Improving database consolidation efficiency using the HP ProLiant DL385 G7 server with low voltage DDR3 memory and Solid State Drives

Technical white paper

Table of contents

Executive summary

As processor, memory, and disk technology have improved, HP ProLiant G7 servers have become ideal platforms for consolidating applications with virtual machines. In particular, the computing power of multi-core processors, large DDR3 low power memory capacity, and the performance advantages of HP 3G SATA hot plug solid state disks make a compelling case for consolidating a number of database instances onto a single virtualized host system. This paper describes tests to characterize database I/O that were done by HP in collaboration with Samsung on two identical HP ProLiant DL385 G7 servers – one with HP 3G SATA solid state disks based on Samsung Green SSD technology, and one with standard 15K RPM enterprise SAS drives. The examples describe tests done with a Microsoft® SQL Server I/O simulator tool, but the techniques and results likely would be similar for other database management software. The results of the tests clearly demonstrate the superior performance of servers with HP 3G SATA solid state disks over servers with traditional rotating media drives.

Target audience: This paper is intended for anyone with an interest in consolidating database instances into virtual machines that use a very high performance storage subsystem on the host server.

The HP ProLiant DL385 G7 server

The DL385 G7 server accommodates up to two AMD Opteron 6100 series processors (formerly code-named Magny-Cours) with up to 12 cores per processor. It has 24 memory slots that can support up to 256 GB of DDR3 memory. Internal direct attached storage capability is provided by 8 hot plug drive slots for 2.5" SFF drives, with the ability to accommodate up to 16 total drives when an optional drive expander cage is added.

HP Qualified Options – Low voltage memory

HP Qualified memory options using Samsung Green DDR3 technology bring new levels of performance to HP ProLiant servers and push the envelope in key areas like lower power consumption, higher speed and bandwidth. Low voltage 1.35V DDR3 memory technology provides up to a 73% power savings over traditional DDR2 at 2x the bandwidth, resulting in power savings of up to 38% at the server level[1]. Figure 1 is a graph provided by Samsung that illustrates the progression of power savings that can be achieved when upgrading to systems using HP low voltage DDR3 memory technology.

[1] http://www.samsung.com/global/business/semiconductor/Greenmemory/images/download/greenpartners/03_1_1_Green%20Partners_Enterprise_HP/01/hp_thermal_logic.pdf

