



LAB VALIDATION REPORT

IBM DS5000 Series Real-World Performance for Virtualized Systems

By Brian Garrett
With Claude Bouffard

September, 2008

Table of Contents

Table of Contents	i
Introduction	1
<i>Background</i>	1
<i>The IBM DS5000 Series</i>	2
ESG Lab Validation	3
<i>A Mixed Real-world Benchmark Methodology</i>	3
<i>The Application Workloads</i>	4
<i>The Physical Test Bed</i>	5
<i>The Drive Layout</i>	5
<i>Configuring Virtual Machines</i>	6
<i>The Results</i>	7
<i>Raw Aggregate Throughput</i>	7
<i>Virtual Machine Utilization</i>	9
<i>Mixed Real-world IOPS Scalability</i>	9
<i>Handling Throughput Spikes with Ease</i>	11
<i>Mixed Application-level Performance Scalability</i>	12
ESG Lab Validation Highlights	16
Issues to Consider	16
ESG Lab's View	17
Appendix	18

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 ESG Lab report was sponsored by LSI.

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

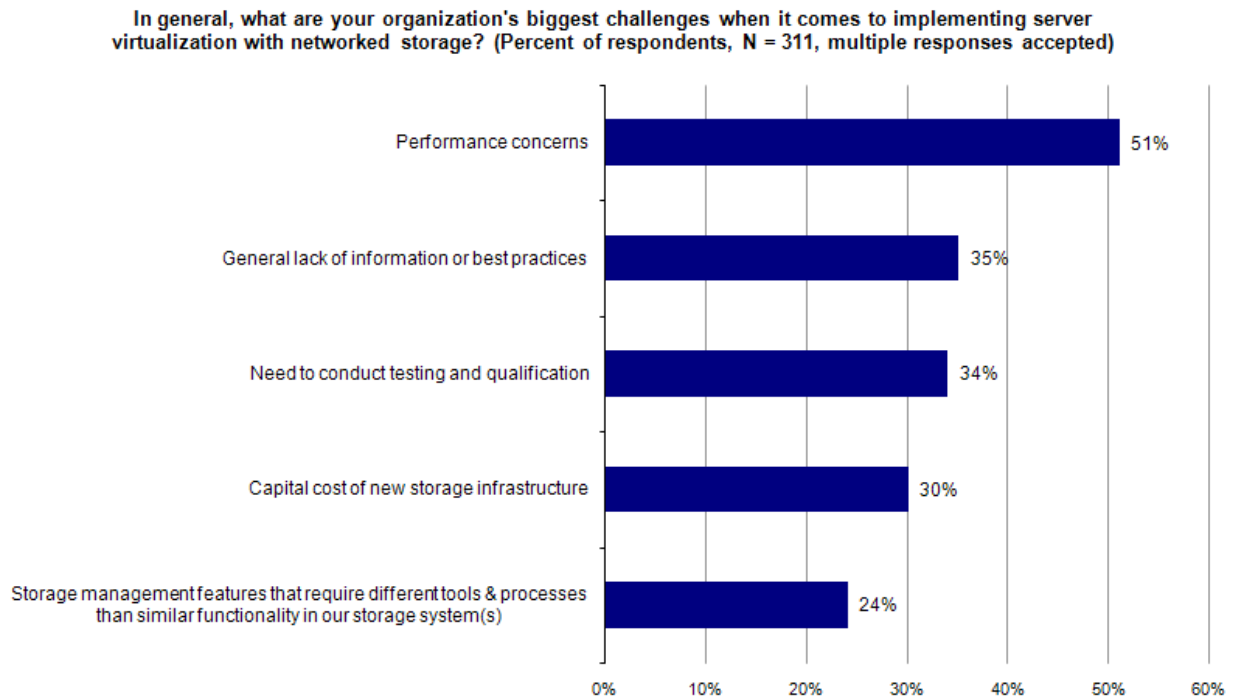
Networked storage is being deployed in conjunction with server virtualization by a growing number of organizations to consolidate, reduce costs and improve the flexibility and availability of mission-critical applications including databases and e-mail. ESG research indicates that IT managers looking 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 which was designed to assess the real-world performance capabilities of an IBM SAN-attached DS5300 storage system deployed in a highly virtualized, consolidated data center.

Background

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 server virtualization, 52% of organizations had already deployed the technology, 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




¹ ESG Research Report, *The Impact of Server Virtualization on Storage*, December, 2007

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. 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 IBM DS5000 Series

IBM has recently announced the System Storage DS5300—a turbo-charged SAN-attached storage system with up to four times the performance of the previous generation DS4800 and eight times the performance of the DS4700. The DS5300 is designed to meet the demanding performance requirements of real-world enterprise-class storage environments. With high performance that is optimized for mixed workloads, the DS5300 is designed for modular scalability (capacity and/or performance), high availability and advanced functionality including copy service and remote replication. As shown in Figure 2, the DS5300 is a dual controller system supporting up to 16 4 Gbps Fibre Channel host interfaces, up to 256 drives (FC or SATA), up to 16GB of cache and gigabytes per second of internal bandwidth.

