

Lab Validation Report

IBM System Storage DS5020/DS3950 Express and IBM BladeCenter HS22

Real-World Mixed Workload Performance in VMware Environments

By Brian Garrett

November 2009

Contents

Introduction.....	3
The Challenges.....	3
The Solution.....	4
The Results	5
ESG Lab Validation	6
VMmark.....	6
A Mixed Real-world Storage Benchmark Methodology.....	7
Mixed Workloads.....	8
Test Bed.....	9
Drive Layout.....	10
Configuring Virtual Machines	11
The Results	12
<i>Characterization</i>	<i>12</i>
<i>Virtual Machine Utilization</i>	<i>14</i>
<i>Mixed Real-world IOPS Scalability</i>	<i>14</i>
<i>Handling Throughput Spikes with Ease.....</i>	<i>16</i>
<i>Mixed Application Performance Scalability.....</i>	<i>17</i>
<i>DS3950 Express Performance Analysis</i>	<i>21</i>
ESG Lab Validation Highlights.....	22
Issues to Consider	22
The Bigger Truth	23
Appendix	24

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 VMware. This ESG Lab report was sponsored by IBM.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. Copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482.0188.

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 BladeCenter HS22 servers deployed in a highly virtualized, consolidated 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.

Figure 1. Server Virtualization and Networked Storage Challenges



Source: Enterprise Strategy Group, 2009.

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

¹ Source: ESG Research Report, *The Impact of Server Virtualization on Storage*

deployed. That users would be so concerned with the performance of their infrastructures makes sense given the fact that 46% of virtualization users report that they currently run “Tier 1” applications on virtual machines and 33% plan to in the future.


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 BladeCenter HS22 servers:** The IBM BladeCenter HS22 is a highly efficient server with extraordinary scalability that delivers the ability to add more processing power, memory, or IO needed in virtualized environments.
- VMware® vSphere™:** Building on the power of VMware Infrastructure, VMware vSphere transforms IT infrastructures into a private cloud which enables the automated delivery of IT infrastructure as a service.
- QLogic 8 Gb SAN Switch Module and CFFh and CIOvy form factor FC Expansion cards for IBM BladeCenter:** Providing up to 20 ports of SAN connectivity, the QLogic 8 Gb Fibre Channel switch and expansion card infrastructure 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 BladeCenter H supports up to 14 blade servers, each populated with up to eight Intel Xeon 5500 processor cores and 98 GB of DDR-3 RAM. 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.

Figure 2. IBM Server and Storage Highlights



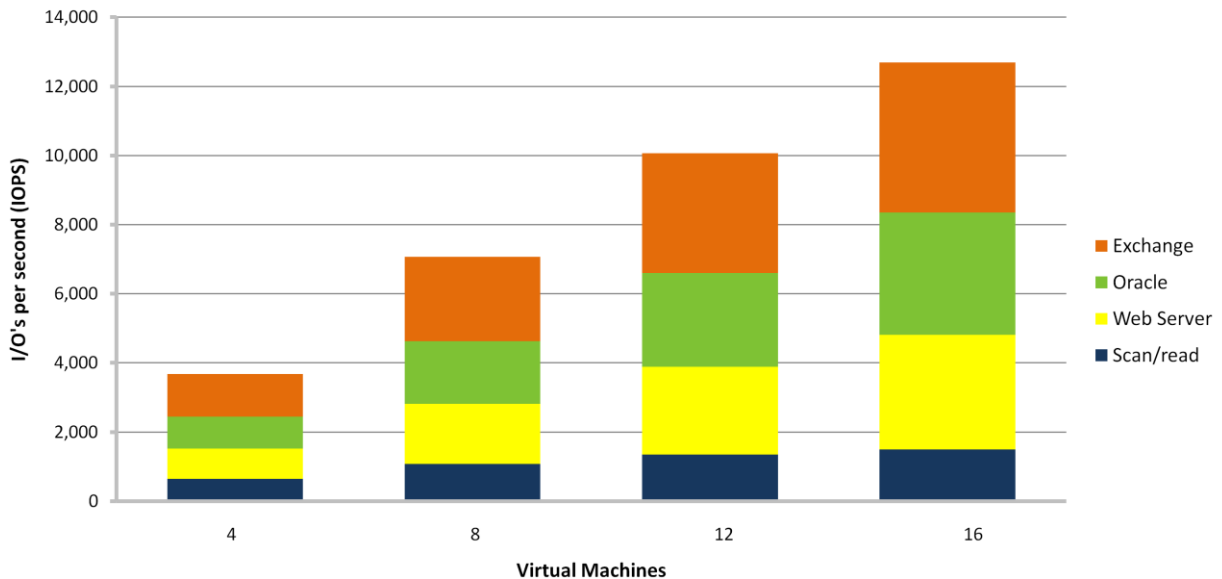
The Results

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 VMware vSphere-enabled virtual server environment powered by a pair of IBM BladeCenter HS22 servers. In particular, this report explores how:

- A single BladeCenter HS22 achieved an excellent VMmark mixed workload score of **24.05@17 Tiles**.
- A single DS5020 attached to a pair of BladeCenter HS22 servers running a mix of real-world application workloads in 16 virtual machines supports up to:
 - **8,680 mailboxes** using the Microsoft Exchange Jetstress utility
 - **and 3,593 small database IOs per second** using the Oracle Orion utility
 - **and 433 MB/sec of throughput** for large OLAP Oracle Orion operations
 - **and 3,317 simulated web server IOPs**
 - **and 374 MB/sec of throughput** for simulated backup/scan/index jobs
 - with the predictably fast response times and scalability
- Within a vSphere enabled infrastructure, the DS5020 Express achieved a maximum aggregate throughput of 3.1 GB/sec during bandwidth intensive throughput testing and 1.21 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 real-world performance capabilities of the IBM DS5020 Express were assessed by ESG Lab at an IBM facility located in Tuscon, Arizona. The methodology presented in this report was designed to assess the performance capabilities of a single IBM DS5020 Express storage system shared by multiple virtual servers running a mix of real-world application workloads. The cooperation of VMware, IBM and QLogic was key to the success of this project. In particular, this project benefitted from VMware's expertise in helping customers plan for the deployment of business-critical applications in virtual server environments and IBM's long heritage of success in the modular storage systems market.


VMmark

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.

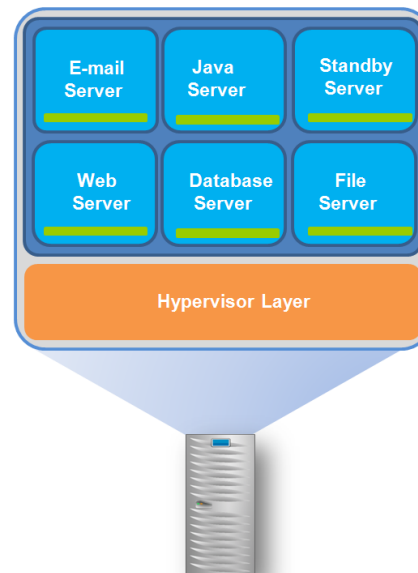
Traditional benchmarks running a single application workload can't help IT managers understand what happens when a mix of applications are deployed together in a virtual server environment. To overcome these limitations, VMware created a mixed workload benchmark called VMmark. VMmark uses a tile-based scheme for measuring application performance and provides a consistent methodology that captures both the overall scalability and individual application performance of a virtual server solution. As shown in Figure 4, compared to a traditional benchmark, which tests a single application running on a single physical server, VMmark measures performance as a mix of application workloads are run in parallel within virtual machines deployed on the same physical server.

