Lab Validation Report

IBM System Storage DS5020/DS3950 Express and IBM System x3950 M2

Mixed Workload Performance in Microsoft Hyper-V R2 Environments

By Brian Garrett

May 2010

Contents

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This report was made possible through cooperative testing by Enterprise Strategy Group, IBM System x, IBM System Storage, QLogic and Microsoft. This ESG Lab report was sponsored by IBM.

Introduction

Networked storage is being deployed in conjunction with server virtualization by a growing number of organizations interested in consolidation, reduced costs, and improved flexibility and availability of mission-critical applications including databases and e-mail. ESG research indicates that IT managers looking to reap the benefits of server and storage consolidation are concerned about performance. This ESG Lab report presents the results of a new performance benchmark methodology designed to assess the real-world performance capabilities of a SAN attached IBM System Storage DS5020 Express/DS3950 Express storage system and IBM System x3950 M2 servers deployed in a highly virtualized Microsoft Windows Server 2008 Hyper-V R2 enabled data center.

The Challenges

The use of server virtualization technology is on the rise among organizations of all sizes and in all industries around the world. In a recent ESG survey of current and planned users of the technology, 52% of organizations had already deployed, while 48% plan to do so.[1] Given the impressive economic benefits of server virtualization, the glut of affordable and under-utilized processing power, and growing power and cooling issues in the data center, ESG predicts that the brisk adoption of server virtualization will continue for the foreseeable future.

ESG research indicates that the vast majority (87%) of organizations that have deployed server virtualization have done so in conjunction with networked storage. Compared to islands of direct attached hard drives, utilization is greatly increased when applications share a pool of networked storage. Applications deployed on virtual machines sharing a pool of storage are more mobile and available than those deployed on direct attached hard drives.

While the benefits of server virtualization and networked storage are clearly compelling, IT managers are faced with a number of challenges as they try to manage a consolidated mix of real-world applications running on a virtualized infrastructure. As shown in Figure 1, the top two concerns are performance and a general lack of information and best practices. This holds true across organizations of all sizes, regardless of the number of virtual servers deployed.

Figure 1. Server Virtualization and Networked Storage Challenges

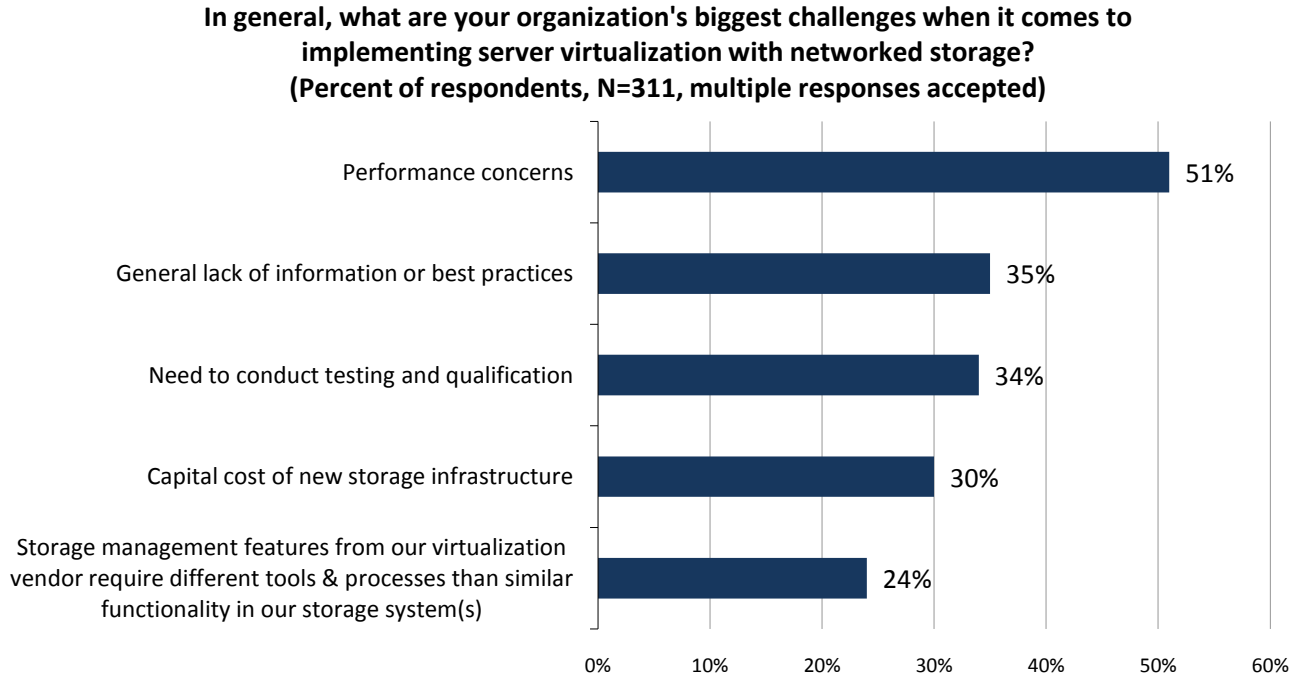

Source: Enterprise Strategy Group, 2007.

[1] Source: ESG Research Report, *The Impact of Server Virtualization on Storage*, December 2007.

The Solution

This ESG Lab report examines the performance of real-world application workloads running in a virtualized and consolidated IT environment that leverages the following technologies:

- **IBM System Storage DS5020 Express and DS3950 Express storage systems:** With high performance that is optimized for mixed workloads, the DS5020 Express and the DS3950 Express were designed for modular scalability (capacity and/or performance), high availability, and advanced functionality including copy services and remote replication.
- **IBM System x3950 M2 servers:** The IBM System x3950 M2 is a rack-optimized server with extraordinary performance and scalability that is ideally suited for virtualized environments.
- **Microsoft Windows Server 2008 R2**: With a built-in and updated Hyper-V platform and Live Migration, organizations can increase the efficiency and flexibility of their server infrastructure with Windows Server 2008 R2.
- **QLogic 8 Gb FC Dual-port HBA for IBM System x:** The QLogic 8 Gb Dual-port HBA for IBM System x is designed to deliver sustained high throughput and reliability in highly available virtual environments.

The capabilities of the IBM servers and storage that were used during this evaluation are summarized in Figure 2. The IBM System x3950 M2 supports up to 96 Intel Xeon 7400 processor cores, 1 TB of DDR-2 memory, and 28 PCI-Express x8 expansion slots. The IBM System Storage DS5020 Express and the DS3950 Express support up to 112 drives (FC, SATA, or mixed configurations) and are equipped with up to 4 GB of cache and 8 GB/sec of internal bandwidth. The DS3950 Express is a variant of the DS5020 Express that is available in two pre-configured models. The DS5020 Express can be custom configured and adds support for full disk encryption (FDE). The DS5020 Express supports up to eight FC host interfaces and the DS3950 Express supports up to four FC host interfaces.

Figure 2. IBM Server and Storage Highlights

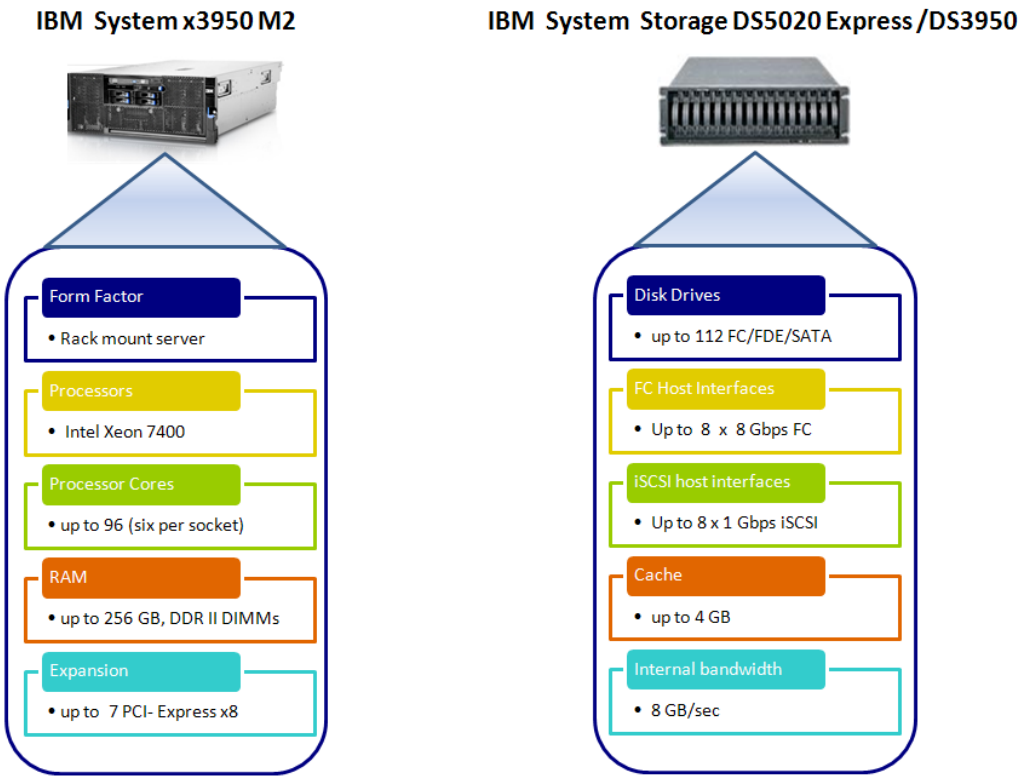

The Results

This report examines the performance capabilities of IBM System Storage DS5020 Express and DS3950 Express storage systems running a mix of real-world applications in a Microsoft Hyper-V R2 enabled virtual server environment powered by a pair of IBM System x3950 M2 servers. In particular, this report explores how:

- A single DS5020 attached to a pair of System x3950 M2 servers running a mix of real-world application workloads in 16 Microsoft Hyper-V R2 enable virtual machines supports up to:
  - **18,750 mailboxes** using the Microsoft Exchange 2010 Jetstress utility
  - *and* **3,475 small database IOs per second** using the Microsoft SQLIO utility
  - *and* **650 MB/sec of database throughput** for large scan operations using the SQLIO utility
  - *and* **3,106 simulated web server IOPs**
  - *and* **413 MB/sec of throughput** for simulated backup/scan/index jobs
  - with the predictably fast response times and scalability
- Within a Microsoft enabled virtual server infrastructure, the DS5020 Express achieved a maximum aggregate throughput of 3.3 GB/sec during bandwidth intensive performance testing and 1.1 GB/sec during mixed application workload testing.

The predictably fast, mixed workload performance scalability of the virtualized environment tested by ESG Lab is summarized in Figure 3. The results will be explored in detail later in this report, but for now it should be noted that the performance of the DS5020 Express scaled extremely well as a mix of real-world application workloads ran in parallel on up to 16 virtual machines.