FIGURE 2. DS5300 HIGHLIGHTS



This report examines the enterprise-class performance capabilities of the turbo-charged DS5300 including IBM’s claim that it is ideally suited to handle the demanding performance requirements of mixed real-world applications deployed in a virtual server environment. In particular, this report demonstrates how a single DS5300 supports:

- An impressive 6.2 GB/sec of sustained aggregate throughput.
- A mix of business-critical applications running on sixteen virtual machines deployed on two physical servers.
- Up to 17,512 e-mail users using the Microsoft Exchange JetStress utility **and** 9,164 small database I/Os per second using the Oracle Orion utility **and** 4,551 simulated web server IOPs and 425 MB/sec of throughput for large reads using the Iometer utility.
- Optimal consolidation and growth with fast response times and excellent performance scalability.

ESG Lab Validation

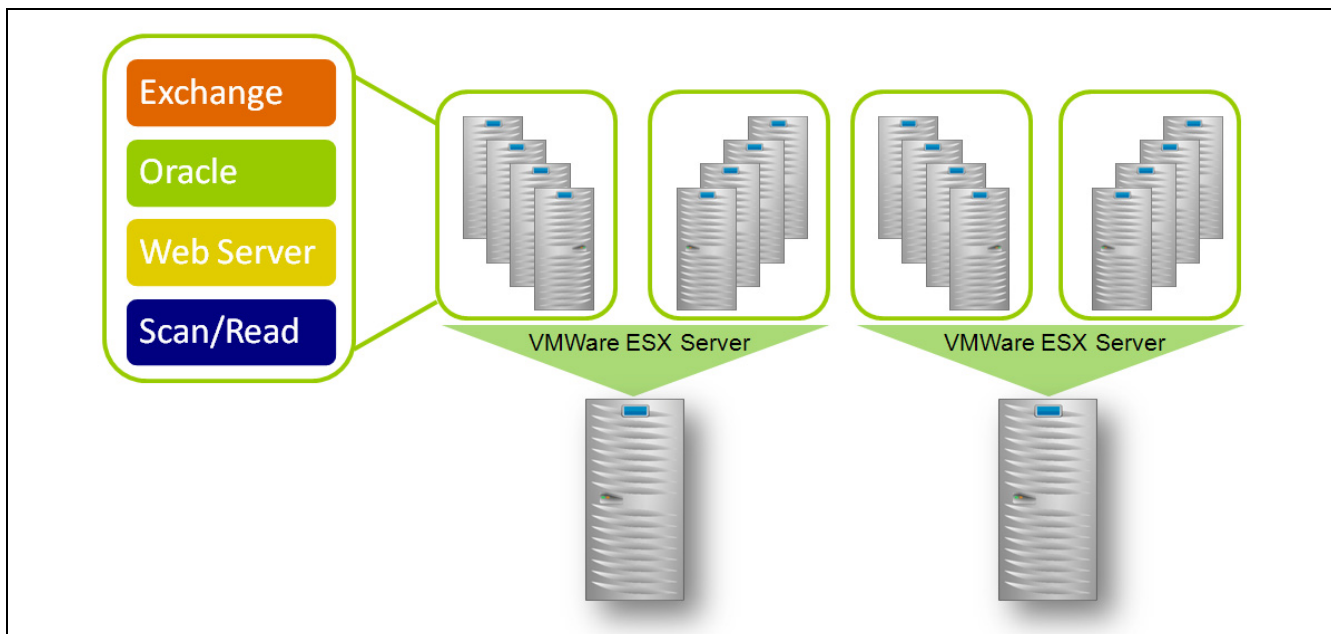
The real-world performance capabilities of the IBM DS5300 were assessed by ESG Lab via hands-on testing at an IBM facility located in Gaithersburg, Maryland. The methodology presented in this report was designed to assess the performance capabilities of a single IBM DS5300 storage system shared by multiple virtual servers running a mix of real-world application workloads. The cooperation of VMware, IBM and LSI 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 in partnership with LSI.

A Mixed Real-world 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 on-line 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. 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. As an example, the high-end IBM x3850 servers used during this ESG Lab Validation have an excellent published VMmark score of 13.5 tiles.

FIGURE 3. A TILE-BASED, STORAGE-FOCUSED 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 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 during this ESG Lab Validation. As shown in Figure 3, each tile is composed of a mixture of four application workloads. Two physical servers, each configured with eight virtual machines, was used to measure performance as the number of active tiles was increased from one to four.