Figure 4. Traditional Benchmarking vs. VMmark Tile-Based Benchmarking

Traditional Benchmarking



VMmark Tile-Based Benchmarking



The novel VMmark tile concept is simple, yet elegant. A tile is defined as a mix of industry standard benchmarks that emulate common business applications (e.g., e-mail, database, web server). The number of tiles running on a

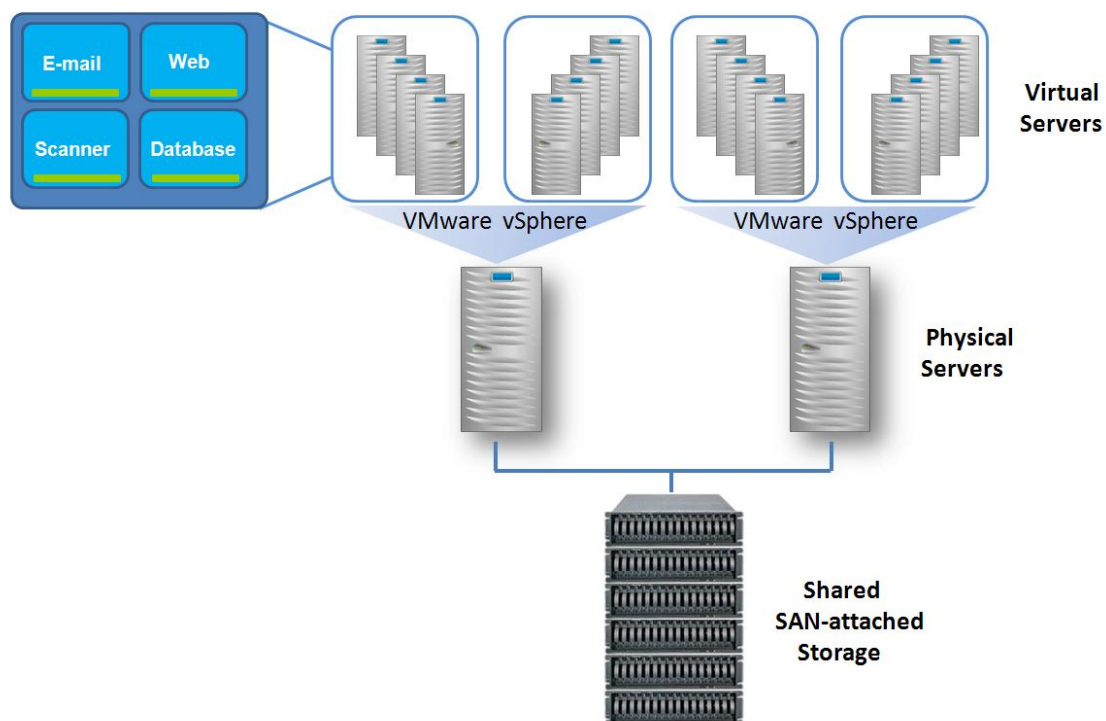
single machine is increased until the server runs out of performance. A score is derived so that IT managers can compare servers with a focus on their performance capabilities when running virtualized applications.

The IBM BladeCenter HS22 used during this ESG Lab Validation has an excellent published VMmark score of 24.05@ 17 Tiles.² At a high level, this means that the IBM BladeCenter HS22 did 24.05 times more work than the dual processor, single core server that VMware used as a reference when VMmark was first released in 2007.

A Mixed Real-world Storage Benchmark Methodology

While VMmark is well suited for understanding the performance of a mix of application running on a single server, it was not designed to assess what happens when a mix of applications are run on multiple servers sharing a single storage system. VMmark tends to stress server internals more than it does the storage system. The methodology presented in the balance of this report was designed to stress the storage system more than the servers. Taking a cue from the VMmark methodology, a tile-based concept was used. As shown in Figure 5, each tile 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 tiles was increased from one to four.

Figure 5. ESG Lab Tile-Based Storage Benchmarking




The difference between the server-focused VMmark benchmarking and storage-focused ESG Lab benchmarking is shown in Figure 6. Note how VMmark testing is performed with a single server, often attached to multiple storage systems. As a matter of fact, the IBM BladeCenter HS22 VMmark results presented earlier in this report were achieved with a pair of IBM System Storage DS4700 arrays.³ In other words, when vendors publish VMmark results, they make sure there is plenty of storage available so they can record the highest VMmark server score. This provides IT managers with a fair comparison of the performance capabilities of competitive server technologies.

² The full disclosure for the IBM BladeCenter HS22 report is available at <http://www.vmware.com/files/pdf/vmmark/VMmark-IBM-2009-06-30-HS22.pdf>. To learn more about VMmark, including a full list of published results, go to <http://www.vmware.com/products/vmmark/>

³ The IBM System Storage DS4700 is the previous generation of the IBM System Storage DS5020 examined by ESG Lab in this report.

ESG Lab storage-focused benchmarking uses a different approach. Instead of testing with a single server and more than enough storage, multiple servers are attached to a single storage system. Rather than running application level benchmarks which stress the CPU and memory of the server, lower level industry standard benchmarks are used with a goal of measuring the maximum mixed workload capabilities of a single storage system.

Figure 6. Server-focused VMmark vs. Storage-focused ESG Lab Benchmarking



Mixed Workloads

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

- **E-Mail:** The Microsoft Jetstress utility was used to generate e-mail traffic. Similar to the Microsoft LoadSimm utility used in the VMmark benchmark, Jetstress simulates the activity of typical Microsoft Exchange users as they send and read e-mails, make appointments, and manage to-do lists. The Jetstress utility is however a more light-weight utility than LoadSimm. Using the underlying Jet Engine database, Jetstress was designed to focus on storage performance.
- **Database:** The Orion utility from Oracle was used to generate database traffic. Much like Jetstress, Orion is a lightweight tool that is ideally suited for measuring storage performance. Orion 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 an Oracle database. Orion is typically used to measure two types of database activity: response-time sensitive online transaction processing (OLTP) and bandwidth sensitive online analytic processing (OLAP).
- **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.⁴ Iometer tests were performed on Windows physical drives running over VMware raw mapped devices.

⁴ Web server Iometer (www.sourceforge.net/projects/iometer) workload definitions are included in a results file excerpt as Figure 13.


- **Scan/read:** The Iometer utility was used to generate a single stream of 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 tile 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, with the Jetstress e-mail test taking the longest to complete (approximately three hours). The settings for each of the industry standard benchmarks are documented in the appendix.

Test Bed

VMware vSphere version 4.0 was installed on a pair of powerful IBM BladeCenter HS22 blades, each with a pair of quad-core processors and a QLogic 8 Gb Fibre Channel CFFh and CIOv expansion card providing four ports of connectivity per blade. A DS5020 Express with 112 15K RPM FC drives was connected to the servers through four QLogic 8 Gb SAN Switch Module for IBM BladeCenter as shown in Figure 7.

Figure 7. 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 Oracle, web server and scan/read workloads ran against four drive RAID-10 volumes. The operating system volumes (Vmdk/OS) 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 spread evenly over two VMware host groups with a multipath policy of most recently used (MRU).⁵