Figure 1. Power savings with HP Qualified memory based on Samsung Green DDR3

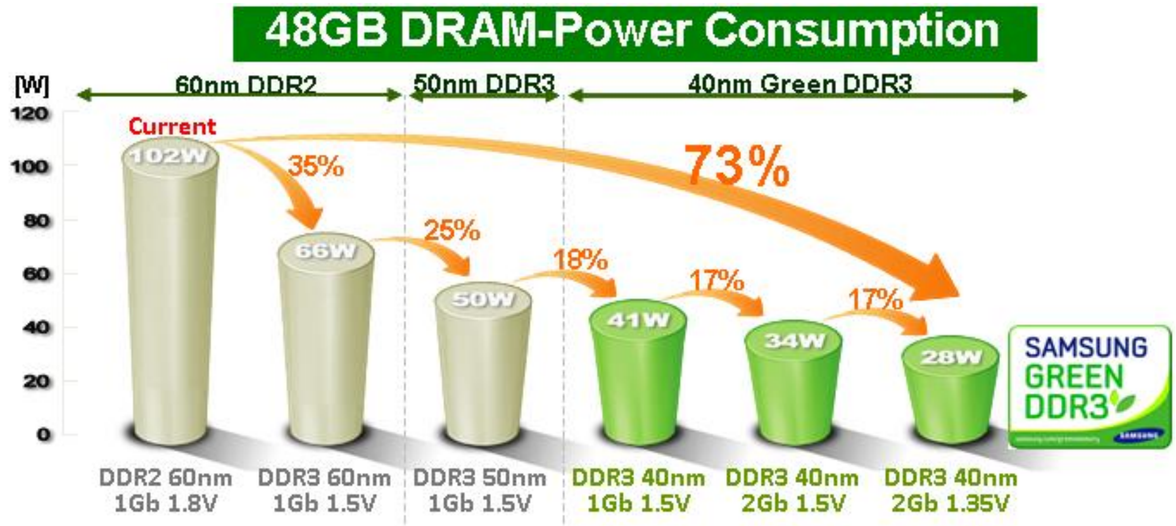

HP 3G SATA solid state drives

HP 3G solid state SATA drives for ProLiant servers using Samsung Green SSD technology feature a 3 Gb/second SATA interface and reliable NAND Single Level Cell flash technology. They are available in both 60GB and 120GB capacities as SFF and LFF hot plug devices on the HP universal drive carrier for general use across the ProLiant server portfolio. These drives deliver higher performance, lower latency and lower power consumption when compared with traditional rotating media. They are ideal for server applications that typically over-provision rotating media spindle count to achieve required performance levels, providing up to 3X the performance of 15K SAS drives with 1/3 the number of drives, resulting in lower overall cost, power and TCO. Figure 2 is a graph provided by Samsung that depicts the IOPS/Watt performance advantages of HP 3G SATA solid state drives.

Figure 2. IOPS/Watt advantage of HP 3G SATA SSDs based on Samsung Green technology

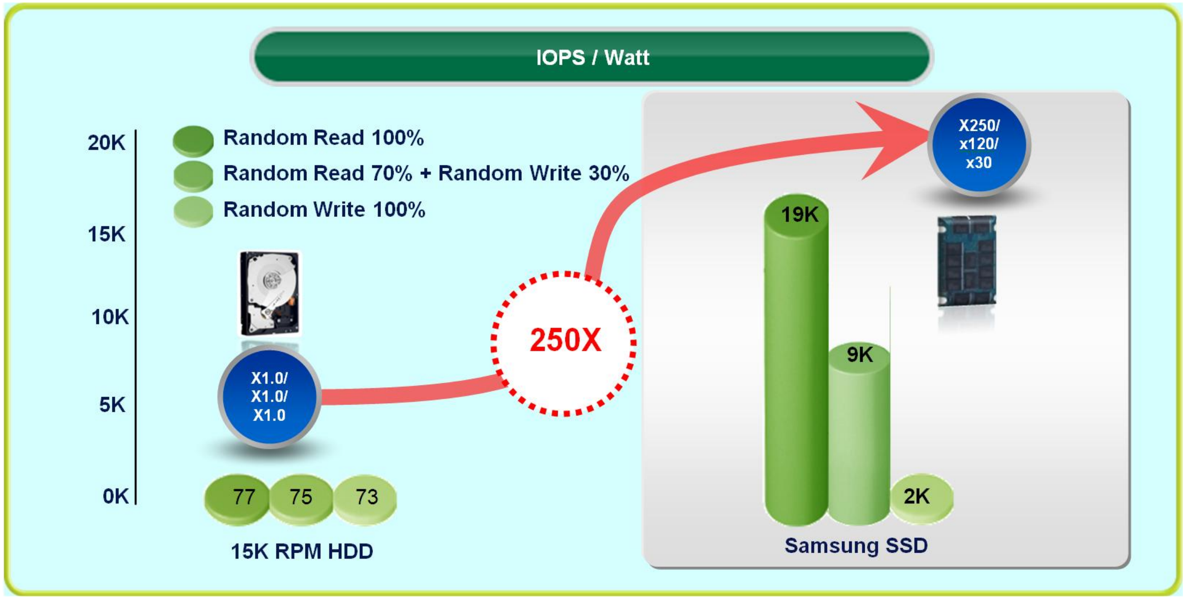

Note

- All SSD test results are based on Samsung measurements. Performance results can vary under different test conditions and firmware versions.
- 15K RPM HDD power consumption level is based on available specs.
- Test Program- IOMETER with queue depth 32

The case for virtualized database consolidation

Over time, many companies accumulate multiple database instances on a number of physical servers. This presents both a system management challenge, and space, power, and cooling challenges. Many IT groups have elected to consolidate multiple databases onto a smaller number of physical systems by using virtual machine technology, such as VMware vSphere 4.

The DL385 G7 server provides an ideal host platform for this based on its combination of powerful CPUs, large memory footprint, and large disk complement. The 256 GB memory capacity of this server allows you to provision each virtual database instance with a large amount of dedicated memory for higher database performance. For example, you could consolidate up to 12 database servers currently deployed on older dedicated servers into virtual machines on a single server, providing each with 2 virtual CPUs, more than 20 GB of memory, and 80 GB of mirrored disk storage[2].