The Application Workloads

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

- **E-Mail:** The Microsoft JetStress utility was used to generate e-mail traffic. Similar to the Microsoft LoadSim 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 LoadSim. Using the underlying Jet Engine database, JetStress was designed to focus on storage performance.²
- **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 on-line transaction processing (OLTP) and bandwidth sensitive on-line analytic processing (OLAP).³
- **Web Server:** The industry standard Iometer utility was used to generate web server traffic. The I/O definition was composed of random reads of various block sizes as documented in the Appendix. 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.
- **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 I/O profile of interactive database and e-mail applications. As most experienced database and storage administrators have learned, a throughput intensive burst in I/O traffic can drag down the performance for interactive applications causing end-users to complain about performance. 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 I/O utilization increased.⁵

Each of the four workloads ran in parallel with the JetStress e-mail test taking the longest to complete (approximately three hours). The Iometer workloads were stopped manually after the JetStress utility had finished.

² JetStress parameters including 32 threads per storage group are included in the JetStress output shown in the Appendix as Figure 11.

³ Orion parameters including a small I/O size of 8 KB and a large I/O size of 1 MB are included in the Orion output shown in Figure 12.


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

⁵ The scan/read workload definition is included in a results file excerpt as Figure 14.

The Physical Test Bed

VMware ESX Server 3.5 software was installed on a pair of powerful IBM xSeries 3850 servers, each with four quad-core 3 GHz processors and 128 GB of RAM. Each server had four dual port 4 Gbps FC host bus adapters connected to a Cisco MDS-9513 FC SAN switch. A DS5300 with 256 15K RPM FC drives was connected to the servers via 16 4 Gbps FC ports as shown in Figure 4.

FIGURE 4. ESG LAB TEST BED



The Drive Layout

The DS5300 drive configuration is summarized in Table 1. Two Microsoft Exchange storage groups and two Oracle databases were tested within each tile. Exchange database volumes were configured over eight drive RAID-10 groups. Simulating a pair of database applications with different performance and cost requirements, one of the Oracle databases was configured using RAID-10 and the second was configured with RAID-5. The web server and scan/read volumes were configured using a 7+1 RAID-5 layout. Volume ownership was balanced across the dual controllers within the DS5300 and distributed evenly over the 16 host interfaces.⁶

TABLE 1. DRIVE CONFIGURATION


Application	Number of LUNs	Number of Drives	Usable Capacity (GB)
Exchange DB	8	64	8,640
Exchange Log	8	40	8,704
Oracle	8	52	4,352
Web Server	4	32	2,176
Scan/Read	4	32	2,176
Vmdk/OS	4	20	2,176
Total	36	240	26,048

⁶ For more detail, please refer to Figures 15 and 16 in the Appendix. The balances of the 256 drives were not configured (8) or defined as hot spares (8).

Configuring Virtual Machines

The configuration of one of the sixteen virtual machines is shown in Figure 5. Each of the sixteen virtual machines was mapped to a quad-core CPU, 16 GB of RAM, a virtual disk over VMFS for the operating system and one or more mapped raw LUNs. DS5300 disk capacity was used for all storage capacity including VMware virtual disk files (VMDK), Windows 2003 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 that was used for JetStress e-mail testing is shown in Figure 5. Note how four mapped raw LUNs were configured: two for the Exchange database volumes and two for the Exchange log volumes.

FIGURE 5. VIRTUAL SERVER CONFIGURATION



Why This Matters

ESG research indicates that the top concern when implementing networked storage platforms to support server virtualization is performance. Of 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 (response time or throughput). Server benchmarks have typically tested only one server running a CPU intensive workload that doesn't stress storage. So that IBM customers can understand how a DS5300 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 zero to sixty are a good first pass indicator of a car's performance. 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.

