rbrw”) will cause AIX 5L to clear memory pages when they are no longer referenced, whereas the “direct I/O” option (“dio”) will cause I/O requests to bypass all kernel buffers.

4.0 Test Workload

This series of tests used the open source filesystem performance benchmark IOzone, which can be downloaded at no cost at <http://www.iozone.org>. The use of an openly available performance tool makes the test environment and results easy to reproduce.

IOzone version 3.254 was downloaded in source form and compiled with the GNU C Compiler (GCC) on AIX 5L with large file support. Below is an outline of the test methodology; please see Appendix A for detailed information on the IOzone options and shell scripts used:

- Between 1 and 12 IOzone tasks simultaneously wrote a different 8-GB or 100-GB file to the same filesystem, /perf2, in 4-KB to 1-MB blocks
- Sequential write I/O throughput was measured in MB/s
- The filesystem was unmounted to eliminate filesystem caching and measure strictly virtual storage throughput.
- Next, between 1 and 12 IOzone tasks each read one of the previously written 8-GB or 100-GB files
- Sequential read I/O throughput was measured in MB/s

The test script used a "warm-up" run, followed by 2 test runs. The throughput results are averages of the 2 test runs. Each run first writes and re-writes the file, then reads and re-reads it. Typically, higher throughput is registered on the re-write and re-read operations, because only the file data are changed and filesystem structures are not created for it.

PLEASE NOTE: this test environment and methodology were designed specifically to measure AIX 5L virtual storage throughput under various configurations in order to provide recommendations for designing solutions. AIX 5L I/O performance was deliberately constrained in several important ways to eliminate confounding factors and focus on the differences between the test conditions:

- The test files (8 or 100 GB) were much larger than the amount of system memory in the AIX 5L LPAR (2 GB)
- The test files (8 or 100 GB) were much larger than the amount of memory in the i5/OS machine pool (1 GB), where virtual I/O is processed at i5/OS V5R3
- The test filesystem was unmounted between write and read tests to eliminate filesystem caching
- The test filesystem was mounted with the "rbrw" and "dio" options, eliminating multiple references to the same memory page and bypassing the AIX 5L kernel buffers

In a production environment, the factors that were removed from the test would normally improve performance. Therefore, the test results do not illustrate the maximum I/O throughput that AIX 5L is capable of when using virtual storage from i5/OS. Instead, they provide support to recommendations on which configurations yield the optimal performance relative to one another.

5.0 Result Summaries by Test Scenario

Three main scenarios were tested. Scenario 1 examined the effect of introducing I/O parallelism by adding more VSCSI connections (and NWSDs) between AIX 5L and

i5/OS, and of upgrading from i5/OS V5R3 to V5R4. Scenario 2 examined disk contention in i5/OS by hosting storage for one vs. three AIX 5L LPARs in the System Auxiliary Storage Pool (ASP). Scenario 3 tested the effect of isolating hosted storage for three AIX 5L LPARs in separate user ASPs vs. hosting them all in the System ASP. Two additional scenarios were also tested: Scenario 4 examined the effect of larger block or record sizes and Scenario 5 tested the effect of increasing disk utilization in AIX 5L by running more simultaneous IOzone tasks. All results show only sequential *read* I/O throughput in MB/s.

5.1 Scenario 1: More NWSDs and i5/OS V5R3 vs. V5R4

Virtual Storage Configuration	i5/OS V5R3	i5/OS V5R4
1 NWSD, 1 virtual disk x 100 GB, 8 IOzone tasks	41 MB/s	66.9 MB/s
1 NWSD, 10 virtual disks x 10 GB, 8 IOzone tasks	32.2 MB/s	78.8 MB/s
10 NWSDs, 10 virtual disks x 10 GB, 8 IOzone tasks	61.6 MB/s	80.7 MB/s

In this test, 8 IOzone tasks wrote and read a different 8-GB file in 1-MB blocks to the same filesystem. At i5/OS V5R3, increasing the number of NWSDs (and VSCSI connections) from 1 to 10 provided a 50% boost in performance (41 → 61.6 MB/s). Adding just NWSSTGs decreased performance in this case (41 → 32.2 MB/s), because disk utilization (not shown) dropped from almost 100% to about 25%.