Database applications often impose a high random access workload on the storage system that can be largely read biased. The very high random read/write performance advantages of HP 3G SATA solid state disks are an ideal match for database applications with these characteristics.

[2] Presumes use of the optional drive expander cage with a total of sixteen 120 GB SSD drives configured as a RAID 1+0 set.

Table 1 depicts an example derived from statistics generated by the HP Sizing Tool for VMware vSphere and the HP BladeSystem Power Sizer that reflects the rack space, power, and cooling benefits that could be achieved from such a consolidation.

Table 1. Example of savings with database consolidation on a DL385 G7 server with VMware ESX

	Old non-virtualized	**New virtualized, Green**
Server model	DL385 G5	DL385 G7
Processors (2 CPUs)	AMD Opteron 2352	AMD Opteron 6174
Memory	32 GB DDR2	256 GB LV-RDIMM DDR3
Disk	8x146 GB 15K RPM SAS	8x120 GB SSD
Number of Servers	5	1
Total Rack Space	10U	2U
Total Power, 75% workload	2170 Watts	667 Watts
Total BTU, 75% workload	7550 BTU/hr	2275 BTU/Hr
Average power savings[3]		$1320/year
Average cooling savings		$1050/year

Comparison test of SSD vs. HDD performance

To quantify the actual system performance advantages of a system with HP 3G SATA SSDs, two servers were configured identically with two AMD Opteron 6174 twelve core processors and 256 GB of HP Qualified memory using Samsung Green DDR3 technology. Each server contained eight 16 GB RDIMMs and sixteen 8 GB RDIMMs to fully populate the 24 available memory slots and achieve the maximum capacity of 256 GB. Memory was a mixture of standard voltage (1.5V) and low voltage (1.35V) DDR3 RDIMMs running at 1.5V with a memory speed setting of 800 MHz[4]. One of these servers was populated with eight 120 GB HP hot plug 3G solid state drives configured as one large RAID 1+0 volume using the integrated HP Smart Array P410i disk controller. The other was populated with eight 146 GB HP hot plug 15K RPM enterprise SAS drives configured in the same way.

VMware ESX 4.0 was installed on each server, and four virtual machines were configured on each server as shown in Table 2. A sample database was installed on each virtual machine data drive.

Table 2. VMware ESX virtual machine configuration for database hosting

Virtual CPUs	2
Virtual memory	32 GB
Virtual operating system drive (C:\ )	40 GB
Virtual data drive (E:\)	40 GB
Operating system	Win2K3 Enterprise Edition

[3] Power and cooling savings based on 24 hour, 365 day operation at an average power cost of $0.10/KWh.

[4] When a mixture of standard and low voltage RDIMMs are used, memory speed can be set to 800 MHz by system BIOS and all memory runs at 1.5V. When all slots are populated with low voltage 1.35V DIMMs, memory speed is reduced to 667 MHz, but power savings are even greater than depicted in Table 1 because all memory runs at 1.35V.

To simulate the database access patterns of a Microsoft SQL Server database application, a copy of the Microsoft SQLIO tool was installed in each of the four VMs on each server. This command line tool provides a characterization of the server storage subsystem to help determine the overall performance capabilities in a SQL Server application environment. For the purpose of this demonstration, a random read test was configured with 8 KB block sizes typical of a transaction processing application. The SQLIO tool runs a series of tests at various disk I/O queue depths to determine when the storage system reaches saturation. The output of the test was directed to a text file for later analysis with Microsoft Office Excel. Figure 3 shows the command lines in the batch file used.

Figure 3. SQLIO batch file for 8 KB random reads

```
sqlio -kR -t2 -s120 -dE -o1 -frandom -b8 -BH -LS Testfile.dat
sqlio -kR -t2 -s120 -dE -o2 -frandom -b8 -BH -LS Testfile.dat
sqlio -kR -t2 -s120 -dE -o4 -frandom -b8 -BH -LS Testfile.dat
sqlio -kR -t2 -s120 -dE -o8 -frandom -b8 -BH -LS Testfile.dat
sqlio -kR -t2 -s120 -dE -o16 -frandom -b8 -BH -LS Testfile.dat
sqlio -kR -t2 -s120 -dE -o32 -frandom -b8 -BH -LS Testfile.dat
```

For the command lines in this file:

-kR	configures a read test
-t2	configure two threads (one per vCPU)
-s120	configures a 2 minute test period
-dE	configures a test on virtual drive E:\
-o(n)	configures outstanding I/O requests (1-32)
-frandom	configures random I/O
-b8	configures an 8 KB block size
-BH	configures disk drive hardware cache buffering only
-LS	configures capture of latency information
Testfile.dat	is the target database file

Figure 4 shows the sample output from the command with queue depth 1 followed by the execution command for queue depth 2 (the first two lines in the batch file) for the SSD system.