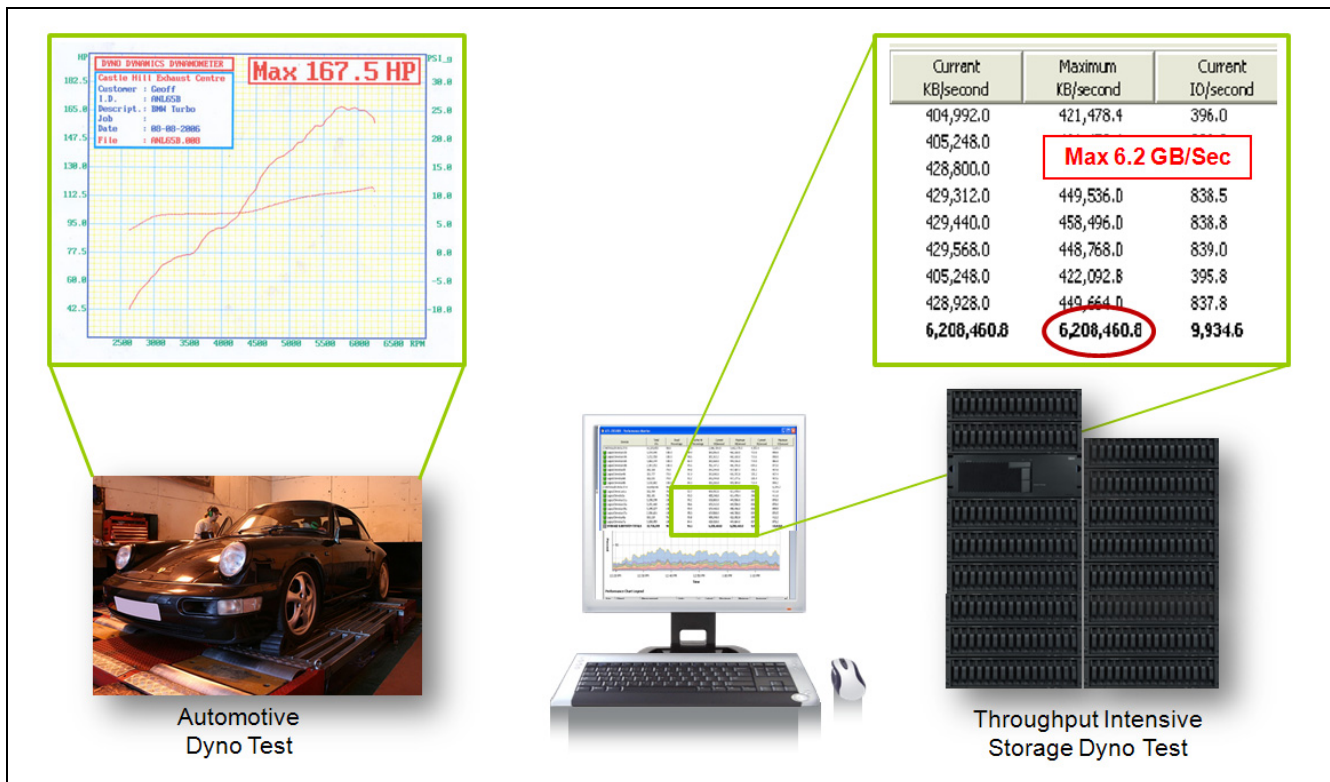
Raw Aggregate Throughput

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 on ten entry-level IBM x335 physical servers running Microsoft Windows operating systems. Half of the drives used later in the mixed, real-world tests were exercised (128 disk drives).

A total of ten servers with twenty 4 GB FC ports were connected through a Cisco MDS 9513 switch to the DS5300 with sixteen active 4 Gbps host ports. A total of 16 LUNs were exercised. Each of the LUNs was configured over a RAID-5 group of 15K RPM drives configured with a 4+1 parity scheme. Each of the Windows servers exercised two LUNs distributed across both DS5300 controllers.

An Iometer profile of 1 MB sequential reads and 1 MB sequential writes were used for this first pass analysis of the raw aggregate throughput capabilities of the DS5300. A similar round of tests performed against the same test bed using an IBM p595 server running the AIX operating system produced similar results. Similar to a dynamometer horsepower rating for a car, the maximum throughput reported by the DS5000 console was used to quantify the power of a turbo-charged DS5300 storage engine.

FIGURE 6. AUDITED TEST RESULTS



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 available bandwidth that the system can take advantage of. Throughput can be measured on a stream or aggregate basis. A stream is represented by one application or user communicating through one I/O 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 confirmed that the DS5300 has an impressive aggregate throughput rating of 6.2 GB/sec for reads and 5.85 MB/sec for writes. These results are approximately four times better than the ratings of the previous generation IBM DS4800.