Table 1: Drive Configuration

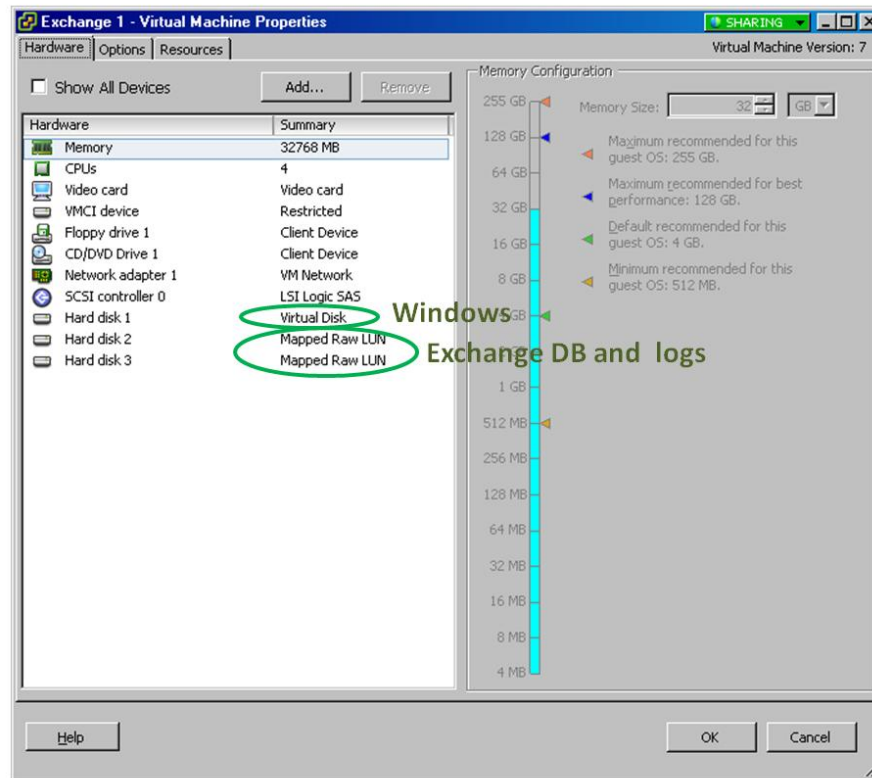
Application	Number of LUNs	Number of Drives	Usable Capacity (GB)
Exchange DB	4	32	4,356
Exchange Log	4	16	2,231
Oracle	4	16	2,231
Web Server	4	16	2,231
Scan/Read	4	16	2,231
Vmdk/OS	4	16	3,346
Total	24	112	16,626

⁵ For more detail, see the Appendix.

Configuring Virtual Machines

Each of the Microsoft Exchange machines was configured with four virtual CPU cores, 32 GB of RAM, a virtual disk over VMFS for the operating system, and two mapped raw LUNs. DS5020 Express disk capacity was used for all storage capacity including VMware virtual disk files (VMDK), Windows Server 2008 operating system images, application executables, and application data. All of the application data volumes under test were configured as mapped raw LUNs (also known as raw device mapped, or RDM volumes). The configuration of one of the four virtual machines used for e-mail testing is shown in Figure 8. Note how three hard disks have been configured. One virtual disk for the operating system and two mapped raw LUNs for the Exchange database and Logs.

Figure 8. Virtual Server Configuration



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 virtual server environment, this benchmark was designed to assess how real-world applications behave when running on multiple virtualized servers sharing a single storage system.

The Results

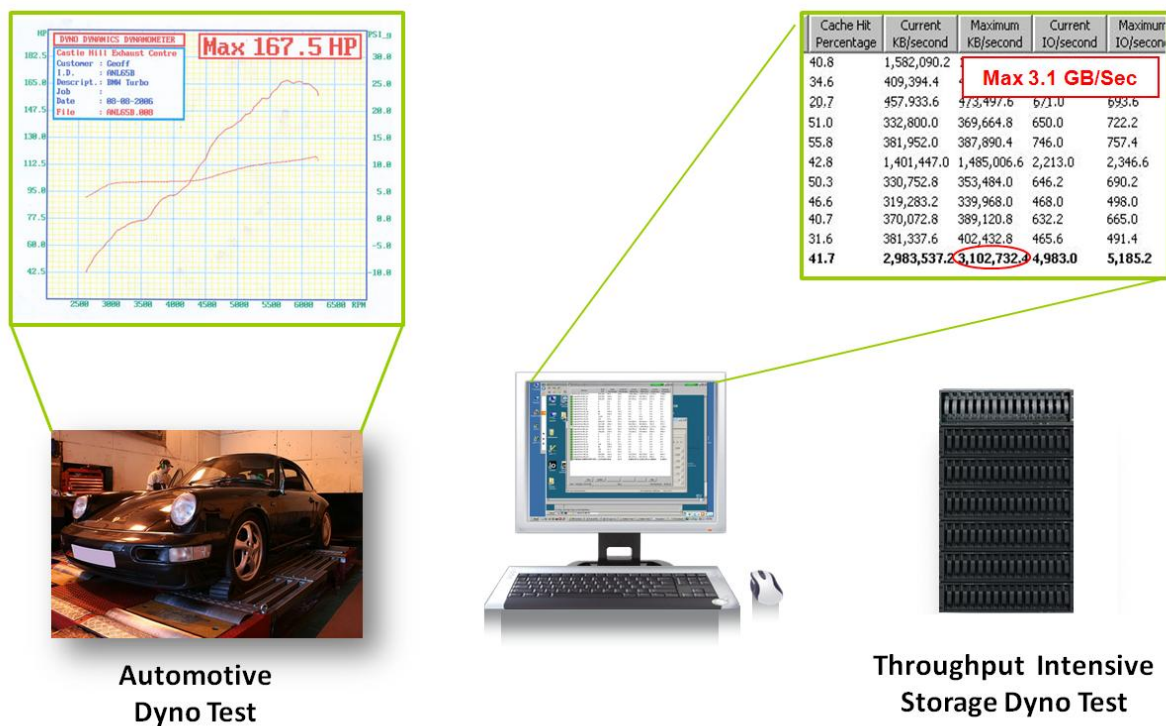
In a way, storage system benchmark testing is like an analysis of the performance of a car. Specifications including horsepower and acceleration from 0 to 60 are a good first pass indicator of a car’s performance. But while specifications provide a good starting point, there are a variety of other factors 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 the eight 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.⁶ 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 9, ESG Lab recorded a maximum throughput of 3.1 GB/sec.

Figure 9. 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.

⁶ The configuration and methodology that was used during characterization testing is described in the Appendix.

- 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 from eight logical devices through eight FC interfaces.⁷
- ESG Lab recorded a peak aggregate throughput of 3.1 GB/sec in a VMware vSphere environment.
- Forty-two percent 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 VMware vSphere environment achieved 3.1 GB/sec of aggregate large block sequential read throughput.

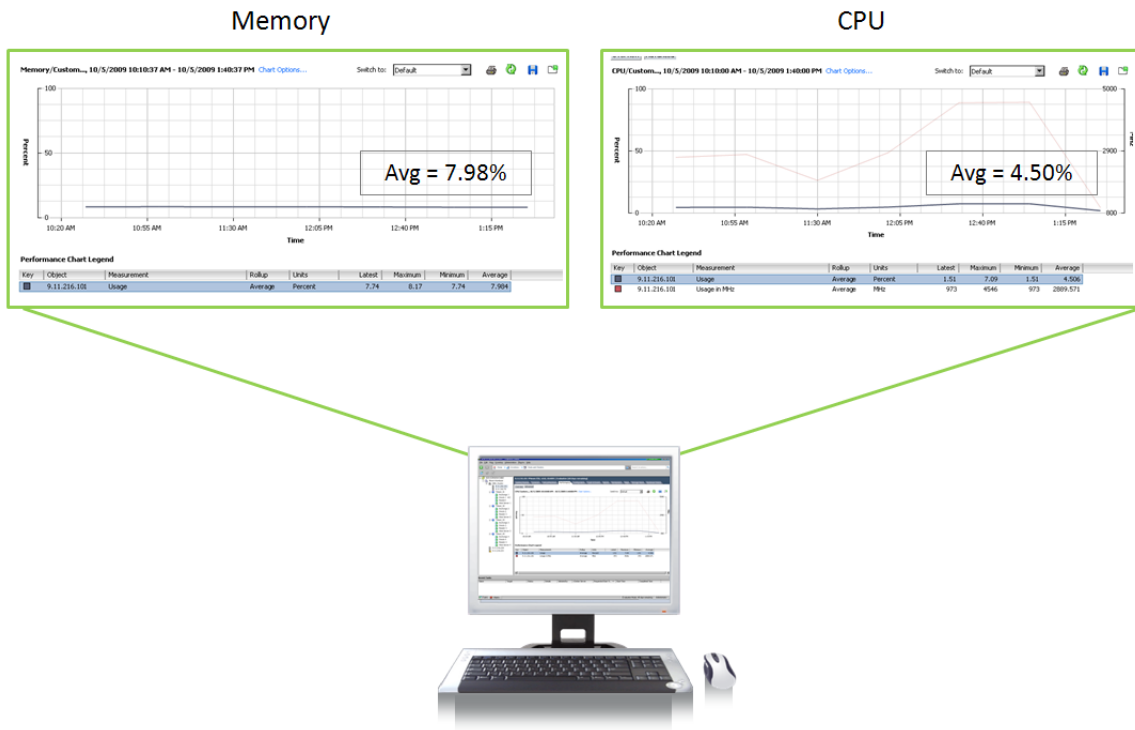
In ESG Lab's experience, these are excellent results for a dual controller modular storage system. As a matter of fact, these results provide an excellent early indication that the DS5020 Express is well suited for virtual server consolidation and mixed real-world business applications.