Figure 4. Execution command and output

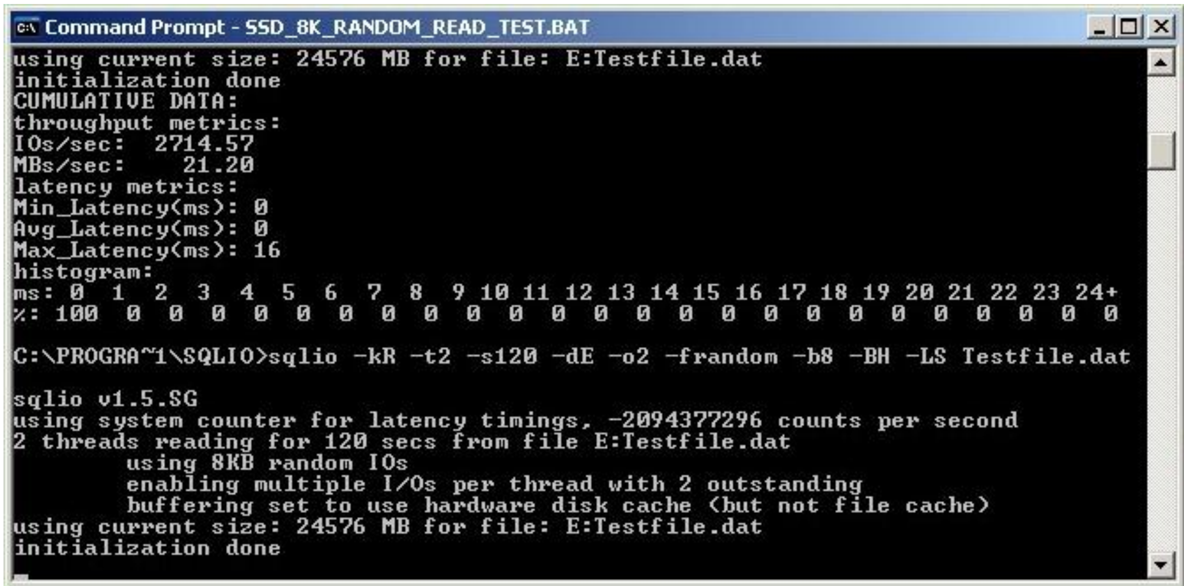

```
Command Prompt - SSD_8K_RANDOM_READ_TEST.BAT
using current size: 24576 MB for file: E:Testfile.dat
initialization done
CUMULATIVE DATA:
throughput metrics:
IOs/sec:  2714.57
MBs/sec:    21.20
latency metrics:
Min_Latency(ms): 0
Avg_Latency(ms): 0
Max_Latency(ms): 16
histogram:
ms: 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24+
%: 100  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0

C:\PROGRA~1\SQLIO>sqlio -kR -t2 -s120 -dE -o2 -frandom -b8 -BH -LS Testfile.dat

sqlio v1.5.SG
using system counter for latency timings, -2094377296 counts per second
2 threads reading for 120 secs from file E:Testfile.dat
        using 8KB random IOs
        enabling multiple I/Os per thread with 2 outstanding
        buffering set to use hardware disk cache (but not file cache)
using current size: 24576 MB for file: E:Testfile.dat
initialization done
```

Test results

Figure 5 shows a summary of the test results obtained with the test described above. You can see that the 4 VMs on the system with SSDs consistently produced a random read rate starting at an average of about 2300 I/Os per second at a queue depth of one, rising to a maximum of nearly 8000 I/Os per second at a queue depth of 8. At greater queue depths, the storage system is saturated and no further performance is possible, but the SSD system continues to support almost 8000 I/Os per second even at a queue depth of 32.