Why This Matters

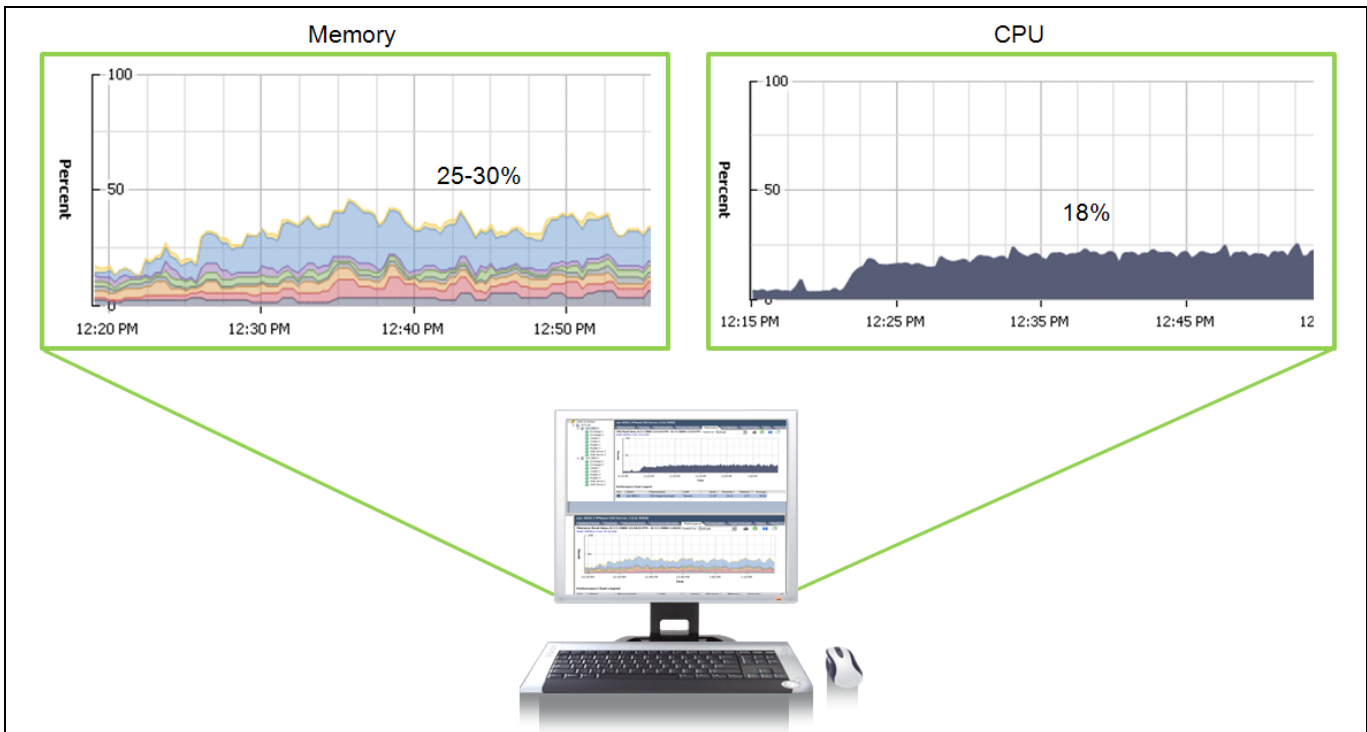
A storage system needs a strong engine and well-designed modular architecture to perform predictably in a mixed real-world environment. One measure of the strength of a storage controller engine is its maximum aggregate throughput. ESG Lab confirmed that a DS5300 system with half the drives used during the mixed workload tests presented later in this report can sustain an excellent 6.2 MB/sec of aggregate large block sequential read throughput.

In ESG Lab's experience, this is an extremely impressive result for a dual controller modular storage system. As a matter of fact, this result indicates that the DS5000 should be well suited for virtual server consolidation and mixed real-world business applications, it is definitely well suited for clustered computing, video editing and scientific applications with extreme bandwidth requirements.

Virtual Machine Utilization

Having finished the low level throughput testing using ten entry-level physical servers, The DS5300 was reconfigured for mixed real-world testing using a pair of high end IBM x3850 servers as documented previously in this report. 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 DS5300. Memory and CPU utilization as reported by the VMware Infrastructure Manager is shown in Figure 7.

FIGURE 7. SYSTEM COMPONENTS (DASHBOARD)



These screenshots were taken during the peak activity phase of the four tile test. With memory utilization under 50 percent and CPU utilization under 25 percent, there was no obvious bottleneck between virtualized applications and the IBM DS500.