Figure 3. DS5020 Express Mixed Workload Scalability

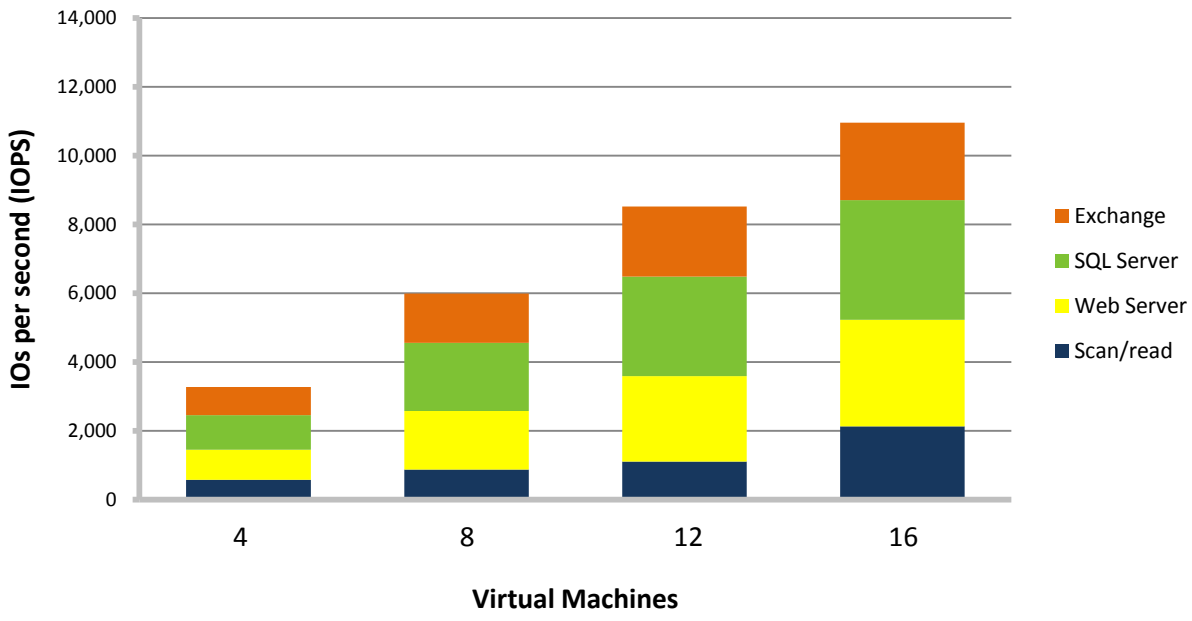

The balance of this report explores how mixed workload testing was accomplished, what the results mean, and why they matter to your business.

ESG Lab Validation

The methodology presented in this report was designed to assess the performance capabilities of a single IBM DS5020 Express storage system shared by multiple Microsoft Hyper-V R2 enabled virtual machines running a mix of real-world application workloads. The cooperation of Microsoft, IBM, and QLogic played a vital role in the success of this project. In particular, this project benefitted from Microsoft's expertise in helping customers plan for the deployment of business-critical applications in consolidated environments and IBM's long heritage of success in the modular storage systems market.

A Mixed Real-world Storage Benchmark Methodology

Conventional server benchmarks were designed to measure the performance of a single application running on a single operating system inside a single physical computer. SPEC CPU2000 and CPU2006 are well known examples of this type of server benchmarking tool. Much like traditional server benchmarks, conventional storage system benchmarks were designed to measure the performance of a single storage system running a single application workload. The SPC-1 benchmark, developed and managed by the Storage Performance Council with IBM playing a key role, is a great example. SPC-1 was designed to assess the performance capabilities of a single storage system as it services an online interactive database application.

The IBM DS5020 used during this ESG Lab Validation has a published SPC-1 result of 26,080 SPC-1 IOPS with a price-performance value of $8.08 per SPC-1 IOP. [2] This is an excellent result for a modular storage system of this class. The DS5020 also has a published Microsoft Exchange Solution Review Program (ESRP) report which indicates that the DS5020 supports up to 28,000 simulated Microsoft Exchange 2007 mailboxes.

While each of these results can be used to predict how a server and storage solution will perform when running a single application workload, neither was designed to help IT managers understand what happens when a mix of applications is deployed together in a virtual server environment.

Compared to a traditional benchmark which tests a single application running on a physical server, ESG Lab tested a mix of virtualized applications running at the same time. The methodology shown in Figure 4 was designed to assess the performance of a mix of application workloads, including database and e-mail applications, running on hypervisor enabled virtual servers sharing a single pool of SAN attached storage.

Figure 4. Traditional Benchmarking vs. Mixed Workload Benchmarking

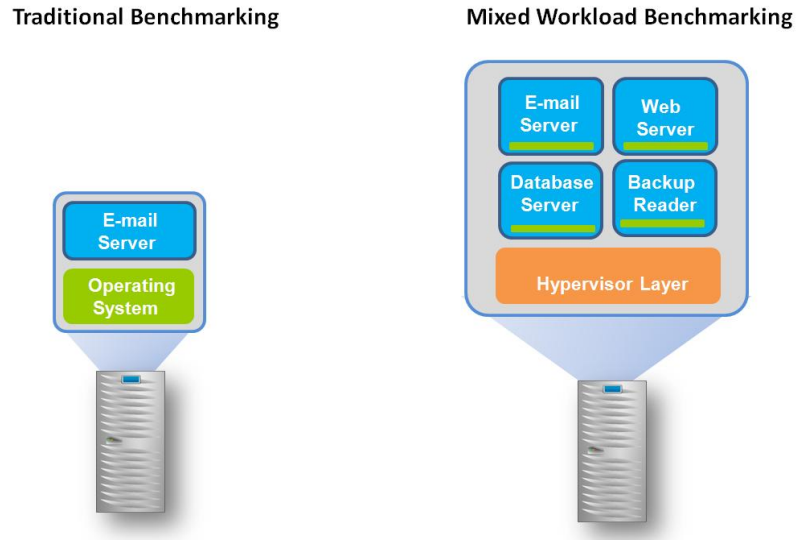

[2] http://www.storageperformance.org/benchmark_results_files/SPC-1/IBM/A00081_IBM_DS5020-Express/a00081_IBM_DS5020-Express_SPC1_executive-summary-r1.pdf

A cell concept was used during ESG Lab mixed workload testing. As shown in Figure 5, each cell is composed of a mixture of four application workloads. Two physical servers, each configured with eight virtual machines, were used to measure performance as the number of active cells was increased from one to four.

Figure 5. ESG Lab mixed Workload Benchmarking

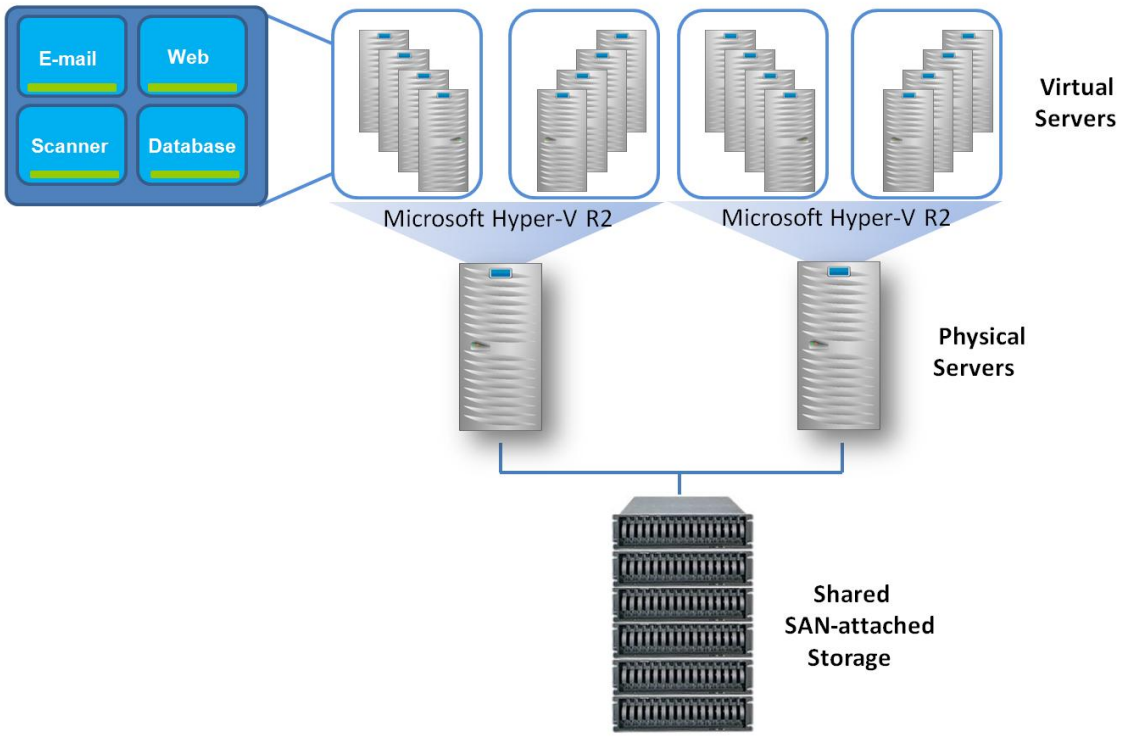

Mixed Workloads

Industry standard benchmarks were used to emulate the IO activity of four common business application workloads:

- **E-Mail:** The Microsoft Jetstress 2010 utility was used to generate e-mail traffic. Jetstress simulates the activity of typical Microsoft Exchange users as they send and read e-mails, make appointments, manage to-do lists, etc. The Jetstress utility is a lightweight tool that is ideally suited for predicting storage performance in a Microsoft Exchange environment. The JetStress utility uses the jet engine database that Microsoft Exchange uses internally. JetStress 2007 is the tool that was used in the IBM DS5020 Microsoft ESRP test that was mentioned previously in this report. ESG Lab tested with JetStress 2010 running in Hyper-V R2 enabled virtual machines running Microsoft Server 2008 R2 as a guest operating system.
- **Database:** The Microsoft sqlio utility was used to generate database traffic. Much like Jetstress, sqlio is a lightweight tool that is ideally suited for measuring storage performance. Sqlio was designed to help administrators understand the performance capabilities of a storage system, either to uncover performance issues or to size a new database installation without having to create and run a SQL Server database. Sqlio was used to simulate two types of database activity: response-time sensitive online transaction processing (OLTP) and bandwidth sensitive decision support systems (DSS).
- **Web Server:** The industry standard Iometer utility was used to generate web server traffic. The IO definition was composed of random reads of various block sizes. The web server Iometer profile used for this test was originally distributed by Intel, the author of Iometer. Iometer has since become an open source project. Iometer tests were performed on Hyper-V R2 enabled virtual hard disks running within virtual machines running Windows 2008 R2 as the guest operating system. Windows NTFS file systems

were not configured for access as a drive letter. Instead, Iometer testing was performed at the underlying Windows physical drive level (e.g. , Physical Drive 3).