By contrast, the 4 VMs on the system with high performance rotating media produce a maximum random read rate of about 500 I/Os per second at a queue depth of 2, but beyond that the system collapses and is unable to support any workload at all.

Figure 5. I/O per second results for SSD vs. HDD systems

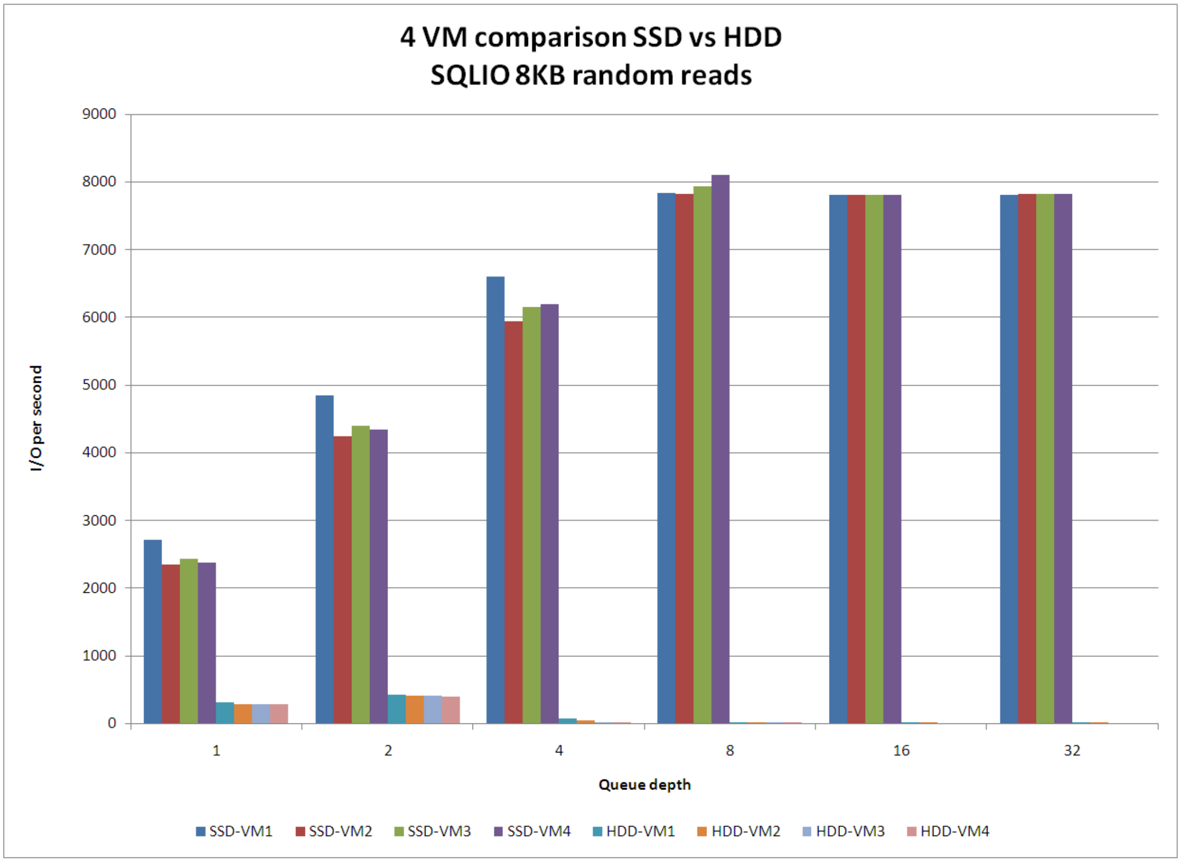

Figure 6 shows the corresponding latency results for the 8 KB random read test on a log scale. Here you can see that the virtual disks in the HDD system exhibit an initial latency of about 6 ms, rising to over 8 ms at a queue depth of 2. At a queue depth of 4, the latencies of the virtual disks in the 4 VMs begin to diverge, ranging from below 100 ms to over 1 second. Comparing this with the I/Os per second curve, this was the point where the system was unable to sustain the SQLIO workload. At a queue depth of 8, the latencies skyrocket into multiple seconds and were not plotted beyond this point.

By contrast, the virtual disks in the SSD system exhibit no noticeable latency up to a queue depth of 16, beyond which they simply become comparable to those of the virtual disks on the HDD system. These throughput and latency plots clearly demonstrate the superior performance of SSDs in Microsoft SQL Server database applications that are biased toward high random read workloads. Although not included in this paper, comparable results were obtained with a workload based on 64 KB block sizes.