Upgrading from i5/OS V5R3 to V5R4 with a single large NWSSTG increased throughput by 63% (41 → 66.9 MB/s). At i5/OS V5R4, increasing the number of NWSSTG did provide a small boost of about 18% (66.9 → 78.8 MB/s) despite the lower disk utilization, because multiple I/O operations were dispatched to multiple NWSSTGs simultaneously. With the i5/OS V5R4 virtual storage architecture enhancements, multiple operations were also dispatched to each NWSSTGs. Accordingly, increasing the number of NWSDs (and Virtual SCSI server/client adapter pairs) provided a minimal boost of about 4% (78.8 → 80.7 MB/s).

This test clearly illustrated the performance benefits of the i5/OS V5R4 virtual storage architecture. For i5/OS environments that cannot be upgraded to i5/OS V5R4 yet, the recommendation is to increase the number of VSCSI client and server adapter pairs in AIX 5L and i5/OS, which would also result in a greater number of NWSDs. At i5/OS V5R4, most environments can use the simpler setup of 1 NWSD and several large NWSSTGs. The test demonstrated a large improvement in performance at V5R4 with a single NWSSTG, as well, but is extremely rare to use a single large disk for the operating system and applications in AIX 5L.

5.2 Scenario 2: One vs. Three AIX 5L LPARs Hosted in the System ASP

Virtual Storage Configuration	i5/OS V5R3
10 NWSDs, 10 NWSSTGs x 10 GB, 1 AIX 5L LPAR, ASP 1	61.6 MB/s
10 NWSDs, 10 NWSSTGs x 10 GB, 3 AIX 5L LPARs, ASP 1	33.1 MB/s <i>in each LPAR</i>

In this test, 8 IOzone tasks wrote and read a different 8-GB file in 1-MB blocks to the same filesystem in one and then three AIX 5L LPARs. In the second case, the NWSSTGs for all three AIX 5L LPARs were created using the same RAID 5 set of 10 drives (see Section 3 for a detailed description). When i5/OS hosted virtual storage for three AIX 5L LPARs in the System ASP, performance in each one decreased by 46% (61.6 → 33.1 MB/s). This test supports the existing rule of thumb that hosting more than one AIX 5L LPAR in the same ASP in i5/OS, using the same physical drives, can decrease performance

5.3 Scenario 3: Three AIX 5L LPARs in the System ASP vs. Separate ASPs

Virtual Storage Configuration	i5/OS V5R3
10 NWSDs, 10 NWSSTGs x 10 GB, 3 AIX 5L LPARs, ASP 1	33.1 MB/s in each LPAR
10 NWSDs, 10 NWSSTGs x 10 GB, 3 AIX 5L LPARs, 3 ASPs	59 MB/s in each LPAR

In this test, 8 IOzone tasks wrote and read a different 8-GB file in 1-MB blocks to the same filesystem in three AIX 5L LPARs. In the first test, all three AIX 5L LPARs were hosted in the System ASP. In the second test, each was hosted in a separate user ASP, using a RAID 5 set of 10 drives. Hosting each AIX 5L LPAR in a separate ASP provided a performance boost of 78% (33.1 → 59 MB/s). This test supports the existing recommendation to configure virtual storage for each hosted AIX 5L LPAR in a separate ASP. Notice that the result for 3 AIX 5L LPARs in separate ASPs (59 MB/s) is very close to that for a single AIX 5L LPAR from Scenario 2 (61 MB/s). This also shows that adding a hosted AIX 5L LPAR does not significantly reduce performance in existing hosted LPARs if the new virtual storage is isolated in a separate ASP.