- **Scan/read:** The Iometer utility was used to generate a single stream of bandwidth intensive sequential read traffic. Operations that tend to generate this type of large block sequential traffic include scan and index operations, long running data base queries, backup operations, bulk data uploads, and copies. One 256 KB sequential read workload was included in each cell to add a throughput intensive component to the predominantly random IO profile of interactive e-mail, database, and web server applications. As most experienced database and storage administrators have learned, a throughput intensive burst in IO traffic can drag down the performance for interactive applications, causing performance problems for end-users. Adding a few streams of throughput intensive scan/read traffic was used to determine whether interactive performance would remain predictably responsive as the amount of mixed IO utilization increased.

Each of the four workloads ran in parallel for approximately two hours. The Jetstress e-mail test ended with a mandatory database consistency check which completed roughly an hour after the rest of the mixed workloads had completed. The settings for each of the industry standard benchmarks are documented in the appendix.

Why This Matters

ESG research indicates that the top concern when implementing networked storage platforms to support server virtualization is performance. According to 51% of respondents who had already deployed server virtualization and networked storage, performance was by far the top customer concern.

Storage benchmarks have historically focused on one type of workload (e.g., database or e-mail) and one key performance metric (e.g., response time or throughput). Server benchmarks have typically tested only one server running a CPU intensive workload that doesn't stress storage. To help IBM customers understand how a DS5020 Express performs in a Microsoft Hyper-V enabled virtual server environment, this benchmark was designed to assess how real-world applications behave when running on multiple virtual servers sharing a single storage system.

Test Bed

Microsoft Windows Server 2008 R2 was installed on a pair of powerful IBM System x3950 M2 servers, each with four six-core processors and a pair of QLogic 8 Gb Fibre Channel host bus adapters. A DS5020 Express with 112 15K RPM FC drives was connected to the servers through QLogic dual port 8 Gbps host bus adapters and a QLogic FC switch as shown in Figure 6.

Figure 6. ESG Lab Test Bed

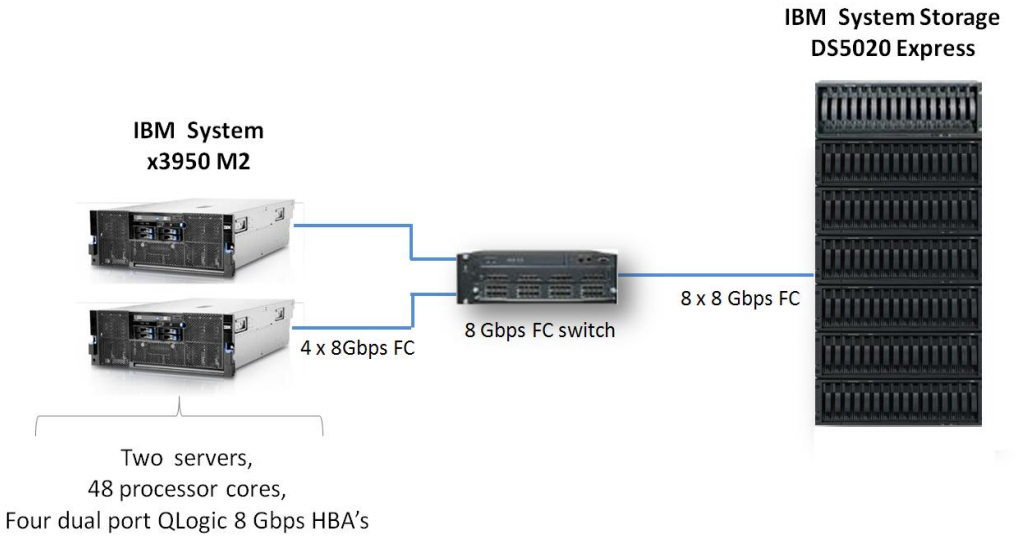

Drive Layout

The DS5020 drive configuration is summarized in Table 1. Four Exchange database volumes were configured. Each of the Exchange database volumes was configured with an eight drive RAID-10 database volume and a four drive RAID-10 log volume. The SQL Server, web server, and scan/read workloads ran against four drive RAID-10 volumes. The guest operating system volumes were configured using a 3+1 RAID-5 layout. Volume ownership was balanced across the dual controllers within the DS5020 Express and distributed evenly over the eight host interfaces. The volumes were configured for access over two redundant host interfaces and accessed through the IBM redundant disk array controller (RDAC) driver. [3]

Table 1: Drive Configuration

Application	Number of LUNs	Number of Drives	Usable Capacity (GB)
Exchange DB	4	32	4,356
Exchange Log	4	16	2,231
SQL Server	4	16	2,231
Web Server	4	16	2,231
Scan/Read	4	16	2,231
Guest O/S	4	16	3,346
Total	24	112	16,626

[3] For more detail, see the Appendix.

Configuring Virtual Machines

Each of the Microsoft Jetstress virtual machines was configured with two logical 2.67 GHz CPU cores and 12 GB of RAM. DS5020 Express disk capacity was used for all storage capacity including the Windows Server 2008 R2 guest operating system, application executables, and application data. All of the application data volumes under test were configured as fixed virtual hard disks (VHDs). The configuration of one of the four virtual machines used for e-mail testing is shown in Figure 7. Note how three hard disks have been configured: one for the guest operating system and two for the Exchange database and logs.

Figure 7. Virtual Server Configuration

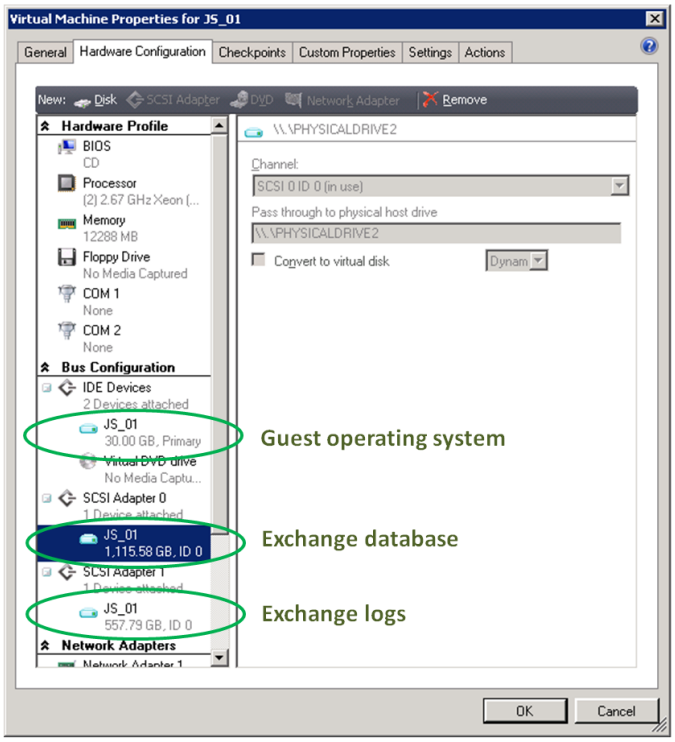

The Results

In a way, storage system benchmark testing is like an analysis of the performance of a car. Specifications including horsepower and acceleration from 0 to 60 are a good first pass indicator of a car's performance. But while specifications provide a good starting point, there are a variety of other factors that should be taken into consideration including the condition of the road, the skill of the driver, and gas mileage ratings. Much like buying a car, a test drive with real-world application traffic is the best way to determine how a storage system will perform in real-world conditions.

Characterization

Performance analysis began with an examination of the low level aggregate throughput capabilities of the test bed. This testing was performed using the Iometer utility running within eight of the virtual machines that were used later during mixed workload testing. The eight virtual machines accessed DS5020 Express storage through eight 8 Gbps FC interfaces.

Iometer access definitions which measured the maximum throughput from disk were used for this first pass analysis of the underlying capabilities of the DS5020 Express. Similar to a dynamometer horsepower rating for a car, maximum throughput was used to quantify the power of a turbo-charged DS5020 Express storage engine. As shown in Figure 8, ESG Lab recorded a maximum throughput of 3.3 GB/sec.