⁷ For more on the characterization configuration and methodology please see the Appendix.

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. Memory and CPU utilization as reported by the VMware Infrastructure Manager are shown in Figure 10.

Figure 10. Low Memory and CPU Utilization



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

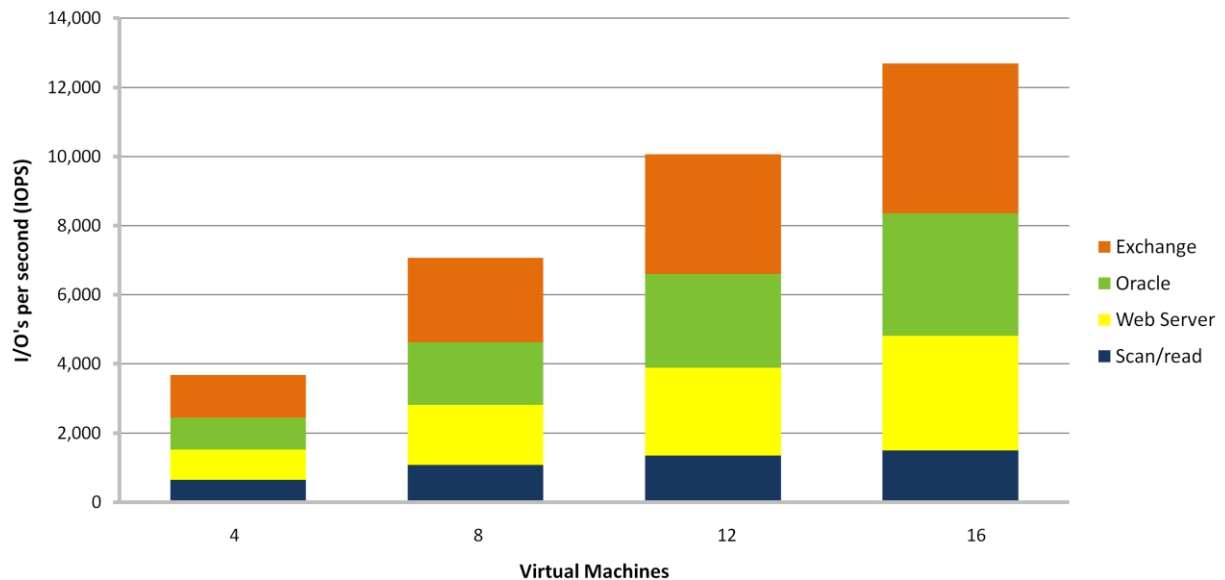
Mixed Real-world IOPS Scalability

I/Os 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.

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 measured IOPS performance 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. The IOPS scalability during the peak period of mixed workload activity is shown in Figure 11.

Figure 11. DS5020 Express Mixed Workload Scalability



What the Numbers Mean

- IOPS varied throughout the mixed workload test with peaks occurring during the Orion small IOPs phase and towards the end as the Jetstress utility as it performed a database consistency check.
- A peak of 19,046 and a steady state of 12,688 IOPS were recorded during the four tile run.
- IOPS scaled in a near-linear fashion as mixed real-world application traffic increased from four through sixteen virtual servers.

Why This Matters

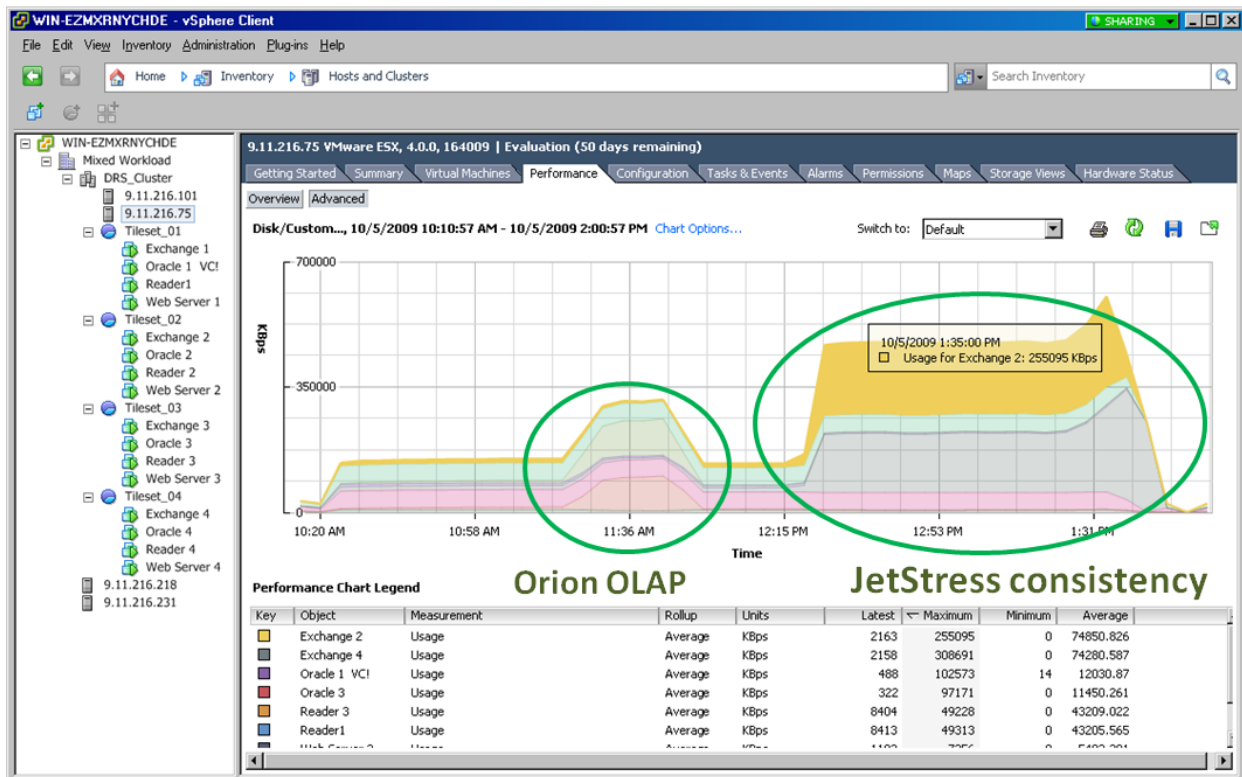
Predictable performance scalability is a critical concern when a mix of applications shares a storage system. A burst of IO activity in one application (e.g., a database consistency check) can lead to poor response times, lost productivity, and, in the worst case, lost revenue.

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.

Handling Throughput Spikes with Ease

As noticed during IOPS monitoring, peaks of throughput activity could be correlated to the periodic behavior of real-world applications. Two bursts of aggregate throughput were observed: the first during the Oracle large MBPS test, which simulates a throughput intensive OLAP application, and the second during the Jetstress database consistency check. A VMware vSphere view of mixed workload performance on one of the HS22 blades is shown in Figure 12.

Figure 12. Peak Throughput (One Server, Four Active Tiles, Stacked VM View)



What the Numbers Mean

- An aggregate throughput level of 1.21 GB/sec was recorded as mixed, real-world applications were run on 16 virtual machines sharing a single DS5020 Express storage system (605 MB/sec for one of the two physical servers is shown in Figure 12).
- As throughput intensified during the Oracle Orion OLAP test phase, bandwidth utilization for other mixed workloads operating in parallel remained steady.