5.4 Scenario 4: Larger Block or Record Sizes

Virtual Storage Configuration	i5/OS V5R3
6 NWSDs, 6 NWSSTGs x 10 GB, 1 task, record size = 4KB	15.8 MB/s
6 NWSDs, 6 NWSSTGs x 10 GB, 1 task, record size = 16KB	17.9 MB/s
6 NWSDs, 6 NWSSTGs x 10 GB, 1 task, record size = 64KB	18.1 MB/s
6 NWSDs, 6 NWSSTGs x 10 GB, 1 task, record size = 256KB	41.4 MB/s
6 NWSDs, 6 NWSSTGs x 10 GB, 1 task, record size = 1MB	43.5 MB/s

In this test, a single IOzone task wrote and read an 8-GB file using different record sizes to a filesystem consisting of 6 NWSSTGs. As in most performance studies, larger block or record sizes yielded better performance results, because fewer I/O operations were necessary to process the same amount of data. Similar results were obtained when the filesystem contained between 1 and 5 NWSSTGs.

This example also illustrates a trend that was observed in the course of testing. Notice that the throughput number with 6 NWSDs, 6 NWSSTGs and a record size of 1 MB is only slightly higher than that with 1 NWSD and 1 NWSSTG in Section 5.1. The difference between the two cases is the disk utilization in AIX 5L (not shown). In this

case, a single IOzone task kept the 6 NWSSTGs comprising the filesystem at only 15-20% utilized, whereas with 8 IOzone tasks and a single NWSSTG (Section 5.1), the disk utilization was nearly 100%. As the detailed results in Appendix B show, a single IOzone task was not enough to keep the NWSSTGs busy and I/O throughput *fell* as the number of NWSSTGs increased. That is the reason why the main three scenarios (Sections 5.1, 5.2 and 5.3) were tested with 8 simultaneous IOzone tasks. Note that “utilization” in this case means how *busy* the disk units were, and not how *full* they were.

5.5 Scenario 5: Higher Number of Simultaneous IOzone Tasks

Virtual Storage Configuration	i5/OS V5R4
1 NWSD, 10 NWSSTGs x 10 GB, 8 tasks	78.8 MB/s
1 NWSD, 10 NWSSTGs x 10 GB, 12 tasks	81.1 MB/s
10 NWSDs, 10 NWSSTGs x 10 GB, 8 tasks	80.7 MB/s
10 NWSDs, 10 NWSSTGs x 10 GB, 12 tasks	82.3 MB/s

In this test, 8 or 12 IOzone tasks simultaneously wrote and read an 8-GB file using different record sizes to a filesystem consisting of 10 NWSSTGs. Two different configurations were tested: in the first one all NWSSTGs were attached to a single VSCSI client and server pair in AIX 5L and i5/OS; in the second one, each NWSSTG had a VSCSI connection and NWSD. In both cases, a small but measurable (1%) boost in I/O throughput indicated that higher disk utilization in AIX 5L led to better performance.

6.0 Summary and Suggestions for Future Testing

This series of tests clearly demonstrated the benefits of the improved virtual storage architecture in i5/OS V5R4 when compared to that in i5/OS V5R3. It also provided support for several long-standing rules of thumb regarding virtual storage hosted by i5/OS, such as the recommendation to isolate hosted storage for different AIX 5L LPARs in separate user ASPs. Based on the results in Section 5, the following recommendations are provided for designing and configuring virtual storage solutions on System i5 involving AIX 5L and i5/OS:

- At i5/OS V5R3, increasing the number of Virtual SCSI client and server adapters and NWSDs can lead to better performance. A different way to state this conclusion would be that multiple NWSSTGs should be created in i5/OS and each should have its NWSD and VSCSI connection
- Upgrading from i5/OS V5R3 to V5R4 can significantly increase virtual storage performance even for a single large NWSSTG
- Upgrading from i5/OS V5R3 to V5R4 *and* using multiple NWSSTGs can result in even better performance
- Using a separate NWSD for each NWSSTG is much less important at i5/OS V5R4. Most environments can use the simpler configuration of a single NWSD and multiple NWSSTGs

- Isolating hosted storage for different AIX 5L LPARs in separate user ASPs can significantly improve virtual storage performance in each one
- Larger block or record sizes and higher disk utilization in AIX 5L both can improve performance

Please note again that these test were deliberately designed to restrict AIX 5L in order to focus on testing virtual storage throughput. Therefore, the tests illustrate the differences among several configurations, and not the maximum I/O performance that AIX 5L can achieve on System i5 using virtual storage from i5/OS.