Figure 8. Characterizing the DS5020 Engine

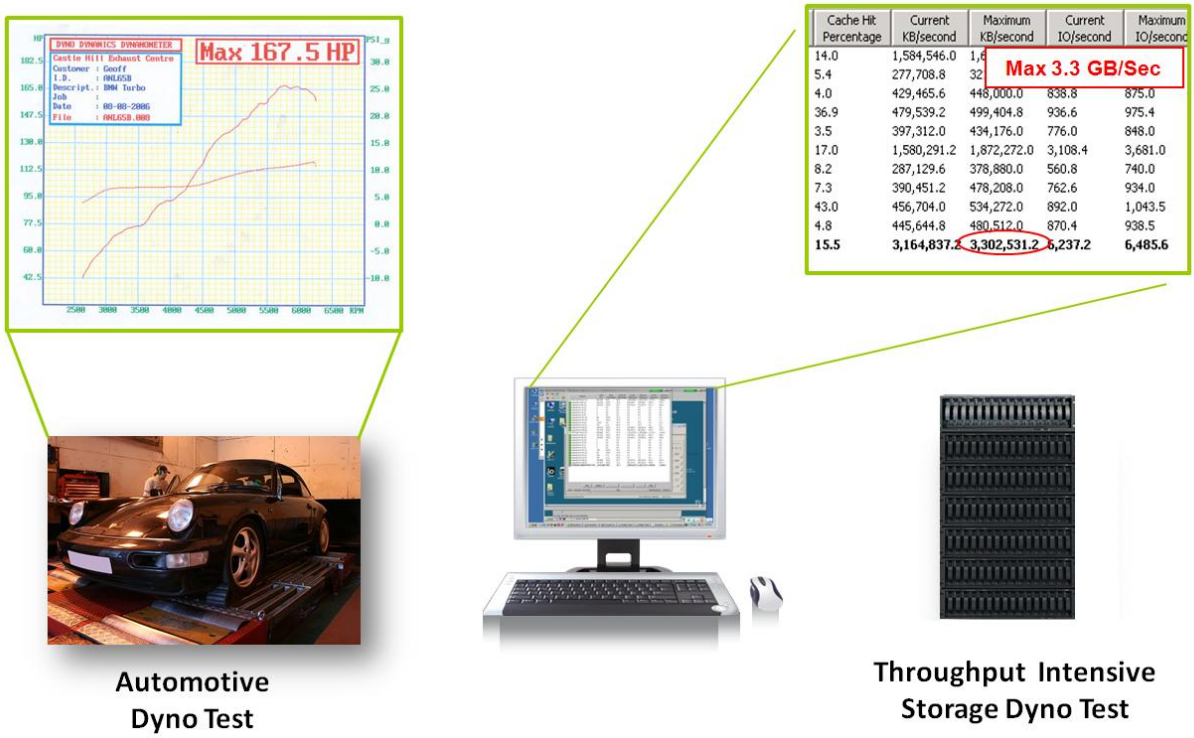

What the Numbers Mean

- Much like the horsepower rating of a car, the throughput rating of a storage system is a good indicator of the power of a storage system's engine.
- Storage throughput is a measure of the bandwidth available to the system. Throughput can be measured on a stream or aggregate basis. A stream is represented by one application or user communicating through one IO interface to one device. Aggregate throughput is a measure of how much data the storage system can move on a whole for all applications and users.
- ESG Lab throughput characterization was performed using the industry standard Iometer utility as 32 streams performed large sequential reads (1 MB) from eight logical devices through eight FC interfaces.
- ESG Lab recorded a peak aggregate throughput of 3.3 GB/sec in a Microsoft Hyper-V R2 enabled virtual server environment.
- Approximately 15% of the throughput was delivered from DS5020 Express cache; the balance was serviced from disk.
- When comparing the performance capabilities of two servers in a virtual server environment, the server with more cache tends to perform better. ESG Lab is confident that a similar pattern holds true for storage systems. A storage system with more cache—and better caching algorithms—should perform better in a virtual server environment.
- ESG Lab characterization testing indicates that the DS5020 Express has more than enough cache and front end bandwidth to meet the needs of virtualized applications requiring up to 112 disk drives for capacity.
- ESG Lab is convinced that the patented caching algorithms of the DS5020 Express provide a significant performance boost during mixed application virtualized application testing.

Why This Matters

A storage system needs a strong engine and well-designed modular architecture to perform predictably in a mixed real-world environment. One measure of the strength of a storage controller engine is its maximum aggregate throughput. ESG Lab testing of the DS5020 Express in a Microsoft Hyper-V R2 environment achieved 3.3 GB/sec of aggregate large block sequential read throughput.

In ESG Lab's experience, these are excellent results for a dual controller modular storage system. As a matter of fact, these results provide an excellent early indication that the DS5020 Express is well suited for virtual server consolidation and mixed real-world business applications.

Virtual Machine Utilization

Mixed application testing began with a quick analysis of server memory and CPU utilization to make sure that there were no bottlenecks between virtualized applications and the DS5020 Express. The memory and CPU utilization within the virtual machine is shown in Figure 9.[4]

Figure 9. Low CPU, Network, and Memory Utilization

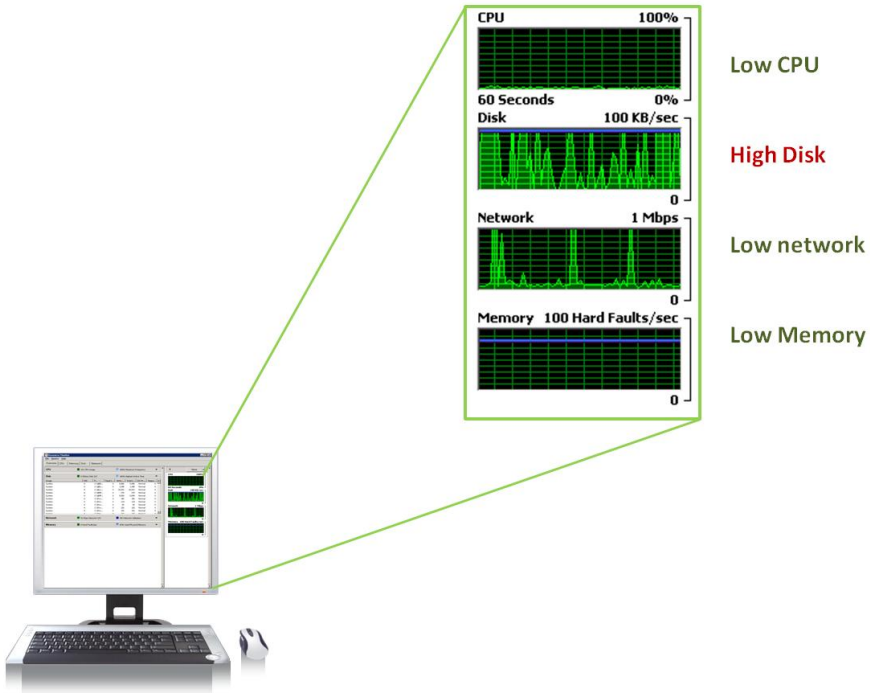

These screenshots were taken during the peak activity phase of the four cell test. With memory and CPU utilization at less than 10%, there was no obvious bottleneck between virtualized applications and the DS5020 Express.

Mixed Real-world IOPS Scalability

IOs per second, or IOPS, is a measure of the number of operations that a storage system can perform in parallel. When a system is able to move a lot of IOPS—from disk and from cache—it will tend to be able to service more applications and users in parallel. Much like the horsepower rating for a car engine, the IOPS rating for a storage controller can be used as an indicator of the power of a storage system engine.

[4] Low CPU and memory utilization on the physical server was nearly identical to the utilization reported within the virtual machine.

While IOPS out of a cache is typically a big number and can provide an indication of the speed of the front end of a storage controller, IOPS from disk is a more useful metric when determining the real-world performance of a storage system servicing a mix of business applications. For example, e-mail and interactive database applications tend to be random in nature and therefore benefit from good IOPS from disk. With that said, a mix of real-world applications tends to have random and sequential IO traffic patterns that may be serviced from disk or from cache.

ESG Lab recorded the number of IOPS as reported by the DS5020 Express as the number of virtual machines running mixed real-world application workloads was increased from four through sixteen. With a mix of random and sequential IO over 112 disk drives, the goal was not to record a big IOPS number; the goal with this exercise was an assessment of the scalability of the DS5020 Express as an increasing number of applications are consolidated onto a single virtualized platform. Mixed workload IOPS scalability is shown in Figure 10.

Figure 10. DS5020 Express Mixed Workload Scalability

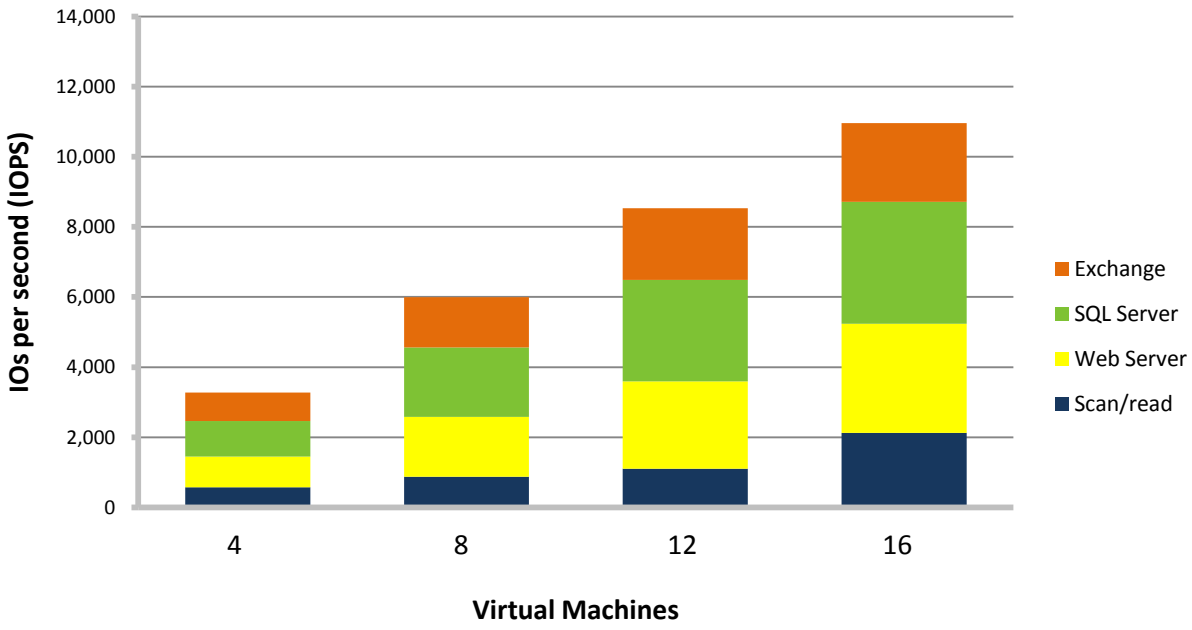

What the Numbers Mean

- IOPS varied throughout the mixed workload test with peaks occurring towards the beginning during the SQL Server small IOPs phase and then again towards the end as the Jetstress utility performed a database consistency check.
- A peak of 16,572 IOPS was reported by the DS5020 performance monitor during the four cell run.
- The average number of IOPS during each three hour run is shown in Figure 10.
- The average number of IOPS scaled in a near-linear fashion as mixed real-world application traffic increased from four through sixteen virtual servers.

Why This Matters

Predictable performance scalability is a critical concern when a mix of applications shares a storage system. A burst of IO activity in one application (e.g., a database consistency check) can lead to poor response times, lost productivity, and, in the worst case, lost revenue.

ESG Lab confirmed that the rate of IOs processed by the DS5020 Express scales extremely well as many applications ran in parallel when running a mix of real-world application workloads.

Mixed Application Performance Scalability

Having looked at the IOPS and throughput ratings of the turbo-charged DS5020 Express engine, here's where the rubber meets the road as we examine performance at the application level. The output from each of the industry standard benchmark utilities was analyzed to determine the performance scalability and responsiveness of real-world applications running in a consolidated virtual environment.

Microsoft Exchange

The IO efficiency and performance efficiency of Microsoft Exchange have improved significantly over the years. Architectural improvements in Exchange 2010 including a new store schema, larger page sizes (8 KB to 32 KB), improved read/write coalescing, improved pre-read support, and increased cache effectiveness have reduced the number of IOs per user up to 70% compared to Exchange 2007.[5] ESG Lab typically uses a value of 0.5 IOPS per mailbox to emulate a typical worker when testing with JetStress 2007. A value of 0.12 IOPS per mailbox was used during JetStress 2010 testing to reflect the 70% reduction in IOPS compared to Exchange 2007.