Why This Matters

Storage benchmarks typically focus on response time sensitive interactive workloads or throughput intensive sequential workloads, yet mixed real-world applications in virtualized environments are usually a mix of both. A burst of activity due to a search and index operation, a database query, a backup job, or a video stream can be extremely throughput intensive. Deploying more storage systems, or more hardware within each storage system, is one way to avoid the potential performance impact of a throughput intensive workload in a mixed environment, but this increases cost and complexity and defeats the goal of shared storage consolidation. ESG Lab observed a peak aggregate throughput of 1.21 GB/sec as a throughput intensive Jetstress e-mail database consistency check was running—while other applications ran in parallel with predictably good response times.

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

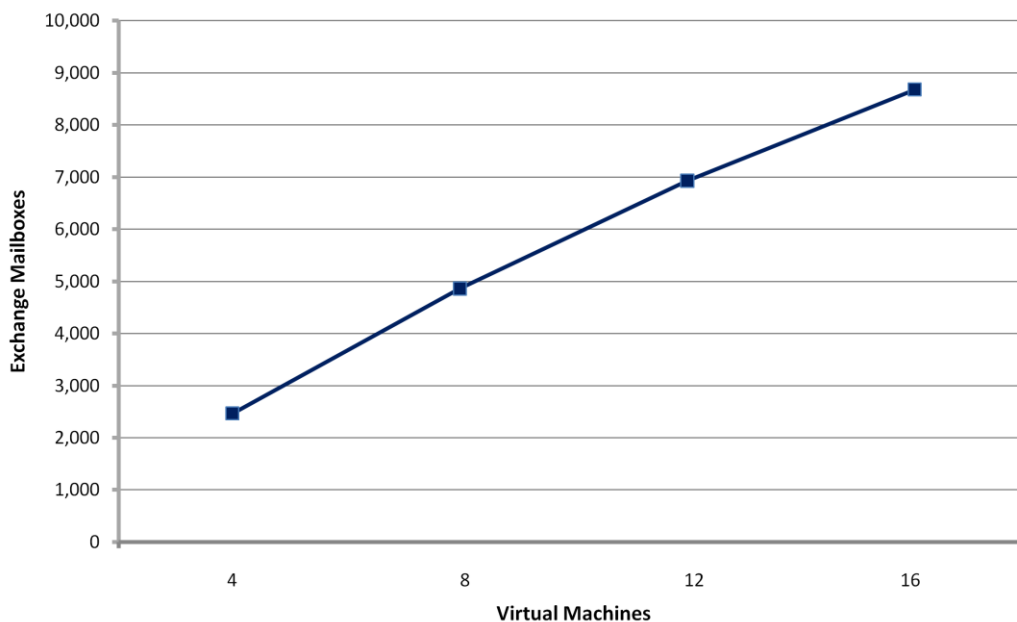
The Microsoft Jetstress tool was used to see how many simulated e-mail users could be supported by the DS5020 Express resources allocated for the Exchange application. The number of IOPS and response time for each database and log volume was recorded at the end of each Jetstress run. A response time goal of 20 milliseconds or less for DB reads and 5 milliseconds or less for log writes 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.

ESG used the following IBM guidelines from an IBM report describing the results of an IBM System Storage DS4800 Mailbox Jetstress Analysis report to interpret the results:

In an enterprise Exchange 2007 environment, performance is usually designed around a 0.5 IOPS user profile, which is equivalent to a very heavy Exchange user. While disk performance varies, generally you should calculate based on a one hundred IOPS per disk metric, which is a conservative starting point, and tune from there for your specific environment.⁸

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.⁹ Based on a 0.5 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 13, Figure 14 and Table 2.

Figure 13. Mixed E-mail Scalability (Number of Mailboxes)



⁸ IBM System Storage DS4800 Exchange Server 2007 15,000 Mailbox JetStress Analysis, David Hartman and David West, November 2007, <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101123>

⁹ A sample JetStress log is included in the Appendix as Figure 11.

Figure 14. Mixed E-mail Scalability (Response Time)

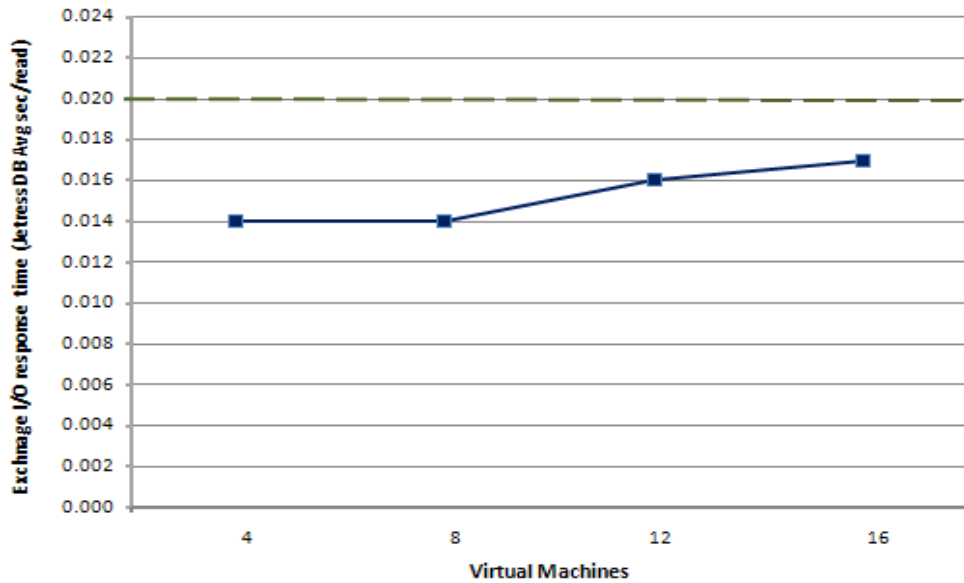


Table 2: Jetstress Performance Results (One Through Four Tiles)

Virtual Machines	Achieved IO per Seconds	Users (0.5 Profile)	DB Avg. Disk Sec/Read
4	1,233	2,466	.014
8	2,431	4,862	.014
12	3,465	6,930	.016
16	4,340	8,680	.017

What the Numbers Mean

- The single tile mixed application test supported 2,466 Exchange users with an average DB disk response time of 14 milliseconds.
- Performance scaled in a near-linear fashion to 8,680 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 14.
- The four tile test, which produced 4,340 IOPS over 32 database drives, delivered 135 IOPS per drive—well above the conservative IBM guideline of 100 IOPS per drive.

Oracle Orion

The Oracle Orion utility was used to measure small transfer (8 KB) IOPS and response time and large transfer (1 MB) throughput. The small results are used to predict the performance and scalability of response time sensitive interactive database applications (e.g. OLTP). The large results are used to predict the performance of throughput intensive database mining and decision support systems (DSS).

ESG used the following guidelines from presentations presented at Oracle OpenWorld in November 2007 to interpret the results:

Target 5-10 millisecond for response time critical IO. Start by assuming 30 IOPS per disk for OLTP and 20 MB/sec per disk in DSS. This is way below the theoretical value but allows for media repair etc.¹⁰

For new or non-existing applications, use business rules or data model transaction profiles flow to understand “what is a transaction,” and then extrapolate for transactions per second or hour. Optionally you can use the numbers we have seen in our consulting gigs. Note that these are just guideline values. Use the following as basic guidelines for OLTP:

Low transaction system – 1,000 IOPS or 200MBytes/s

Medium transaction system – 5,000 IOPS or 600 Mbytes/s

High-end transaction system – 10,000 IOPS or 1Gbytes/s <- almost rarely achievable and usually TPC-C type workloads¹¹

The results for the four tile Orion test are summarized in Table 3. A sample Orion report is shown in the Appendix.