There are other test conditions whose effect would be beneficial to explore. Follow-up studies to this one could examine how partial processors (or shared processor units) in both AIX 5L and i5/OS affect I/O performance; or how the incremental changes in the virtual storage architecture in i5/OS V5R3M5 Licensed Internal Code (LIC) would affect I/O throughput.

Appendix A Detailed Test Methodology

As mentioned in Section 4, each test followed this sequence:

- a) Between 1 and 12 IOzone tasks simultaneously wrote a different 8-GB or 100-GB file to the same filesystem, /perf2, in 4-KB to 1-MB blocks
- b) Sequential write I/O throughput was measured in MB/s
- c) The filesystem was unmounted to eliminate filesystem caching and measure strictly virtual storage throughput.
- d) Next, between 1 and 12 IOzone tasks each read one of the previously written 8-GB or 100-GB files
- e) Sequential read I/O throughput was measured in MB/s.

Steps a) and d) each consisted of 3 consecutive write or read operations – a “warm-up” run and 2 test runs. Each write or read run involved first writing or reading the test files (8 or 100 GB in size), then re-writing and re-reading them. In steps b) and e), only data from the 2 test runs were recorded, then average to arrive at the I/O throughput number for that run. This is an example of a full set of read I/O data for a test case in MB/s:

Virtual Storage Configuration	Read 1	Re-read 1	Read 2	Re-read 2	Read Avg.
10 NWSDs, 10 NWSSTGs, 8 tasks	80.7	80.7	80.3	80.9	80.7

Test Scripts

IOzone can be run from the AIX 5L command line, but in this case, two shell scripts were used to automate testing and ensure consistency – “dotest” and “iotest2.sh.” “dotest” is the script that executes a write or a read test, and it contains the number of test runs after the warm-up run, the log file in which to record the results and the IOzone command line. This is an example line from “dotest” with its preceding comment:

```
# i5/OS V5R4 VSCSI WRITE test with 10 hdisks, 1 NWSD, 2 iterations, 12
processes, s=8GB, r=1MB, dio + rbrw mount options
#./iotest2.sh 2 V5R4_s8G_r1M_t12_10hdisks.1NWSD.write.log "./iozone -
r 1024 -s 8g -i 0 -e -w -t 12 -F /perf2/four /perf2/five /perf2/six
/perf2/seven /perf2/eight /perf2/nine /perf2/ten /perf2/eleven
/perf2/12 /perf2/13 /perf2/14 /perf2/15"
```

In this case, “2” is the number of test runs to perform after the warm up run, “V5R4_s8G_r1M_t12_10hdisks.1NWSD.write.log” is the results log file and “./iozone -r 1024 -s 8g -i 0 -e -w -t 12 -F /perf2/four /perf2/five /perf2/six /perf2/seven /perf2/eight /perf2/nine /perf2/ten /perf2/eleven /perf2/12 /perf2/13 /perf2/14 /perf2/15” is the IOzone command. Notice that the “dotest” example begins by invoking “iotest2.sh.” The purpose of “iotest2.sh” is to take the three input components (number of test runs, log file and IOzone command line), perform a warm-up run, perform the required number of test runs and calculate how long the test run takes. This is “iotest2.sh:”

```
#!/usr/bin/bash
```

```

echo "IOtest:"
echo "  Iterations: $1"
echo "  Log File:   $2"
echo "  Command:    $3"
echo ""
rm $2;
echo "***Warmup."
( echo $3 | sed -e 's;/\&/g' -e 's/$/\&/g'; echo wait) | bash ;
echo "***Start Iterations."
(for ((x=1;x<=$1;x++)); do echo $3 | sed -e 's;/\&/g' -e 's/$/\&/g';
echo wait; done) | timex -s bash >> $2 2>&1
echo "***Done."
exit

```

In the script, “command” is the IOzone command. It is passed to *sed* (the stream editor) for the purpose of possibly executing multiple IOzone commands, then it is passed to *timex* for time calculation and to *bash* for execution. After the warm-up run, the process repeats for each test run, or “iteration” – “(for ((x=1;x<=\$1;x++)); do.” For each test case, “dotest” was executed twice: once for the write portion and the second time for the read portion.