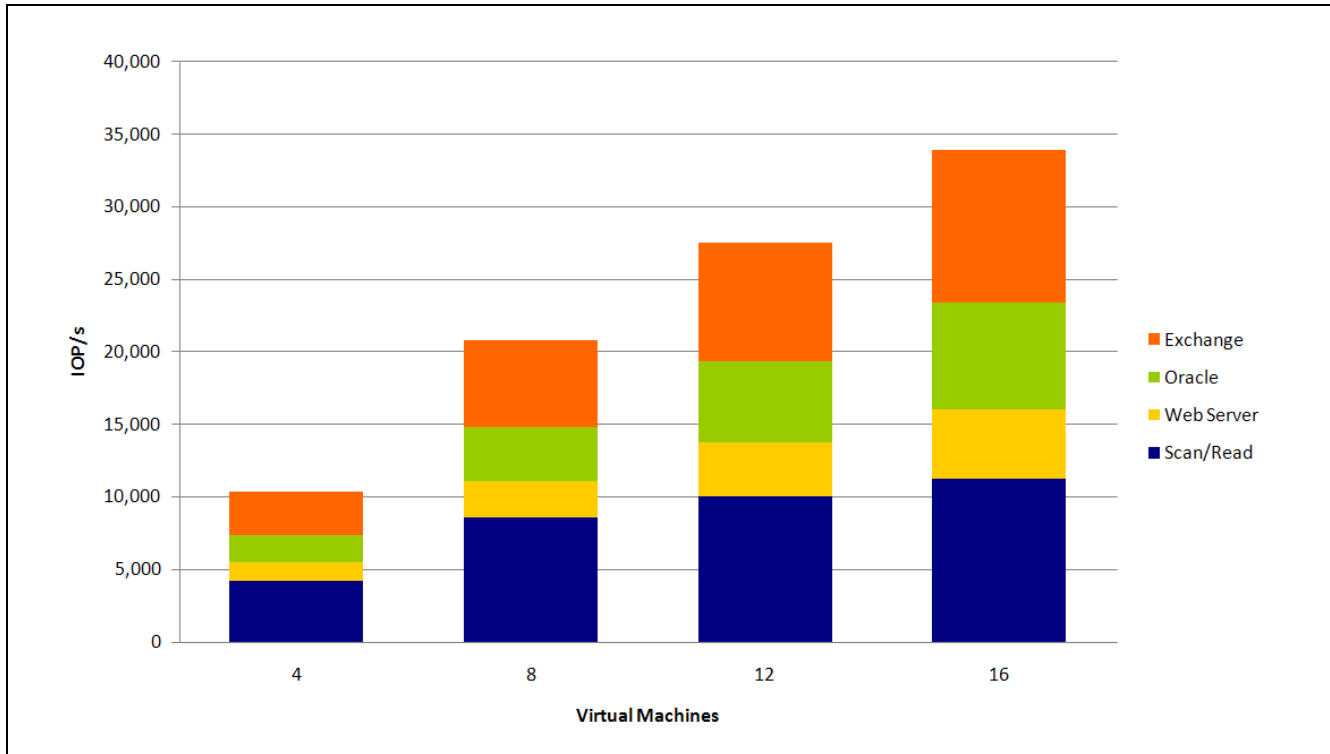
Mixed Real-world IOPS Scalability

I/O's 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 the disk, or out of cache, it will tend to be able to service more applications and users in parallel. Much like the torque 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 I/O traffic patterns that may be serviced from disk or from cache.

ESG Lab measured IOPS performance as reported by the DS5300 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 I/O over hundreds of 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 DS5300 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 8.

FIGURE 8. PEAK ENABLED MIXED WORKLOAD CONSOLIDATION



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 34,438 and an average of 24,714 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

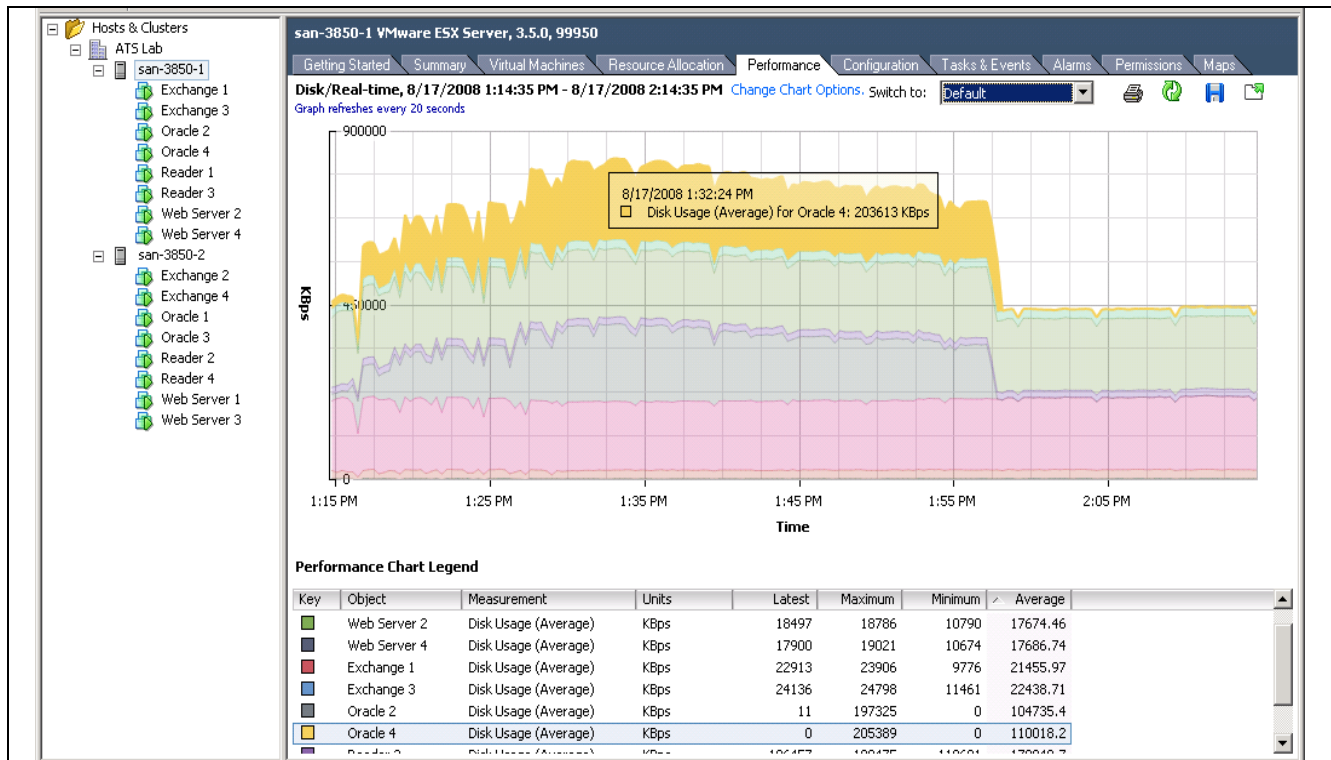
Predictable performance scalability is a critical concern when a mix of applications is sharing a storage system. A burst of I/O activity in one application (e.g., a database consistency check) can lead to poor response times, lost productivity and in the worst case—lost revenue.