Table 3: Orion Four Tile Performance Results

Tile	Small IOPS	Large MBPS	Small Latency (ms)
1	882	108	5.42
2	899	111	5.54
3	875	106	5.43
4	883	108	5.35
Total	3,539	433	5.43

What the Numbers Mean

- The four tile test achieved a grand total of 3,539 small IOPS and 433 large MBPS while the system was simultaneously running a mix of real-world application workloads.
- Using Oracle’s back of the envelope sizing guidelines, this level of IO activity is significantly higher than a typical “*low transaction system*” and nearly represents a “*medium transaction system*.”
- The total number of small IOPS processed during the busy four tile test yielded an excellent rate of 222 small IOPS per drive, which dwarfs the extremely conservative Oracle planning guideline of 30 IOPS per drive.
- Orion reported an average latency of 5.43 milliseconds for the small IOPs workload. Given the Oracle guidance of 5 to 10 milliseconds, ESG Lab believes that these are excellent results—especially given the mix of IO intensive workloads that were being serviced by the DS5020 Express in parallel.

¹⁰ Current trends in Database Performance, Andrew Holdsworth, Oracle OpenWorld, Nov 2007, http://www.oracle.com/technology/dep/perform/pdf/PerfTrends_Holdsworth.pdf

¹¹ Back of the Envelope Database Storage Design, Nitin Vengurlekar, RAC/ASM Development, Oracle Open World, Nov 2007, <http://www.oracle.com/technology/products/database/asm/pdf/back%20of%20the%20env%20by%20nitin%20oow%202007.pdf>

Web Server and Scan/Read

Performance results as reported by the Iometer utility for the web server and scan/read workloads executing within virtual machines during the four tile test are listed in Table 4.

Table 4: Iometer Four Tile Performance Results

Tile	Web Server (IOPs)	Scan/Read (MB/Sec)
1	846	96
2	838	97
3	800	90
4	833	91
Total	3,317	374

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 at least 90 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 virtualization solutions like VMware, 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 BladeCenter HS22 servers with a goal of analyzing the performance difference between the DS5020 Express and the DS3950 Express. The results are shown in Figure 15 and Table 5.

Figure 15. DS3950 Express vs. DS5020 Express

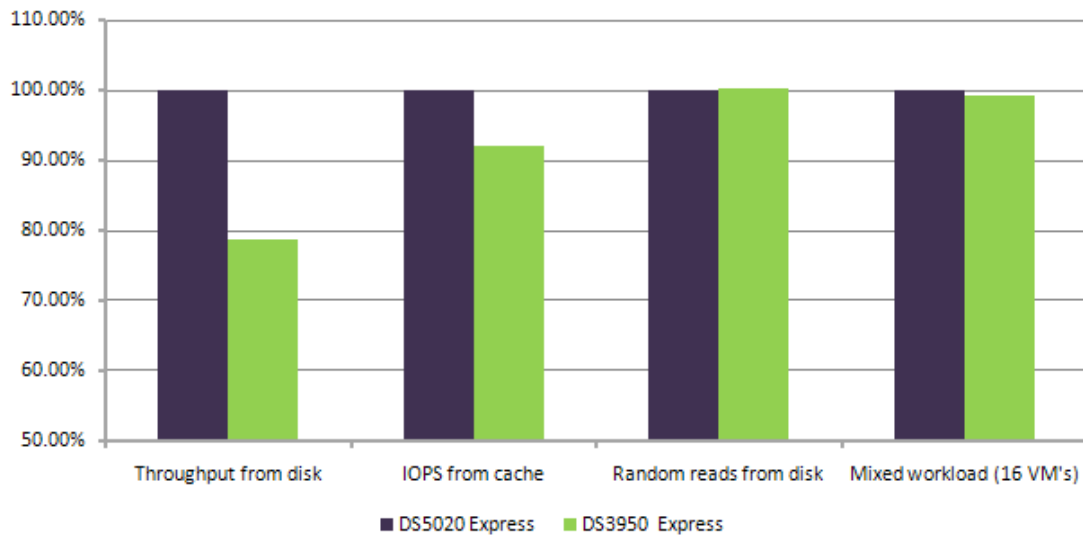


Table 5: DS3950 Express vs. DS5020 Express

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

- ☑ 3.1 GB/sec of aggregate throughput was sustained during characterization testing within a VMware vSphere environment.
- ☑ Mixed real-world application workloads running simultaneously within sixteen virtual machines deployed over two IBM BladeCenter HS22 servers provided the performance needed to concurrently support:
 - **8,680 mailboxes** using the Microsoft Exchange Jetstress utility
 - **and 3,593 small database IOs per second** using the Oracle Orion utility
 - **and 433 MB/sec of throughput** for large OLAP Oracle Orion operations
 - **and 3,317 simulated web server IOPs**
 - **and 374 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 Oracle 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 (16 to 17 millisecond for Jetstress DB reads, 5.35 to 5.54 milliseconds for Oracle Orion small IOPS).
- ☑ 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 VMware 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 decoupling 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. Taking a cue from the VMmark benchmark from VMware, a “tile” concept was used during the design of this test. Each “tile” was composed of four applications, each running in its own virtual machine. The server horsepower of a pair of IBM BladeCenter HS22 servers, with an excellent published VMmark score of 24.05@17 tiles, was used to drive up to four tiles and sixteen virtual applications in parallel. ESG believes that the results of this storage-focused benchmark complement the excellent server-focused results of the IBM BladeCenter HS22 VMmark test.

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 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 a confirmation that the DS5020 Express test bed can deliver up to 3.1 GB/sec of raw aggregate throughput in a VMware vSphere 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 8,680 simulated Exchange users *and* 4,144 Oracle Orion small database IOs per second *and* 433 MB/sec of throughput for large OLAP Oracle Orion operations *and* 4,551 simulated web server IOPs *and* 374 MB/sec of throughput for bandwidth intensive streams of read traffic—all while delivering predictably fast response times. Testing on the DS5020 Express with the same number of drives and less host connections indicates that the IBM System Storage DS3950 Express delivers similar levels of performance for the mix of applications tested by ESG Lab.

ESG Lab is pleased to report that the combination of IBM System Storage DS5020 Express, IBM BladeCenter HS22 servers and QLogic 8 Gb Fibre Channel switch and expansion card infrastructure delivers the performance needed to meet the needs of a mix of real-world business applications running within a VMware vSphere enabled virtual infrastructure.

Appendix

Table 6. Test Bed Overview

Storage	
IBM DS5020 Express, Firmware: 07.60.08.00	8 drive trays, 112 15K RPM FC drives
Server	
Two IBM BladeCenter HS22	CPU type: Intel Xeon x5570 (2 sockets/4 cores @2.93 GHz) CPUs per server: 2 , total CPU cores: 16 RAM: 96 GB DDR-3
SAN	
QLogic 8 Gb SAN Switch Module for IBM BladeCenter	
Host Bus Adapters	
QLogic 8Gb Fibre Channel Expansion Cards (CIOv and CFFh) for IBM BladeCenter	
Virtualization Software and Guest Operating Systems	
Server Virtualization	VMware vSphere ,version 4.0.0
Guest OS	Windows Server 2008 R1, Build 6002, Service Pack 2

Table 7. Bill of Materials

IBM Code	Description	Quantity
8852HC1	BCH Chassis	1
39Y9314	MSIMS	2
43W4401	CISCO 2012 switch	4
44X1911	QLOGIC 8G 20 port switch	4
7870AC1	HS22 BLADE 2*2.93	2
44X1940	QLOGIC 8G FC Expansion Card(CFFh) for IBM Blade Center	2
44X1945	QLOGIC 8G FC Expansion Card (CIOv) for IBM Blade Center	2
1814-20A	IBM DS5020 Storage System	1
1814-52A	IBM DS520 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
7804	VMware ESX Host Attach Kit	1
8701	4 Host Partitions	1

Table 8. Benchmark Utilities/Workload Generators

Characterization	Iometer, version 2006.07.27 Dymomo clients ran within eight guest VM's running on a pair of IBM BladeCenter HS22 servers with Windows Server 2008 R1, SP2 as the 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 raw device mapped volumes in a VMware vSphere 4.0 environment. The HS22 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 I/O's per physical drive.
E-Mail	Microsoft Jetstress, version 08.02.0060.000 <ul style="list-style-type: none"> • Mailboxes – 2,280 • Mailbox size – 300 MB • IOPS per mailbox – 0.5 • Thread – 32 • Log Buffers – 9000 • Min DB cache – 64 MB • Max DB cache – 512 MB • Insert operations – 40% • Delete operations – 30% • Replace operations – 5% • Read operations – 25% • Lazy commits – 55%
Database Workload Generator	Oracle Orion, version 10.2.0.1.0 <ul style="list-style-type: none"> • Small IO size: 8 KB • Large IO size: 1024 KB • IO Types: Small Random, Large Random • Simulated Array Type: RAID 0 • Num_disks: 5 • Stripe Depth: 1024 KB • Write: 30% • Duration for each Data Point: 150 seconds
Web Server	Iometer, version 2006.07.27 Four workers, four outstanding I/Os per physical drive 100% random reads, assorted block sizes ¹²
Scanner/Reader	Iometer, version 2006.07.27 One worker, one outstanding I/O per physical drive 100% 256 KB sequential reads

¹² See Figure 13 for workload details.