In the past, performance of the HDD system would have been increased by provisioning a large number of rotating media drives to obtain performance levels comparable to the SSD system used in this test, even though capacity of the rotating media drives would be greatly underutilized. However, this significantly increased system cost, power, cooling, and rack space considerations. These tests demonstrate that HP 3G SATA solid state drives provide a compelling new alternative that can achieve a lower total cost of ownership than traditional rotating media drives for database applications that require high performance at realistic capacity utilization.

Figure 6. Latency results for SSD vs. HDD systems

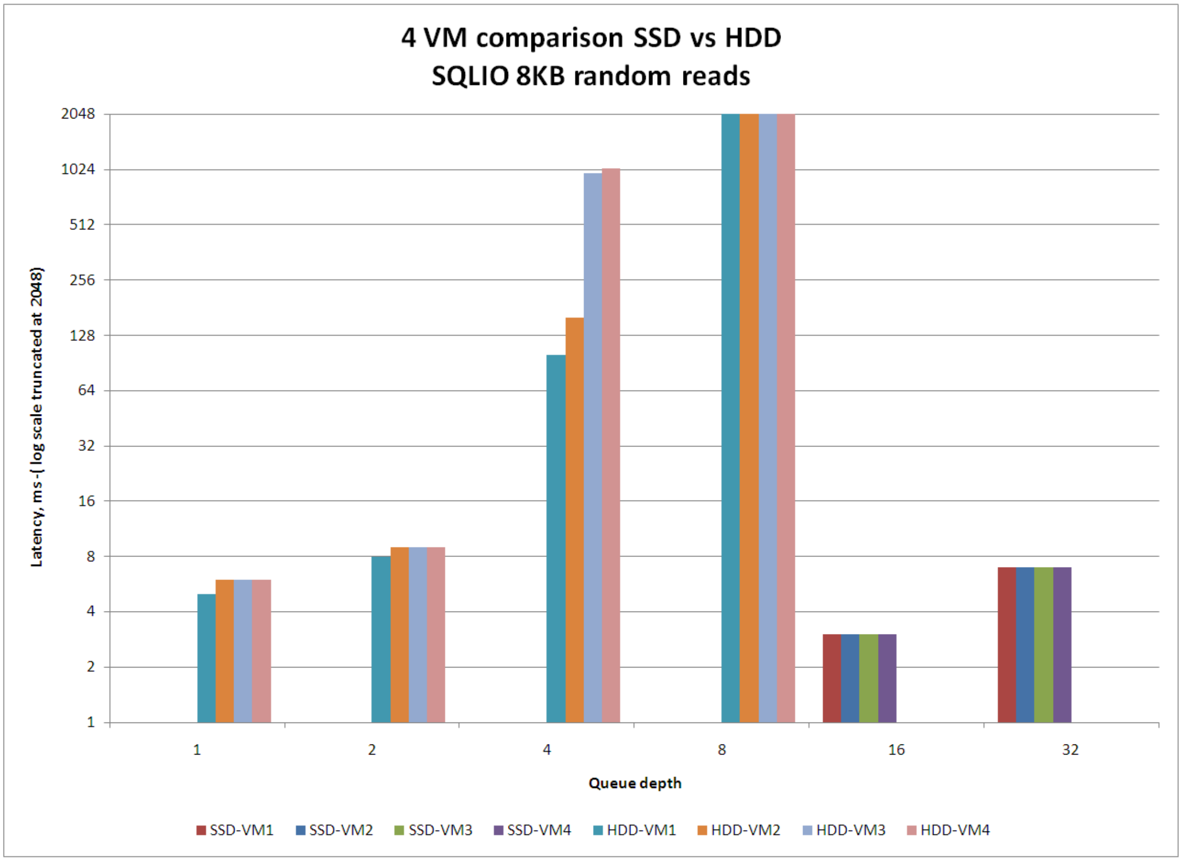

An alternate view of results

As impressive as the graphs are, a more compelling real time display based on Microsoft Windows® Performance Monitor (perfmon) was used to demonstrate the test results at VMworld 2010. In this scenario, a single perfmon window was configured to show the relative performance of a single virtual machine from the SSD system and the corresponding VM from the HDD system side by side while the SQLIO tests were being run.