IOzone Command Line

The following two IOzone commands constitute a test case example:

```

iozone -r 1024 -s 8g -i 0 -e -w -t 12 -F /perf2/four /perf2/five
/perf2/six /perf2/seven /perf2/eight /perf2/nine /perf2/ten
/perf2/eleven /perf2/12 /perf2/13 /perf2/14 /perf2/15"

```

```

iozone -r 1024 -s 8g -i 1 -e -w -t 12 -F /perf2/four /perf2/five
/perf2/six /perf2/seven /perf2/eight /perf2/nine /perf2/ten
/perf2/eleven /perf2/12 /perf2/13 /perf2/14 /perf2/15"

```

They are identical except for the “-i” option, which specifies the type of test to execute: “0” indicates a sequential write test, while “1” indicates a sequential read test. The two IOzone commands must be executed in this order, because the read test uses the files created by the write test. IOzone creates the files during the write test, giving them the names specified after the “-F” option. As mentioned previously, the filesystem “/perf2” was unmounted between each write and read test to ensure that the files were not cached. A brief explanation of the other IOzone options used: “-r” and “-s” determine the test file size and the block or record size IOzone should use; “-e” specifies that the files should be flushed to disk upon creation; “-w” requires that the test files be left in the filesystem, instead of deleting them upon command completion; and “-t” causes IOzone to run in “throughput” mode, with multiple tasks writing to or reading from the same filesystem. The IOzone manual is available at no cost at <http://www.iozone.org>.

Note: in order to run a script that creates files larger than 1 GB as the administrator user, “root,” the file size limit for root must be changed. That limit is specified in the file /etc/security/limits.

Appendix B Complete Test Results

The following tables show the complete set of results by test case, in MB/s:

Comparing fewer vs. more NWSSTGs and NWSDs, 1 task					
Test Cases for i5/OS V5R3	Read 1	Re-read 1	Read 2	Re-read 2	Read Avg.
1 NWSD, 1 NWSSTG x 10 GB, 1 job, record size = 4KB	55	54.9	54.9	54.7	54.88
1 NWSD, 1 NWSSTG x 10 GB, 1 job, record size = 16KB	107.5	107.7	108	107.5	107.68
1 NWSD, 1 NWSSTG x 10 GB, 1 job, record size = 64KB	141.2	141.1	138.6	136.4	139.32
1 NWSD, 1 NWSSTG x 10 GB, 1 job, record size = 256KB	161.6	161.1	162	162	161.68
1 NWSD, 1 NWSSTG x 10 GB, 1 job, record size = 1MB	169.9	168.6	169	169.4	169.23
2 NWSDs, 2 NWSSTGs x 10 GB, 1 job, record size = 4KB	22.6	22.6	22.9	23	22.78
2 NWSDs, 2 NWSSTGs x 10 GB, 1 job, record size = 16KB	31.2	31.1	31.3	31.4	31.25
2 NWSDs, 2 NWSSTGs x 10 GB, 1 job, record size = 64KB	27.3	27.2	27.3	27.3	27.28
2 NWSDs, 2 NWSSTGs x 10 GB, 1 job, record size = 256KB	56	56	55.9	56.1	56
2 NWSDs, 2 NWSSTGs x 10 GB, 1 job, record size = 1MB	57.1	56.8	57.1	56.7	56.93
3 NWSDs, 3 NWSSTGs x 10 GB, 1 job, record size = 4KB	22	22.2	22.4	22.6	22.3
3 NWSDs, 3 NWSSTGs x 10 GB, 1 job, record size = 16KB	29.9	30	30	30	29.98
3 NWSDs, 3 NWSSTGs x 10 GB, 1 job, record size = 64KB	26.9	27	26.9	26.9	26.93
3 NWSDs, 3 NWSSTGs x 10 GB, 1 job, record size = 256KB	55	55	55	55	55
3 NWSDs, 3 NWSSTGs x 10 GB, 1 job, record size = 1MB	55.9	56	55.8	56.1	55.95
... Throughput diminishes because of low disk utilization ...					
6 NWSDs, 6 NWSSTGs x 10 GB, 1 job, record size = 4KB	15.6	15.7	15.8	15.9	15.75
6 NWSDs, 6 NWSSTGs x 10 GB, 1 job, record size = 16KB	17.9	17.9	17.8	18.1	17.93
6 NWSDs, 6 NWSSTGs x 10 GB, 1 job, record size = 64KB	18	18.4	18.3	18	18.18
6 NWSDs, 6 NWSSTGs x 10 GB, 1 job, record size = 256KB	42.2	42.8	40.9	39.7	41.4
6 NWSDs, 6 NWSSTGs x 10 GB, 1 job, record size = 1MB	43.6	43.9	43.9	42.7	43.53