Figure 16. E-mail Results

This is an example of the output created by the Jetstress utility. It shows the performance for one of four Jetstress tests running in parallel. Specifically, this report was created by the Jetstress utility running on a virtual machine within the fourth tile of the four tile test.

Microsoft Exchange Server **Jetstress**

Performance Test Result Report

Test Summary

Overall Test Result **Pass**

Machine Name JS-01
Test Description Blade Center H VMware mixed workload test.
Test Start Time 10/13/2009 3:41:47 PM
Test End Time 10/13/2009 5:43:57 PM
Jetstress Version 08.02.0060.000
Ese Version 08.01.0240.005
Operating System Windows Server (R) 2008 Enterprise without Hyper-V Service Pack 2 (6.0.6002.131072)
Performance Log C:\Downloads\Jetstress\Logs\Performance_2009_10_13_15_41_49.blg
C:\Downloads\Jetstress\Logs\DBChecksum_2009_10_13_17_43_57.blg

Database Sizing and Throughput

Achieved I/O per Second 1077.105
Target I/O per Second 1140
Initial database size 1010127486976
Final database size 1013138997248
Database files (count) 1

Jetstress System Parameters

Thread count 32 (per-storage group)
Log buffers 9000
Minimum database cache 32.0 MB
Maximum database cache 256.0 MB
Insert operations 40%
Delete operations 30%
Replace operations 5%
Read operations 25%
Lazy commits 55%

Disk Subsystem Performance

LogicalDisk	Avg. Disk sec/Read	Avg. Disk sec/Write	Disk Reads/sec	Disk Writes/sec	Avg. Disk Bytes/Write
Database (E:)	0.017	0.015	630.229	481.876	(n/a)
Log (F:)	0.000	0.000	0.000	252.745	5666.955

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.539	0.495	4.089
Available MBytes	30632.917	30629.000	30643.000
Free System Page Table Entries	33559507.904	33559394.000	33559648.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	36980966.400	36954112.000	37036032.000
Pool Paged Bytes	103446109.867	103424000.000	103591936.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/8/2009 12:41:13 PM -- Jetstress testing begins ...
 10/8/2009 12:41:31 PM -- Prepare testing begins ...

```

10/8/2009 12:41:33 PM -- Attaching databases ...
10/8/2009 12:41:33 PM -- Prepare testing ends.
10/13/2009 3:41:47 PM -- Jetstress testing begins ...
10/13/2009 3:41:47 PM -- Prepare testing begins ...
10/13/2009 3:41:48 PM -- Attaching databases ...
10/13/2009 3:41:48 PM -- Prepare testing ends.
10/13/2009 3:41:48 PM -- Dispatching transactions begins ...
10/13/2009 3:41:48 PM -- Database cache settings: (minimum: 32.0 MB, maximum: 256.0 MB)
10/13/2009 3:41:48 PM -- Database flush thresholds: (start: 2.6 MB, stop: 5.1 MB)
10/13/2009 3:41:49 PM -- Database read latency thresholds: (average: 0.02 seconds/read, maximum: 0.05 seconds/read).
10/13/2009 3:41:49 PM -- Log write latency thresholds: (average: 0.01 seconds/write, maximum: 0.05 seconds/write).
10/13/2009 3:41:50 PM -- Operation mix: Sessions 32, Inserts 40%, Deletes 30%, Replaces 5%, Reads 25%, Lazy Commits 55%.
10/13/2009 3:41:50 PM -- Performance logging begins (interval: 15000 ms).
10/13/2009 3:41:50 PM -- Attaining prerequisites:
10/13/2009 3:43:56 PM -- \\MSEExchange Database(JetstressWin)\Database Cache Size, Last: 242573300.0 (lower bound: 241591900.0, upper bound: none)
10/13/2009 5:43:56 PM -- Performance logging ends.
10/13/2009 5:43:56 PM -- JetInterop batch transaction stats: 103218.
10/13/2009 5:43:57 PM -- Dispatching transactions ends.
10/13/2009 5:43:57 PM -- Shutting down databases ...
10/13/2009 5:43:57 PM -- Instance2924.1 (complete)
10/13/2009 5:43:58 PM -- Performance logging begins (interval: 30000 ms).
10/13/2009 5:43:58 PM -- Verifying database checksums ...
10/13/2009 5:57:56 PM -- E: (19% processed)
10/13/2009 5:57:56 PM -- Verifying log checksums ...
10/13/2009 5:57:56 PM -- F:\ (0 logs passed)
10/13/2009 5:57:56 PM -- C:\Downloads\Jetstress\Logs\Performance 2009 10 13 15 41 49.blg has 488 samples.
10/13/2009 5:57:56 PM -- Creating test report ...
10/13/2009 5:57:57 PM -- Volume E: has 0.0165 for Avg. Disk sec/Read.
10/13/2009 5:57:57 PM -- Volume F: has 0.0004 for Avg. Disk sec/Write.
10/13/2009 5:57:57 PM -- Volume F: has 0.0000 for Avg. Disk sec/Read.
10/13/2009 5:57:57 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
10/13/2009 5:57:57 PM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
10/13/2009 5:57:57 PM -- C:\Downloads\Jetstress\Logs\Performance 2009 10 13 15 41 49.xml has 479 samples queried.

```

Figure 17. Database Results

This is an example of the output created by the Oracle Orion utility. It shows the performance for one of eight Orion tests running in parallel. Specifically, this report was created by the Orion utility running on a virtual machine within the fourth tile of the four tile test.

ORION VERSION 10.2.0.1.0

Commandline:

```
-run advanced -testname vmware -num_disks 5 -size_small 8 -size_large 1024 -type rand -simulate raid0 -write 30 -duration 150 -matrix basic
```

This maps to this test:

Test: vmware

Small IO size: 8 KB

Large IO size: 1024 KB

IO Types: Small Random IOs, Large Random IOs

Simulated Array Type: RAID 0

Stripe Depth: 1024 KB

Write: 30%

Cache Size: Not Entered

Duration for each Data Point: 150 seconds

Small Columns:, 0

Large Columns:, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10

Total Data Points: 36

Name: \\.\f: Size: 1924071424

1 FILES found.

Maximum Large MBPS=108.27 @ Small=0 and Large=9

Maximum Small IOPS=882 @ Small=25 and Large=0

Minimum Small Latency=5.42 @ Small=1 and Large=0

Figure 18. Web Server Results

This is an example of the output created by the Iometer utility after a web server test run. This example shows the performance of the four web server tests which ran in parallel during the mixed workload four tile test.