The Microsoft Jetstress 2010 utility was used to see how many simulated e-mail users could be supported by the DS5020 Express during mixed workload testing. The number of IOPS and response time for each database and log volume was recorded at the end of each Jetstress run. A response time goal of 20 milliseconds or less for database reads is required to pass the test. These values are defined by Microsoft as a limit beyond which end-users will feel that their e-mail system is acting slowly. [6]

Microsoft Jetstress logs were used to determine the number of IOPS and response times as the number of active virtual machines was increased from four through sixteen.[7] Based on a 0.12 IOPS user profile, the number of IOPS was used to calculate the number of supported Exchange users. The number of supported mailboxes as the number of virtual machines was increased from four to sixteen is shown in Figure 11, Figure 12, and Table 2.

Figure 11. Mixed E-mail Scalability (Number of Mailboxes)

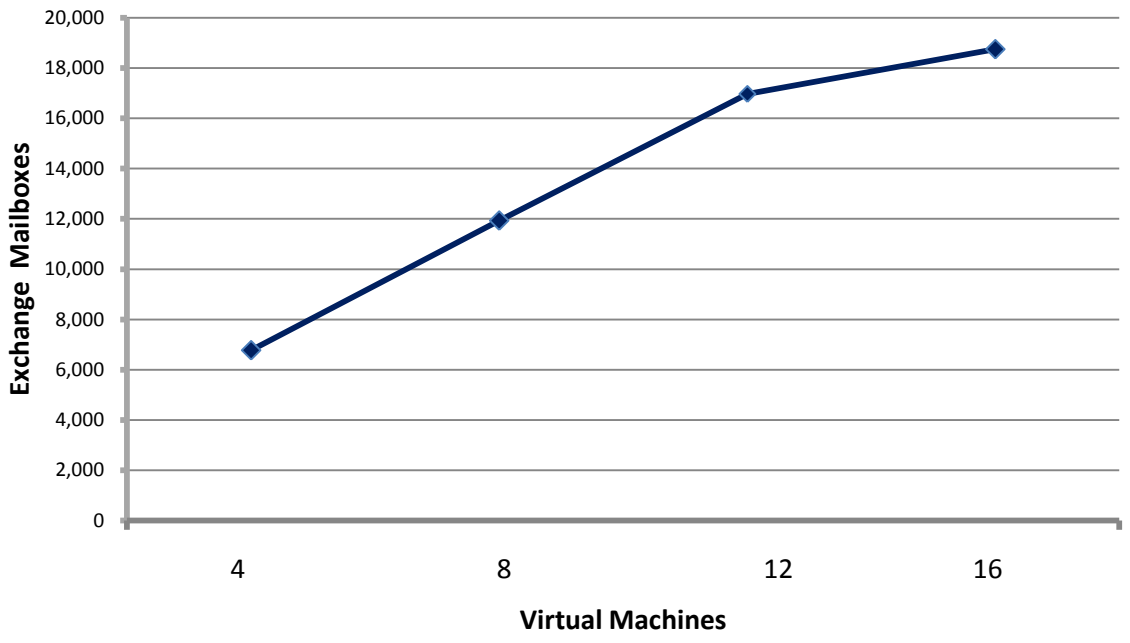

[5] http://download.microsoft.com/download/D/1/B/D1BE3AEC-A9CD-4459-99F1-B28867FAE20B/Exchange2010TCP_parte8.pdf
[6] http://technet.microsoft.com/en-us/library/bb738152(EXCHG.80).aspx
[7] A sample JetStress log is included in the Appendix.

Figure 12. Mixed E-mail Scalability (Response Time)

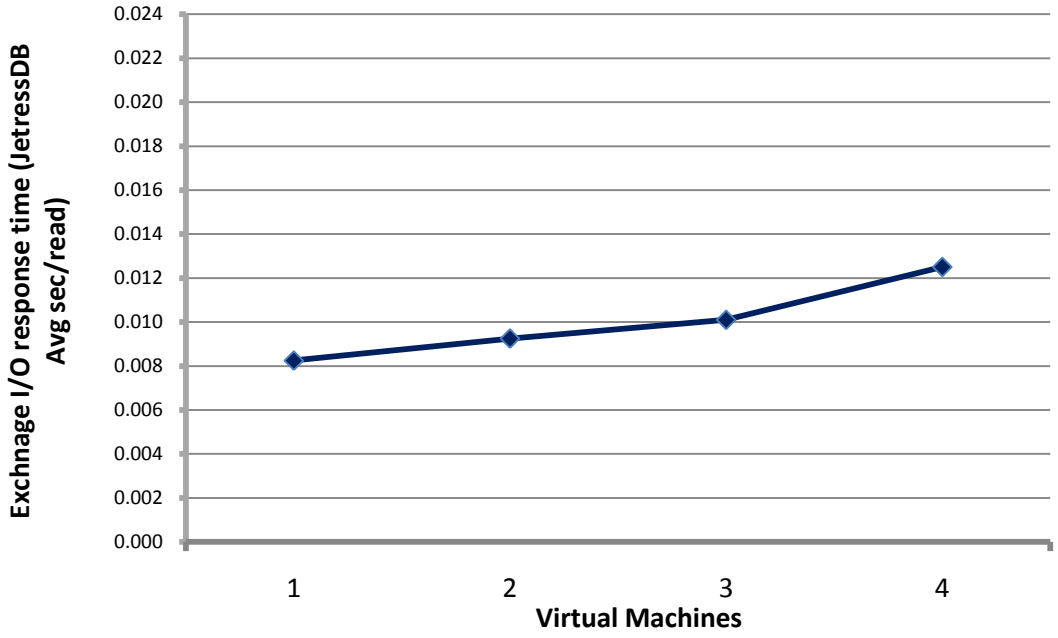

Table 2: Jetstress Performance Results

Virtual Machines	Achieved IO per Seconds	Users (0.12 Profile)	DB Avg. Disk Sec/Read
4	813	6,775	.008
8	1,432	11,933	.009
12	2,037	16,975	.010
16	2,250	18,750	.013

What the Numbers Mean

- The single cell mixed application test supported 6,775 Exchange users with an average DB disk response time of 8 milliseconds.
- Performance scaled in a near-linear fashion to 18,750 users while the DS5020 Express was busy processing and servicing other applications concurrently.
- As the number of simulated e-mail users was increased, the DS5020 Express provided excellent response times that are well within Microsoft's guidelines. For example, the Microsoft guideline for a database read volume is 20 milliseconds as shown by the dotted line in Figure 12.
- The IO efficiency improvements in Exchange 2010 reduce the cost of delivering e-mail support in mixed virtual server environments. In this case, ESG Lab supported up to 18,750 mailboxes on four virtualized Exchange 2010 servers in a mixed workload environment—twice the expected number of supported mailboxes within an Exchange 2007 environment.

Microsoft SQL Server

The Microsoft sqlio utility was used to measure the performance of small (8 KB) random and large (256 KB) sequential read workloads. The small random results were used to predict the performance and scalability of response time sensitive interactive database applications (e.g. OLTP). The large sequential results were used to predict the performance of throughput intensive database mining and decision support systems (DSS).

ESG used the following Microsoft TechNet guidelines to interpret the results:

Disk Latency - Lower values are better but this can vary and is dependent on the size and nature of the IOs being issued. On well-tuned IO subsystems, ideal values would be:

- 4–20 ms for Data on OLTP systems (ideally 10 ms or less).
- 30 ms or less on DSS (decision support system) type.[8]

The results for the four cell sqlio test are summarized in Table 3. A sample sqlio log is shown in the Appendix.

Table 3: Orion Four Cell Performance Results

Cell	OLTP (IOPS)	DSS (MB/sec)	OLTP Latency (ms)	DSS Latency (ms)
1	868	163	6.7	11
2	832	174	6.8	11
3	874	161	6.8	11
4	901	152	6.7	11
	3,475	650	5.75	11

What the Numbers Mean

- The four cell test achieved a grand total of 3,475 OLTP IOPS and 650 MB/sec of aggregate DSS throughput while the system was simultaneously running a mix of real-world application workloads.
- The DS5020 delivered excellent response times that are well under the 20 and 30 millisecond Microsoft TechNet limits.
- Database architects often use a value of 150 IOPS per drive for planning purposes. The total number of small IOPS processed during the busy four cell test yielded an excellent rate of 224 IOPS per drive.
- ESG Lab believes that these are excellent results—especially given the mix of IO intensive workloads being serviced by the DS5020 Express in parallel.

[8] Sqlio is an unsupported tool provided by Microsoft that is available for download at: http://www.microsoft.com/downloads/details.aspx?familyid=9a8b005b-84e4-4f24-8d65-cb53442d9e19&displaylang=en

Web Server and Scan/Read

Performance results for the web server and scan/read workloads executing within virtual machines during the four cell test are listed in Table 4.

Table 4: Iometer Four Cell Performance Results

Cell	Web Server (IOPs)	Scan/Read (MB/sec)
1	788	104
2	786	103
3	779	101
4	753	105
Total	***3,106***	***413***

What the Numbers Mean

- Given the cache friendly, read-only nature of web server IO traffic, ESG Lab believes that these results indicate that the DS5020 Express has the horsepower required to service tens of thousands of simultaneous page requests.
- ESG Lab believes that a file system workload would produce results that are approximately similar to the web server workload used for this test.
- Each of the four scan/read streams sustained more than 100 MB/sec of throughput for the entire duration of the mixed workload test. A stream of this magnitude could service the data needs of a number of simultaneous backup streams, a very aggressive scan and index job, or a throughput intensive database table scan—with no perceivable performance impact on applications that are running parallel.

Much like the electrical system in your home, figuring out how many appliances you can run in parallel before blowing a fuse is not a function of the number of wires behind the walls. What matters more is the design of the circuits used to distribute the right amount of power to appliances when needed. ESG Lab testing indicates that the DS5020 Express engine delivers the right amount of power to virtualized applications when needed.

Why This Matters

Excessive downtime and slow response time can result in the loss of sales, loss of customer goodwill, loss of productivity, loss of competitiveness, and increased costs. With more and more companies running entire suites of business applications on virtual servers, mixed workload scalability with predictable performance is needed.

E-mail is often considered the most significant business application today and, within the world of e-mail, Microsoft Exchange rules the roost. ESG Lab testing confirmed that the DS5020 Express can sufficiently handle a very large number of Exchange users—even as it services other applications and thousands of users with predictably fast response times.

DS3950 Express Performance Analysis

The DS3950 Express supports up to four FC host interfaces as compared to the DS5020 Express, which supports up to eight FC host interfaces. Otherwise, the components and architecture of the DS3950 Express are exactly the same as the DS5020 Express. ESG Lab tested the DS5020 Express with only four active FC host interfaces connected to the IBM System x3950 M2 servers with a goal of analyzing the performance difference between the DS5020 Express and the DS3950 Express. The results are shown in Figure 13 and Table 5.