Comparing fewer vs. more NWSSTGs and NWSDs, 8 tasks					
Test Cases for i5/OS V5R3	Read 1	Re-read 1	Read 2	Re-read 2	Read Avg.
1 NWSD, 1 NWSSTG x 100 GB, record size = 1MB	42.4	42.3	41.3	38	41
1 NWSD, 10 NWSSTGs x 10 GB, record size = 1MB	33.8	30.5	32.7	31.9	32.23
10 NWSD, 10 NWSSTGs x 10 GB, record size = 1MB	60.2	60	61	65.1	61.58

Comparing i5/OS hosting 1 AIX 5L LPAR vs. 3 in ASP 1					
Test Cases for i5/OS V5R3	Read 1	Re-read 1	Read 2	Re-read 2	Read Avg.
10 NWSD, 10 NWSSTGs x 10 GB, r. size = 1MB, 1 LPAR	60.2	60	61	65.1	61.58
10 NWSD, 10 NWSSTGs x 10 GB, r. size = 1MB, 3 LPARs	37.5	34.7	29.8	30.5	33.13

Comparing i5/OS hosting 3 AIX 5L LPARs in ASP 1 vs. separate ASPs					
Test Cases for i5/OS V5R3	Read 1	Re-read 1	Read 2	Re-read 2	Read Avg.
10 NWSD, 10 NWSSTGs x 10 GB, r. size = 1MB, 3 LPARs in ASP 1	37.5	34.7	29.8	30.5	33.13
10 NWSD, 10 NWSSTGs x 10 GB, r. size = 1MB, 3 LPARs in sep. ASPs	58.7	58.2	57.9	61.3	59.03

Comparing fewer vs. more NWSDs and NWSSTG, and more fewer vs. more tasks					
Test Cases for i5/OS V5R4	Read 1	Re-read 1	Read 2	Re-read 2	Read Avg.
1 NWSD, 1 NWSSTGs, 8 tasks	64.1	68.8	69.6	65.1	66.9
1 NWSD, 1 NWSSTGs, 12 tasks	65.6	62	61.9	62.7	63.1
1 NWSD, 10 NWSSTGs, 8 tasks	75.8	78.8	80.7	79.9	78.8
1 NWSD, 10 NWSSTGs, 12 tasks	80.8	80.8	81.2	81.5	81.1
10 NWSDs, 10 NWSSTGs, 8 tasks	80.7	80.7	80.3	80.9	80.7
10 NWSDs, 10 NWSSTGs, 12 tasks	82.7	82	81.9	82.4	82.3

For More Information

To learn more about the IBM System i5 contact your IBM representative, IBM Business Partner or visit the following Web sites:

<http://www-03.ibm.com/systems/i/>
<http://www-1.ibm.com/servers/eserver/series/aix/>

©Copyright IBM Corporation 2005

IBM Systems Group
Route 100
Somers, NY 10589

Produced in the United States of America
June 2005
All Rights Reserved

IBM, the IBM logo, the e-business logo, AIX 5L, Domino, i5, i5/OS, iSeries, Lotus, Operating System/400, POWER, POWER5, Virtualization Engine, WebSphere and xSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows and Windows Server are trademarks of Microsoft Corporation in the United States, other countries or both.

Other company, product and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, IBM warranty terms apply.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

NOTES:

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

Information is provided “AS IS” without warranty of any kind.