ESG Lab confirmed that the rate of I/Os processed by the DS5000 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, there were peaks of throughput activity that 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. The peak recorded shortly after the Orion OLAP phase is shown in Figure 9.

FIGURE 9. PEAK THROUGHPUT (ONE SERVER, FOUR ACTIVE TILES, STACKED VM VIEW)



What the Numbers Mean

- An aggregate throughput level of 1.6 GB/sec was recorded as mixed, real-world applications were run on 16 virtual machines sharing a single DS5300 storage system (800 MB/sec for one of the two physical servers is shown in Figure 9).
- 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. Yet this increases cost and complexity and defeats the goal of shared storage consolidation.

ESG Lab observed a peak aggregate throughput of 1.6 GB/sec as a throughput intensive Oracle Orion OLAP test was running—even as Exchange e-mail traffic ran with predictably good response times.

Mixed Application-level Performance Scalability

Having looked at the IOPS and throughput ratings of the turbo-charged DS5300 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

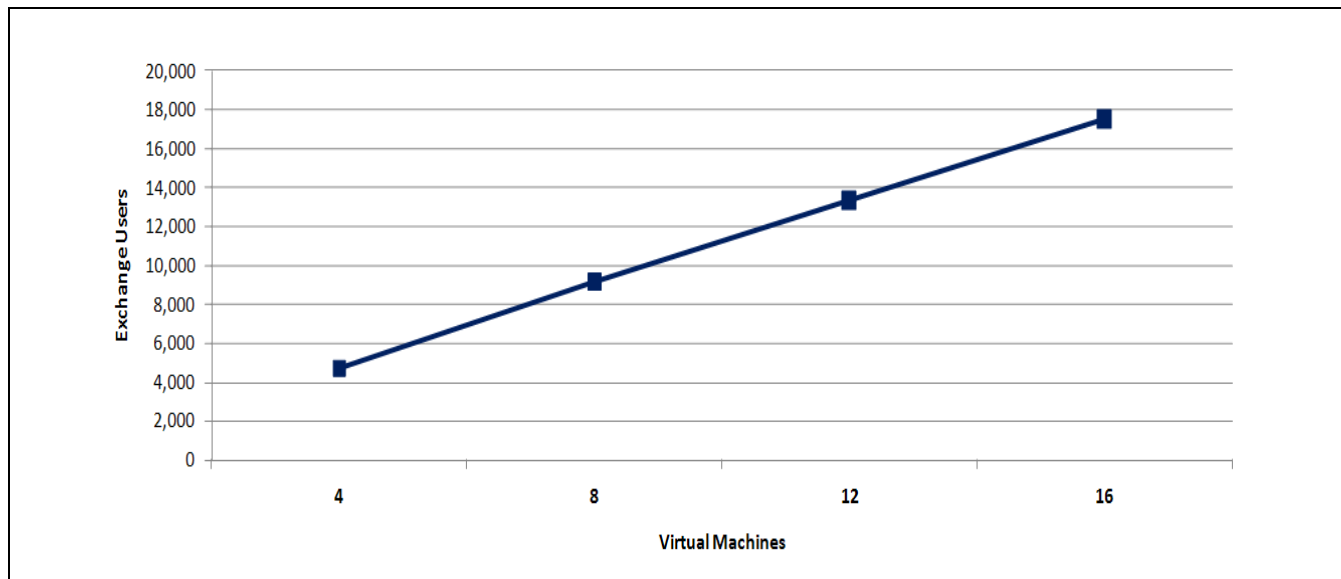
The Microsoft JetStress testing tool was used to see how many simulated e-mail users could be supported by the DS5300 resources allocated for the Exchange application. The number of IOPS and response time for each database and log volume was recorded at the end of each JetStress run. A response time goal of 20 milliseconds or less for DB reads was required to pass the test. This value is 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.⁷

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.⁸ Based on a 0.5 IOPS user profile, the number of IOPS was used to calculate the number of supported Exchange users. Exchange user scalability as the number of tiles was increased from one to four is shown in Figure 10 and Table 2.

FIGURE 10. CONSOLIDATED EXCHANGE RESULTS



⁷ 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>

⁸ A sample JetStress log is included in the Appendix as Figure 11.