Figure 13. DS3950 Express vs. DS5020 Express

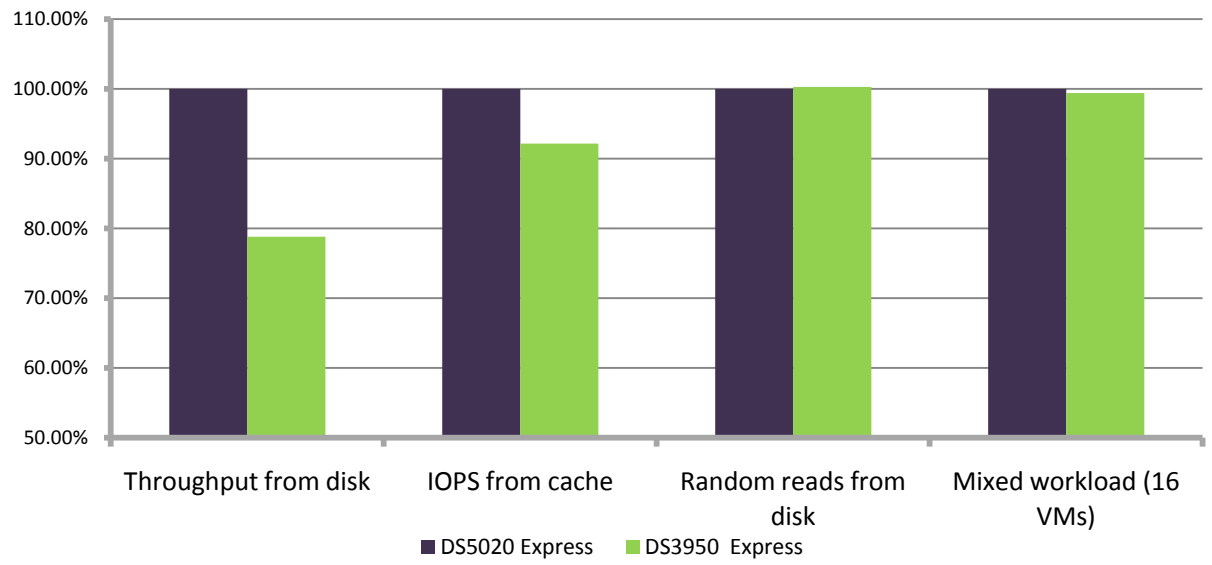

Table 5: DS3950 Express vs. DS5020 Express

Workload	DS5020 Express (8 paths)	DS3950 Express (DS5020 Express, 4 paths)
Throughput from disk	100%	78.82%
IOPS from cache	100%	92.16%
Random reads from disk	100%	100.26%
Mixed Workload (16 VM's)	100%	99.39%

What the Numbers Mean

- The simulated DS3950 Express with four active host paths delivered 78.82% of the throughput and 92.17% of the IOPS from cache compared to the DS5020 Express with eight active paths.
- The mostly random ESG Lab mixed workload performed roughly the same. This is due to the fact that the most important performance consideration for mostly random business application workloads is the number of disk drives operating in parallel. In this case, the same number of drives was tested (112).

Why This Matters

The IBM System Storage DS3950 Express is a cost effective alternative to the DS5020 Express for the mixed application workloads tested by ESG Lab. For environments with more bandwidth intensive requirements (e.g., backup to disk, video streaming, lots of virtual servers), the DS5020 Express with twice the host bandwidth—or the DS5300 with four times the host bandwidth—is a more appropriate solution.

ESG Lab Validation Highlights

- ☑ 3.3 GB/sec of aggregate throughput was sustained during characterization testing within a Microsoft Hyper-V R2 environment.
- ☑ A single DS5020 attached to a pair of System x3950 M2 servers running a mix of real-world application workloads in 16 Microsoft Hyper-V R2 enable virtual machines supports up to:
  - **18,750 mailboxes** using the Microsoft Exchange 2010 Jetstress utility
  - *and* **3,475 small database IOs per second** using the Microsoft SQLIO utility
  - *and* **650 MB/sec of database throughput** for large scan operations using the SQLIO utility
  - *and* **3,106 simulated web server IOPs**
  - *and* **413 MB/sec of throughput** for simulated backup/scan/index jobs
  - with the predictably fast response times and scalability
- ☑ Excellent IOPs per drive were recorded (e.g., 224 for the SQL Server OLTP test).
- ☑ As the number of virtual machines sharing a single DS5020 was increased, performance scaled in a near linear fashion with predictably fast response times (8 to 13 milliseconds for Jetstress DB reads).
- ☑ The DS5020 had horsepower to spare for rebuilds and advanced functions including copy services and remote replication.

Issues to Consider

- ☑ Generally accepted best practices and predominantly default Microsoft and IBM System Storage settings were used during the design of this test. As expected after any benchmark of this magnitude, deep analysis of the results indicates that tuning would probably yield slighter higher absolute results. Given that the goal of this test was not to generate a big number, ESG Lab is confident that the results presented in this report meet the objective of estimating performance scalability and responsiveness as a growing number of virtual machines share a consolidated pool of DS5020 Express storage.
- ☑ For applications requiring extreme performance beyond that which is provided by FC and SATA drives, ESG Lab believes that the DS5020 Express is an ideal architecture for the selective use of solid state disk (SSD) devices. While mixed workload testing was not performed with SSD devices, ESG Lab is confident that SSD devices could be used to improve the performance of highly referenced database indexes and temp files.
- ☑ The test results/data presented in this document are based on industry-standard benchmarks deployed together in a controlled environment. Due to the many variables in each production data center environment, it is still important to perform capacity planning and testing in your own environment to validate a storage system configuration.

The Bigger Truth

Server virtualization is being deployed by a growing number of organizations to lower costs, improve resource utilization, provide non-disruptive upgrades, and increase availability. Each benefit is fundamentally enabled by de-coupling servers, applications, and data from specific physical assets. Storage virtualization takes those very same benefits and extends them from servers to the underlying storage domain—bringing IT organizations one step closer to the ideal of a completely virtualized IT infrastructure.

While the benefits of a completely virtualized infrastructure are obvious to most IT managers, performance is a real concern. Server, storage, and network administrators are looking for answers to a number of questions:

- Can we meet performance service level agreements for a mix of business-critical applications?
- Does the storage system have the horsepower to serve mixed, real-world applications?
- Can the storage system scale to accommodate future growth and consolidation?

IBM approached ESG Lab with an ambitious goal of answering these questions. A performance benchmark was designed to measure the performance capabilities of a storage system subjected to an IO intensive mix of virtualized business applications running in a Microsoft Windows Server 2008 R2 environment. A cell concept was used during the design of this test. Each cell was composed of four applications, each running in its own virtual machine. The server horsepower of a pair of IBM System x3950 M2 servers was used to drive up to four cells and sixteen virtual applications in parallel.

IBM has more than a decade of experience delivering modular FC-attached storage systems designed to meet the cost-optimized performance demands of medium-sized organizations, mid-tier applications, remote departments, and near-line applications. The IBM DS5000 series builds on the heritage of the previous generation DS4000 series disk system with more than 87,000 systems and 511 petabytes shipped to date. The engine under the hood of the DS5000 Series has been turbo-charged to meet the real-world performance demands of virtualized applications. With twice the host bandwidth and three times the internal bandwidth of the previous generation DS4700, the DS5020 Express is designed to deliver the high performance, low latency, and balanced scalability needed to meet the demanding performance needs of a mix of real-world applications sharing a consolidated infrastructure.

ESG Lab testing began with the configuration and deployment of Microsoft enabled virtual machines. Bare metal Hyper-V R2, which is built into Windows Server 2008 R2, provided the underlying virtualization technology. Wizard driven configuration with a familiar look and feel made it easy to configure, clone, and manage the virtualized environment. The combination of high performance IBM DS5020 storage, the enhanced high availability of Windows Server 2008 R2 (e.g. Live Migration), and the improved IO performance of Hyper-V R2 and Exchange 2010 created a solid foundation for mixed workload testing.

The first set of tests revealed that the DS5020 Express can deliver up to 3.3 GB/sec of raw aggregate throughput in a Microsoft Hyper-V R2 environment. This result was an early indicator that the IBM DS5020 Express has the internal bandwidth and processing power needed to serve a mix of real-world application workloads. The results of the mixed workload tests were even more impressive. A single DS5020 Express simultaneously supported 18,750 simulated Exchange 2010 users *and* 3,475 SQL Server OLTP IOPs a*nd* 675 MB/sec of SQL Server DSS throughput *and* 3,206 simulated web server IOPs *and* 413 MB/sec of throughput for bandwidth intensive read traffic—all while delivering predictably fast response times.

ESG Lab is pleased to report that the combination of IBM System Storage DS5020 Express, IBM System x3950 M2 servers, and QLogic 8 Gb Fibre Channel host bus adapters can be used with confidence to meet the performance needs of a mix of real-world business applications running in a Microsoft Windows Server 2008 R2 environment with built-in Hyper-V R2 virtual server support.

Appendix

Table 6. Test Bed Overview

Storage	
IBM DS5020 Express, Firmware: 07.60.08.00	8 drive trays, 112 15K RPM FC drives
Servers	
Two IBM System x3950 M2	CPU type: Intel Xeon 7740 (6 cores per socket @2.67 GHz) Total CPU cores: 48 Total RAM: 512 GB DDR-2
Host Bus Adapters	
QLogic 8 Gb FC Dual-port HBA for IBM System x	
Virtualization Software and Guest Operating Systems	
Host OS	Microsoft Server 2008 R2, Hyper-V role enabled
Guest OS	Windows Server 2008 R2

Table 7. Bill of Materials

IBM Code	Description	Quantity
7233AC2	IBM System x3950 M2	2
42D0510	Qlogic 8G Fiber Channel HBA	4
1814-20A	IBM DS5020 Storage System	1
1814-52A	IBM EXP520 Expansion Drive Enclosure	6
4201	300GB 15K FC HDD	112
7393	33-64 Disk Drive Attach License	1
7394	65-112 Disk Drive Attach License	1
2031	Two dual 8 Gbps FC host ports 2031	1
8701	4 Host Partitions	1

Table 8. Benchmark Utilities/Workload Generators

Characterization	Iometer, version 2008.06.18 Dynomo clients ran within eight guest VMs running on a pair of IBM System x3950 M2 servers with Windows Server 2008 R2 being used as both the host and guest operating system. Each of eight LUNS was built using four 15K RPM FC drives configured as a RAID-10 group (32 drives in total). Each of the LUNs was tested as physical drive over VHD volumes in a Microsoft Hyper-V R2 environment. The servers worked in parallel accessing the DS5020 Express through eight FC interfaces negotiated at 8 Gbps. Maximum throughput was measured using 1 MB sequential reads. Four workers, sixteen IOs per physical drive.
E-Mail	Microsoft Jetstress 2010 Beta, version 14.00.0639.012 • Mailboxes – 4,500 • Mailbox size – 200 MB • Target IOPS per mailbox – 0.12 • Thread Count – 6 • Min DB cache – 32 MB • Max DB cache – 256 MB • Insert operations – 40% • Delete operations – 20% • Replace operations – 5% • Read operations – 35% • Lazy commits – 70% • Run Background Database Maintenance – True • Number of Copies per Database – 1
Database Workload Generator	SQLIO, version 1.5 • IO workloads: OLTP read/ DSS scan • IO sizes: 8 KB/ 256 KB • IO Types: Random read /Sequential read • Threads: 1 • File size: 100,000 MB • Duration: 2,700 seconds • Buffering: No
Web Server	Iometer, version 2008.06.18 Four workers, four outstanding IOs per physical drive 100% random reads, assorted block sizes
Scanner/Reader	Iometer, version 2008.06.18 One worker, one outstanding IO per physical drive 100% 256 KB sequential reads

Figure 14. E-mail Results

This is an example of the output created by the Jetstress utility. It shows the performance for the Jetstress running in the first VM (JS01) during the single cell/four VM test.