Figure 7 shows a screen capture of the perfmon window displaying disk I/Os per second from VM1 on the SSD system next to VM1 from the HDD system for the queue depth 2 random read test. Here you can see the same comparative performance as shown in the graphs in Figure 5.

Figure 7. Perfmon disk I/Os per second for 8 KB random reads on SSD vs HDD VMs

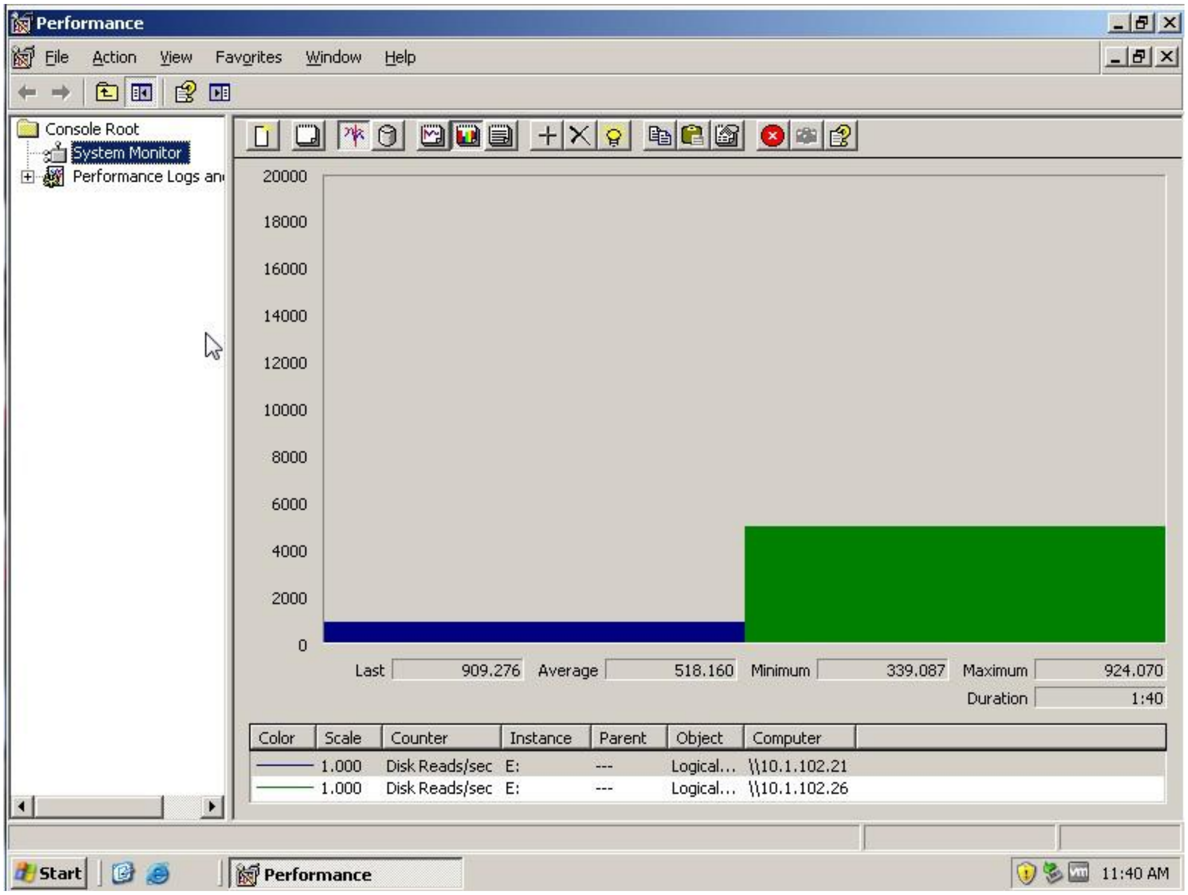

To illustrate that HP 3G SATA SSDs also provide comparable performance improvements for database random write workloads, the batch file shown in Figure 3 was copied and edited to generate an increasing random write workload based on 8 KB block sizes. The corresponding perfmon window comparing performance of a VM on the SSD system with a VM on the HDD system is shown in Figure 8. This clearly shows a comparable performance advantage for the system with SSD storage with random write workloads.