Test Type	Test Description												
0	ESG Lab Characterization												
'Version													
2006.07.27													
'Time Stamp													
2009-10-13 18:11:52:531													
'Access specifications													
'Access specification no default assignment													
Web Server													
1													
'size	% of size	% reads	% random	delay	burst	align	reply						
512	22	100	100	0	0	1	0						
1024	15	100	100	0	0	1	0						
2048	8	100	100	0	0	1	0						
4096	23	100	100	0	0	1	0						
8192	15	100	100	0	0	1	0						
16384	2	100	100	0	0	1	0						
32768	6	100	100	0	0	1	0						
65536	7	100	100	0	0	1	0						
131072	1	100	100	0	0	1	0						
524288	1	100	100	0	0	1	0						
'End access specifications													
'Results													
Target Type	Target Name	Access \$	# Managers	# Workers	# Disks	I/Ops	Read I/Ops	Write I/Ops	MBps	Read MBps	Write MBps		
ALL	All	Web Ser	4	16	16	3317.3	3317.31596	0	50.73714	50.73714	0		
MANAGER	WB-01	Web Server		4	4	846.03	846.027873	0	12.35636	12.356357	0		
PROCESSOR	CPU 0												
WORKER	Worker 1	Web Server			1	211.43	211.430454	0	3.238575	3.238575	0		
DISK	PHYSICALDRIVE:1					211.43	211.430454	0	3.238575	3.238575	0		
WORKER	Worker 2	Web Server			1	211.47	211.473324	0	3.236585	3.236585	0		
DISK	PHYSICALDRIVE:1					211.47	211.473324	0	3.236585	3.236585	0		
WORKER	Worker 3	Web Server			1	211.47	211.468349	0	3.249634	3.249634	0		
DISK	PHYSICALDRIVE:1					211.47	211.468349	0	3.249634	3.249634	0		
WORKER	Worker 4	Web Server			1	211.66	211.655746	0	3.232103	3.232103	0		
DISK	PHYSICALDRIVE:1					211.66	211.655746	0	3.232103	3.232103	0		
MANAGER	WB-02	Web Server		4	4	838.6	838.601953	0	12.83685	12.836852	0		
PROCESSOR	CPU 0												
WORKER	Worker 1	Web Server			1	209.59	209.590166	0	3.217553	3.217553	0		
DISK	PHYSICALDRIVE:1					209.59	209.590166	0	3.217553	3.217553	0		
WORKER	Worker 2	Web Server			1	209.66	209.656621	0	3.205382	3.205382	0		
DISK	PHYSICALDRIVE:1					209.66	209.656621	0	3.205382	3.205382	0		
WORKER	Worker 3	Web Server			1	209.64	209.639302	0	3.204907	3.204907	0		
DISK	PHYSICALDRIVE:1					209.64	209.639302	0	3.204907	3.204907	0		
WORKER	Worker 4	Web Server			1	209.72	209.715865	0	3.20841	3.20841	0		
DISK	PHYSICALDRIVE:1					209.72	209.715865	0	3.20841	3.20841	0		
MANAGER	WB-03	Web Server		4	4	799.85	799.848937	0	12.22021	12.220214	0		
PROCESSOR	CPU 0												
WORKER	Worker 1	Web Server			1	199.95	199.953638	0	3.051417	3.051417	0		
DISK	PHYSICALDRIVE:1					199.95	199.953638	0	3.051417	3.051417	0		
WORKER	Worker 2	Web Server			1	199.84	199.840398	0	3.053159	3.053159	0		
DISK	PHYSICALDRIVE:1					199.84	199.840398	0	3.053159	3.053159	0		
WORKER	Worker 3	Web Server			1	200.09	200.090605	0	3.064349	3.064349	0		
DISK	PHYSICALDRIVE:1					200.09	200.090605	0	3.064349	3.064349	0		
WORKER	Worker 4	Web Server			1	199.96	199.964295	0	3.05129	3.05129	0		
DISK	PHYSICALDRIVE:1					199.96	199.964295	0	3.05129	3.05129	0		
MANAGER	WB-04	Web Server		4	4	832.84	832.837202	0	12.72312	12.723117	0		
PROCESSOR	CPU 0												
WORKER	Worker 1	Web Server			1	208.1	208.097093	0	3.181846	3.181846	0		
DISK	PHYSICALDRIVE:1					208.1	208.097093	0	3.181846	3.181846	0		
WORKER	Worker 2	Web Server			1	208.33	208.334017	0	3.167798	3.167798	0		
DISK	PHYSICALDRIVE:1					208.33	208.334017	0	3.167798	3.167798	0		
WORKER	Worker 3	Web Server			1	208.15	208.147559	0	3.184255	3.184255	0		
DISK	PHYSICALDRIVE:1					208.15	208.147559	0	3.184255	3.184255	0		
WORKER	Worker 4	Web Server			1	208.26	208.258532	0	3.189218	3.189218	0		
DISK	PHYSICALDRIVE:1					208.26	208.258532	0	3.189218	3.189218	0		
'Time Stamp													
2009-10-13 22:11:57:875													

Figure 19. Scan/Read Results

This is an example of the output created by the Iometer utility after a scan/read test run. It shows the performance of the four scan/read tests which ran in parallel during the mixed workload four tile test.

'Time Stamp										
2009-10-13 18:11:26:125										
'Access specifications										
'Access specification name	default assignment									
scanner	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	# Managers	# Workers	# Disks	I/Ops	Read I/Ops	Write I/Op	MBps	
ALL	All	scanner	4	4	4	1493.762	1493.762	0	373.4405	
MANAGER	BK-01	scanner		1	1	385.4812	385.4812	0	96.3703	
PROCESSOR	CPU 0									
WORKER	Worker 1	scanner			1	385.4812	385.4812	0	96.3703	
DISK	PHYSICALDRIVE:1					385.4812	385.4812	0	96.3703	
MANAGER	BK-03	scanner		1	1	387.4058	387.4058	0	96.85146	
PROCESSOR	CPU 0									
WORKER	Worker 1	scanner			1	387.4058	387.4058	0	96.85146	
DISK	PHYSICALDRIVE:1					387.4058	387.4058	0	96.85146	
MANAGER	BK-04	scanner		1	1	354.7564	354.7564	0	88.6891	
PROCESSOR	CPU 0									
WORKER	Worker 1	scanner			1	354.7564	354.7564	0	88.6891	
DISK	PHYSICALDRIVE:1					354.7564	354.7564	0	88.6891	
MANAGER	BK-02	scanner		1	1	366.1187	366.1187	0	91.52967	
PROCESSOR	CPU 0									
WORKER	Worker 1	scanner			1	366.1187	366.1187	0	91.52967	
DISK	PHYSICALDRIVE:1					366.1187	366.1187	0	91.52967	
'Time Stamp										
2009-10-13 22:11:31:234										

Figure 20. 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: 3030303934203030313235204A93DE36

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

Storage Subsystem world-wide identifier (ID): 60080E500017B6BA00000004A93DE34

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	OS_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 VMware_01	Optimal
BK_03	10	B	Host Group VMware_01	Optimal
JS_01	2	A	Host Group VMware_01	Optimal
JS_03	8	A	Host Group VMware_01	Optimal
JS_2L	3	A	Host Group VMware_01	Optimal
JS_4L	9	A	Host Group VMware_01	Optimal
OR_02	1	A	Host Group VMware_01	Optimal
OR_04	7	A	Host Group VMware_01	Optimal
OS_01	0	A	Host Group VMware_01	Optimal
OS_03	6	A	Host Group VMware_01	Optimal
WB_02	5	A	Host Group VMware_01	Optimal
WB_04	11	B	Host Group VMware_01	Optimal
BK_02	4	B	Host Group VMware_02	Optimal
BK_04	10	A	Host Group VMware_02	Optimal
JS_02	2	B	Host Group VMware_02	Optimal
JS_04	8	B	Host Group VMware_02	Optimal
JS_1L	3	B	Host Group VMware_02	Optimal
JS_3L	9	B	Host Group VMware_02	Optimal
OR_01	1	B	Host Group VMware_02	Optimal
OR_03	7	B	Host Group VMware_02	Optimal
OS_02	0	B	Host Group VMware_02	Optimal
OS_04	6	B	Host Group VMware_02	Optimal
WB_01	5	B	Host Group VMware_02	Optimal
WB_03	11	A	Host Group VMware_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.**