Microsoft Exchange Server Jetstress

Performance Test Result Report

Test Summary

Overall Test Result	**Pass**
Machine Name	JS01
Test Description	
Test Start Time	1/28/2010 5:22:18 AM
Test End Time	1/28/2010 7:22:52 AM
Jetstress Version	14.00.0639.012
Ese Version	14.00.0639.019
Operating System	Windows Server 2008 R2 Enterprise (6.1.7600.0)
Performance Log	C:\temp\tools\Performance_2010_1_28_5_22_21.blg C:\temp\tools\DBChecksum_2010_1_28_7_22_52.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	813.711
Target Transactional I/O per Second	480
Initial Database Size (bytes)	992877805568
Final Database Size (bytes)	994471641088
Database Files (Count)	1

Jetstress System Parameters

Thread Count	7 (per database)
Minimum Database Cache	32.0 MB
Maximum Database Cache	256.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%
Run Background Database Maintenance	True
Number of Copies per Database	1

Database Configuration

Instance2424.1 Log Path: E:\
Database: F:\Jetstress001001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance2424.1	8.254	1.900	548.585	265.127	32826.777	34212.313	0.000	0.432	0.000	148.120	0.000	4933.599

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance2424.1	27.325	261789.344

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance2424.1	0.000	0.000

Total I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance2424.1	8.254	1.900	575.910	265.127	43690.383	34212.313	0.000	0.432	0.000	148.120	0.000	4933.599

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	3.621	2.445	5.203
Available MBytes	10763.621	10756.000	10765.000
Free System Page Table Entries	33559477.263	33557902.000	33562510.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	28662502.400	28643328.000	28696576.000
Pool Paged Bytes	74141747.200	74125312.000	74227712.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log1/27/2010 2:17:51 PM -- Jetstress testing begins ...
1/28/2010 5:22:20 AM -- Attaching databases ...
1/28/2010 5:22:20 AM -- Prepare testing ends.
1/28/2010 5:22:20 AM -- Dispatching transactions begins ...
1/28/2010 5:22:20 AM -- Database cache settings: (minimum: 32.0 MB, maximum: 256.0 MB)
1/28/2010 5:22:20 AM -- Database flush thresholds: (start: 2.5 MB, stop: 5.1 MB)
1/28/2010 5:22:21 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 50 msec/read).
1/28/2010 5:22:21 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 50 msec/write).
1/28/2010 5:22:24 AM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
1/28/2010 5:22:24 AM -- Performance logging begins (interval: 15000 ms).
1/28/2010 5:22:24 AM -- Attaining prerequisites:
1/28/2010 5:22:50 AM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 244183000.0 (lower bound: 241591900.0, upper bound: none)
1/28/2010 7:22:51 AM -- Performance logging ends.
1/28/2010 7:22:51 AM -- JetInterop batch transaction stats: 93010.
1/28/2010 7:22:51 AM -- Dispatching transactions ends.
1/28/2010 7:22:51 AM -- Shutting down databases ...
1/28/2010 7:22:52 AM -- Instance2424.1 (complete)
1/28/2010 7:22:53 AM -- Performance logging begins (interval: 30000 ms).
1/28/2010 7:22:53 AM -- Verifying database checksums ...
1/28/2010 8:03:31 AM -- F: (100% processed)
1/28/2010 8:03:31 AM -- Performance logging ends.
1/28/2010 8:03:31 AM -- C:\temp\tools\DBChecksum_2010_1_28_7_22_52.blg has 81 samples.
1/28/2010 8:03:32 AM -- C:\temp\tools\DBChecksum_2010_1_28_7_22_52.html is saved.
1/28/2010 8:03:32 AM -- Verifying log checksums ...
1/28/2010 8:03:32 AM -- E:\ (11 log(s) processed)
1/28/2010 8:03:32 AM -- C:\temp\tools\Performance_2010_1_28_5_22_21.blg has 481 samples.
1/28/2010 8:03:32 AM -- Creating test report ...
1/28/2010 8:03:33 AM -- Instance2424.1 has 8.3 for I/O Database Reads Average Latency.
1/28/2010 8:03:33 AM -- Instance2424.1 has 0.4 for I/O Log Writes Average Latency.
1/28/2010 8:03:33 AM -- Instance2424.1 has 0.4 for I/O Log Reads Average Latency.
1/28/2010 8:03:33 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
1/28/2010 8:03:33 AM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
1/28/2010 8:03:33 AM -- C:\temp\tools\Performance_2010_1_28_5_22_21.xml has 479 samples queried.

Figure 15. SQL IO Parameters and Results

The following parameters that were used during sqlio testing were derived from Microsoft SQL Server best practices (http://msdn.microsoft.com/en-us/library/cc966412.aspx):

```
echo ***run 1 of 2: OLTP reads***
sqlio -BN -o8 -kR -frandom -b8 -LS -s2700 -Fsqlioparms.txt >>sqlresults.txt
echo ***run 2 of 2: DSS scans ***
sqlio -BN -o8 -kR -fsequential -b256 -LS -s2700 -Fsqlioparms.txt >>sqlresults.txt
sqlioparms.txt:  e:\sqlio_test.dat 1 0x0 100000
```

A sample of the results of sqlio testing follows. This example shows the performance of one of four sqlio tests which ran in parallel during the four cell/16 VM test.

```
 ********** esg lab mixed workload sqlio testing begins**********
Wed 01/27/2010 10:36 AM
***run 1 of 2: OLTP read***
sqlio v1.5.SG
using system counter for latency timings, 3579545 counts per second
parameter file used: sqlioparms.txt
        file e:\sqlio_test.dat with 1 thread (0) using mask 0x0 (0)
1 thread reading for 2700 secs from file e:\sqlio_test.dat
        using 8KB random IOs
        enabling multiple I/Os per thread with 8 outstanding
        buffering set to not use file nor disk caches (as is SQL Server)
using specified size: 100000 MB for file: e:\sqlio_test.dat
initialization done
CUMULATIVE DATA:
throughput metrics:
IOs/sec:   867.36
MBs/sec:     6.77
latency metrics:
Min_Latency(ms): 0
Avg_Latency(ms): 8
Max_Latency(ms): 199
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+
%:  0  0  2  5 10 12 13 11  9  7  6  4  4  3  2  2  1  1  1  1  1  1  0  3
***run 2 of 2: DSS scan***
sqlio v1.5.SG
using system counter for latency timings, 3579545 counts per second
parameter file used: sqlioparms.txt
        file e:\sqlio_test.dat with 1 thread (0) using mask 0x0 (0)
1 thread reading for 2700 secs from file e:\sqlio_test.dat
        using 256KB sequential IOs
        enabling multiple I/Os per thread with 8 outstanding
        buffering set to not use file nor disk caches (as is SQL Server)
using specified size: 100000 MB for file: e:\sqlio_test.dat
initialization done
CUMULATIVE DATA:
throughput metrics:
IOs/sec:   710.38
MBs/sec:   177.59
latency metrics:
Min_Latency(ms): 1
Avg_Latency(ms): 10
Max_Latency(ms): 220
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+
%:  0  0  2  3  4  5  6  6  6  7  7  8  9  8  7  6  4  3  2  2  1  1  0  0  1
********** esg lab mixed workload sqlio testing ends**********
```

Figure 16. Web Server and Scan/Read Results

This is an example of the output created by the Iometer utility during mixed workload testing. This example shows the performance of the web server and scan/read virtual machines during the single cell/4 VM test.

'Test Type	Test Description												
0	esg mixed workload												
'Version													
2008.06.18													
'Time Stamp													
2010-01-28 05:20:50:751													
'Access specifications													
'Access speci	default assignment												
Web Server	1												
'size	% of size	% reads	% random	delay	burst	align	reply						
512	22	100	100	0	1	0	0						
1024	15	100	100	0	1	0	0						
2048	8	100	100	0	1	0	0						
4096	23	100	100	0	1	0	0						
8192	15	100	100	0	1	0	0						
16384	2	100	100	0	1	0	0						
32768	6	100	100	0	1	0	0						
65536	7	100	100	0	1	0	0						
131072	1	100	100	0	1	0	0						
524288	1	100	100	0	1	0	0						
'Access speci	default assignment												
Backup reade	1												
'size	% of size	% reads	% random	delay	burst	align	reply						
262144	100	100	0	0	1	0	0						
'End access specifications													
'Results													
'Target Type	Target Name	Access Specif	# Manager	# Worke	# Disk	IOps	Read IOps	Write IOps	MBps	Read MBps	Write MBps	Transactions	Connections
ALL	All		2	5	5	1452.55	1452.55	0	157.69726	157.69726	0	1452.5504	0
MANAGER	WB01	Web Server		4	4	875.256	875.256	0	13.373714	13.373714	0	875.25621	0
PROCESSOR	CPU 0												
PROCESSOR	CPU 1												
WORKER	Worker 1	Web Server			1	220.213	220.213	0	3.364459	3.364459	0	220.21349	0
DISK	1: "IBM 1814	FAStT 106"				220.213	220.213	0	3.364459	3.364459	0	220.21349	0
WORKER	Worker 2	Web Server			1	217.342	217.342	0	3.320721	3.320721	0	217.34244	0
DISK	1: "IBM 1814	FAStT 106"				217.342	217.342	0	3.320721	3.320721	0	217.34244	0
WORKER	Worker 3	Web Server			1	220.292	220.292	0	3.36571	3.36571	0	220.29226	0
DISK	1: "IBM 1814	FAStT 106"				220.292	220.292	0	3.36571	3.36571	0	220.29226	0
WORKER	Worker 4	Web Server			1	217.408	217.408	0	3.322824	3.322824	0	217.40802	0
DISK	1: "IBM 1814	FAStT 106"				217.408	217.408	0	3.322824	3.322824	0	217.40802	0
MANAGER	BK01	Backup reader		1	1	577.294	577.294	0	144.32355	144.32355	0	577.29418	0
PROCESSOR	CPU 0												
PROCESSOR	CPU 1												
WORKER	Worker 1	Backup reader			1	577.294	577.294	0	144.32355	144.32355	0	577.29418	0
DISK	1: "IBM 1814	FAStT 106"				577.294	577.294	0	144.32355	144.32355	0	577.29418	0
'Time Stamp													
2010-01-28 08:20:56:329													