Figure 8. Perfmon disk I/Os per second for 8KB random writes on SSD vs HDD VMs

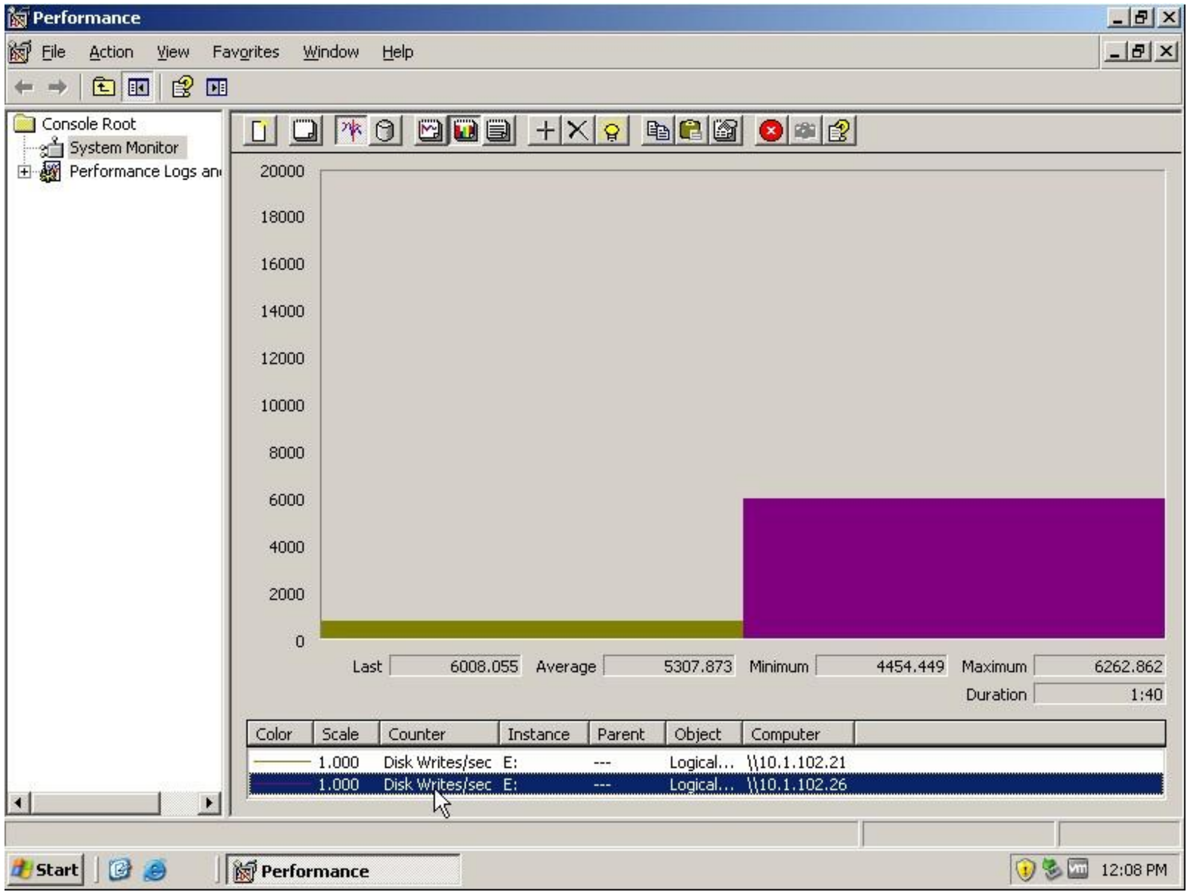

Summary

This paper provides an example of the potential performance improvements and TCO savings that can be achieved by consolidating Microsoft SQL Server databases with VMware ESX on an HP ProLiant DL385 G7 server configured with HP Qualified low voltage memory and HP 3G SATA solid state disks. As always, you should model and test your own unique application environment to quantify the benefits you can achieve.

For more information

To read more about HP and Samsung solutions go to these links:

HP Qualified Options

Samsung Green SSD

Samsung Green DDR3

Video of the VMworld 2010 demonstration

HP solutions for Microsoft SQL Server

HP solutions for VMware ESX

HP and Samsung partnership

To help us improve our documents, please provide feedback at http://h20219.www2.hp.com/ActiveAnswers/us/en/solutions/technical_tools_feedback.html.

4AA3-1779ENW, Created November 2010