Figure 17. DS5020 Express Configuration Details

The following excerpts were extracted from the IBM DS5020 Storage System Profile Summary.

```
PROFILE FOR STORAGE SUBSYSTEM: ESG_DS5020 (Fri Oct 02 05:53:29 PDT 2009)

SUMMARY------------------------------

  Number of controllers:           2

  High performance tier controllers:  Enabled

  Number of arrays:               24

  RAID 6:                      Enabled

  Total number of logical drives used:    25
    Number of standard logical drives:    24
    Number of access logical drives:      1
  Total number of logical drives allowed:  1024

  Drive Limit Management:
    Number of drive slots discovered:  112
    Number of drive slots allowed:     112

  FlashCopy Logical Drives:                    Enabled
    Number of flashcopies used:                  0
    Number of flashcopies allowed:               2
    Number of flashcopies allowed per base logical drive:  2

  Remote Logical Drive Mirroring:  Disabled/Deactivated
    Number of mirrors used:      0
    Number of mirrors allowed:   0

  VolumeCopy:               Disabled
    Number of copies used:    0
    Number of copies allowed:  0

  Number of drives:         112
  Mixed drive types:        Enabled
  Current media type(s):     Hard Disk Drive (112)
  Current interface type(s):  Fibre (112)

  Total hot spare drives:    0
    Standby:             0
    In use:              0

  Drive Security:          Disabled
  Security key identifier:  None

  Storage Partitioning:           Enabled
    Number of partitions used:    2
    Number of partitions allowed:  128

  Number of logical drives allowed per partition:  256

  Access logical drive:  LUN 31,31,31 (see Mappings section for details)
  Default host OS:       DEFAULT (Host OS index 0)

  Current configuration
    Firmware version:              07.60.08.00
    NVSRAM version:                 N1814D20R1060V08
    EMW version:                   10.60.G5.05
```

```
   AMW version:                    10.60.G5.05

  NVSRAM configured for batteries:       Yes

  Start cache flushing at (in percentage): 80
  Stop cache flushing at (in percentage):  80
  Cache block size (in KB):              16


  Media scan frequency (in days):          Disabled
  Failover alert delay (in minutes):       5

  Feature enable identifier:               303030393420303030313235204A93DE36

  Feature pack:                            DS5020 Model 20, 24, 28
  Feature pack submodel ID:                121

  Storage Subsystem world-wide identifier (ID):  60080E500017B6BA000000004A93DE34

CONTROLLERS------------------------------
  Number of controllers: 2

    Controller in Enclosure 85, Slot A

      Status:                Online

      Current configuration
        Firmware version:      07.60.08.00
          Appware version:     07.60.08.00
          Bootware version:    07.60.08.00
        NVSRAM version:        N1814D20R1060V08

      Replacement part number:    37781-03
      Model name:                 4988
      Board ID:                   4988
      Submodel ID:                121
      Product ID:                 1814     FAStT
      Revision:                   1060
      Replacement part number:    37781-03
      Part number:                37781-03
      Serial number:              SQ91100094
      Vendor:                     IBM
      Date of manufacture:        June 2, 2009
      Trunking supported:         No
      Data Cache
        Total present:            1709 MB
        Total used:               1709 MB
      Processor cache:
        Total present:            339 MB
      Cache Backup Device
        Status:                   Optimal
        Type:                     USB flash drive
        Location:                 Controller A, Connector USB 1
        Capacity:                 1,960 MB
        Product ID:               eUSB
        Part number:              Not Available
        Serial number:            200902190239A7D8
        Revision level:           8715
        Manufacturer:             SMART
        Date of manufacture:      Not available
      Host Interface Board
        Status:                   Optimal
        Location:                 Slot 1
```

Type: Fibre channel
Number of ports: 2
Board ID: 0902
Replacement part number: L2-25043-03
Part number: PN L2-25043-03
Serial number: SN SQ91100387
Vendor: VN LSI
Date of manufacture: June 1, 2009
Date/Time: Fri Oct 02 05:54:13 PDT 2009

Associated Logical Drives (* = Preferred Owner):
BK_01*, BK_04*, JS_01*, JS_03*, JS_2L*, JS_4L*, OR_02*, OR_04*, OS_01*, OS_03*, WB_02*, WB_03*

STANDARD LOGICAL DRIVES----------

SUMMARY

Number of standard logical drives: 24

See other Logical Drives sub-tabs for premium feature information.

NAME	STATUS	CAPACITY	RAID LEVEL	ARRAY	MEDIA TYPE	INTERFACE TYPE
BK_01	Optimal	557.791 GB	10	BK_01	Hard Disk Drive	Fibre channel
BK_02	Optimal	557.793 GB	10	BK_02	Hard Disk Drive	Fibre channel
BK_03	Optimal	557.793 GB	10	BK_03	Hard Disk Drive	Fibre channel
BK_04	Optimal	557.793 GB	10	BK_04	Hard Disk Drive	Fibre channel
JS_01	Optimal	1.089 TB	10	JS_01	Hard Disk Drive	Fibre channel
JS_02	Optimal	1.089 TB	10	JS_02	Hard Disk Drive	Fibre channel
JS_03	Optimal	1.089 TB	10	JS_03	Hard Disk Drive	Fibre channel
JS_04	Optimal	1.089 TB	10	JS_04	Hard Disk Drive	Fibre channel
JS_1L	Optimal	557.793 GB	10	JS_1L	Hard Disk Drive	Fibre channel
JS_2L	Optimal	557.793 GB	10	JS_2L	Hard Disk Drive	Fibre channel
JS_3L	Optimal	557.793 GB	10	JS_3L	Hard Disk Drive	Fibre channel
JS_4L	Optimal	557.793 GB	10	JS_4L	Hard Disk Drive	Fibre channel
OR_01	Optimal	557.793 GB	10	OR_01	Hard Disk Drive	Fibre channel
OR_02	Optimal	557.793 GB	10	OR_02	Hard Disk Drive	Fibre channel
OR_03	Optimal	557.793 GB	10	OR_03	Hard Disk Drive	Fibre channel
OR_04	Optimal	557.793 GB	10	OR_04	Hard Disk Drive	Fibre channel
OS_01	Optimal	836.689 GB	5	0S_1	Hard Disk Drive	Fibre channel
OS_02	Optimal	836.689 GB	5	OS_02	Hard Disk Drive	Fibre channel
OS_03	Optimal	836.689 GB	5	OS_03	Hard Disk Drive	Fibre channel
OS_04	Optimal	836.689 GB	5	OS_04	Hard Disk Drive	Fibre channel
WB_01	Optimal	557.793 GB	10	WB_01	Hard Disk Drive	Fibre channel
WB_02	Optimal	557.793 GB	10	WB_02	Hard Disk Drive	Fibre channel
WB_03	Optimal	557.793 GB	10	WB_03	Hard Disk Drive	Fibre channel
WB_04	Optimal	557.793 GB	10	WB_04	Hard Disk Drive	Fibre channel

DETAILS

Logical Drive name: BK_01

Logical Drive status: Optimal

Capacity: 557.791 GB
Logical Drive ID: 60:08:0e:50:00:17:b6:ba:00:00:1a:5e:4a:c0:a8:24
Subsystem ID (SSID): 12
Associated array: BK_01
RAID level: 10

Secure: No

Media type: Hard Disk Drive

Interface type: Fibre channel
Enclosure loss protection: No

Preferred owner: Controller in slot A
Current owner: Controller in slot A

Segment size: 512 KB
Capacity reserved for future segment size changes: Yes
Maximum future segment size: 2,048 KB
Modification priority: High

Read cache: Enabled
Write cache: Enabled
Write cache without batteries: Disabled
Write cache with mirroring: Enabled
Flush write cache after (in seconds): 10.00
Dynamic cache read prefetch: Enabled

Enable background media scan: Enabled
Media scan with redundancy check: Disabled

Pre-Read redundancy check: Disabled

MAPPINGS (Storage Partitioning - Enabled (2 of 128 used))----------

Logical Drive Name	LUN	Controller	Accessible by	Logical Drive status
BK_01	4	A	Host Group Msft_01	Optimal
BK_03	10	B	Host Group Msft_01	Optimal
JS_01	2	A	Host Group Msft_01	Optimal
JS_03	8	A	Host Group Msft_01	Optimal
JS_2L	3	A	Host Group Msft_01	Optimal
JS_4L	9	A	Host Group Msft_01	Optimal
OR_02	1	A	Host Group Msft_01	Optimal
OR_04	7	A	Host Group Msft_01	Optimal
OS_01	0	A	Host Group Msft_01	Optimal
OS_03	6	A	Host Group Msft_01	Optimal
WB_02	5	A	Host Group Msft_01	Optimal
WB_04	11	B	Host Group Msft_01	Optimal
BK_02	4	B	Host Group Msft_02	Optimal
BK_04	10	A	Host Group Msft_02	Optimal
JS_02	2	B	Host Group Msft_02	Optimal
JS_04	8	B	Host Group Msft_02	Optimal
JS_1L	3	B	Host Group Msft_02	Optimal
JS_3L	9	B	Host Group Msft_02	Optimal
OR_01	1	B	Host Group Msft_02	Optimal
OR_03	7	B	Host Group Msft_02	Optimal
OS_02	0	B	Host Group Msft_02	Optimal
OS_04	6	B	Host Group Msft_02	Optimal
WB_01	5	B	Host Group Msft_02	Optimal
WB_03	11	A	Host Group Msft_02	Optimal
Access Logical Drive	31	A,B	Host deimos	Optimal
Access Logical Drive	31	A,B	Host phobos	Optimal
Access Logical Drive	31	A,B	Storage Subsystem	Optimal

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel:508.482.0188 Fax: 508.482.0218 | www.enterprisestrategygroup.com