TABLE 2. JETSTRESS PERFORMANCE RESULTS (ONE THROUGH FOUR TILES)

Tile	Achieved I/O per Seconds	Users (0.5 Profile)	DB Avg Disk Sec/Read
1	2,356	7,713	.01450
2	4,599	9,198	.01525
3	6,670	13,339	.01550
4	8,756	17,512	.01600

What the Numbers Mean

- The single tile mixed application test supported 7,713 Exchange users with an average DB disk response time of 14.5 milliseconds.
- Performance scaled in a near-linear fashion to 17,512 users while the DS5300 was busy processing and servicing other applications concurrently.
- All tests passed with a DB average read response time under Microsoft's guideline of 20 milliseconds.
- The four tile test which produced 8,756 users over 64 database drives delivered 137 IOPS per drive, well above the conservative IBM guideline of 100 IOPS per drive.

This test used two servers and focused solely on storage performance and sizing. The single IBM DS5300 storage array had significant resources remaining and was under-utilized throughout each test run. At 8,752 users per physical server (17,512 users over two physical servers) this result is close to exceeding Microsoft's recommended 10,000 users per server guideline. Microsoft does not recommend more than 10,000 users per server due to the impact of that many users on recovery time service level agreements. In a production environment, it is recommended to stay within Microsoft's recommendations for support and recovery purposes.

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 applications (e.g. DSS).

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

Target 5-10 millisecond response time for disks performance 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 allow for media repair etc.⁹

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 workloads¹⁰

⁹ Current trends in Database Performance, Andrew Holdsworth, Oracle OpenWorld, Nov 2007, http://www.oracle.com/technology/deploy/performance/pdf/PerfTrends_Holdsworth.pdf

¹⁰ 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>

The results for the four tile Orion test are summarized in Table 3. A sample Orion report is shown in the Appendix as Figure 12.

TABLE 3. ORION FOUR TILE PERFORMANCE RESULTS

Tile	Small IOPS		Large MBPS		Small Latency (ms)	
	RAID-10	RAID-5	RAID-10	RAID-5	RAID-10	RAID-5
1	1,582	683	144	80	4.56	6.16
2	1,569	673	140	75	4.60	5.75
3	1,628	704	143	80	4.57	5.85
4	1,622	703	140	81	4.55	5.53
<i>Total</i>	<u>9,164</u>		<u>883</u>		<u>4.6</u>	<u>5.8</u>

What the Numbers Mean

- Two Oracle database applications were tested within each tile. The first ran on RAID-10 capacity and the second on RAID-5 capacity.
- The four tile test achieved a grand total of 9,164 small IOPS and 883 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 I/O activity is significantly higher than a typical “*medium transaction system*” and nearly represents a “*high-end transaction system*” which is “*almost rarely achievable*”.
- The total number of small IOPS processed during the busy four tile test delivered an excellent 176 small IOPS per drive, dwarfing the extremely conservative Oracle guideline of 30 IOPS per drive.
- The RAID-10 DB volumes with an average small latency of 4.6 milliseconds were slightly faster than the RAID-5 volumes at 5.6 milliseconds. Given the Oracle guidance of 5 to 10 milliseconds, ESG Lab believes that these are excellent results – especially given the mix of I/O intensive workloads that were being serviced by the DS5300 in parallel.

Web Server and Scan/Read

Performance results as reported by 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	1,139	103
2	1,135	108
3	1,140	104
4	1,137	112
<u>Total</u>	<u>4,451</u>	<u>425</u>

What the Numbers Mean

- Given the cache friendly, read-only nature of web server I/O traffic, ESG Lab believes that these results indicate that the DS5300 has the horsepower required to service tens of thousands of simultaneous page requests.
- ESG Lab believes that an 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 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. Raw throughput testing provided an early indication that the DS5000 engine was designed to deliver the right amount of power to virtualized applications when needed. Mixed application testing on virtual servers proved it. The DS5000 is ideally suited to support a number of real-world applications and users in a consolidated virtual data center.

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 DS5000 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.

ESG Lab Validation Highlights

- ☑ An impressive 6.2 GB/sec of aggregate throughput was sustained during test bed staging using physical servers and 128 drives.
- ☑ Mixed real-world application workloads running simultaneously within sixteen virtual machines deployed over two IBM x3850 servers provided the performance needed to concurrently support:
 - 17,512 JetStress e-mail users (0.5 profile).
 - 9,164 Orion small database I/Os per second and 883 large MBPS.
 - 4,551 simulated web server IOPs.
 - 425 MB/sec of streaming read traffic.
- ☑ Excellent IOPs per drive (e.g., 167 for the Oracle OLTP test).
- ☑ As the number of virtual machines sharing a single DS5300 was increased, performance scaled in a near linear fashion with predictably fast response times (16 ms for JetStress DB reads, 4.6 to 5.8 ms for Oracle Orion small IOPS).
- ☑ The DS5300 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 DS5300 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 DS5300 storage.
- ☑ Compared to the previous generation DS4000 Series controllers, the DS5300 doubles the maximum number of supported 4 Gbps FC host interfaces to 16 and increases the maximum number of supported drives from 244 to 256. Given the turbocharged horsepower of the DS5000 series storage engine as shown in this report, future support for more drives and host interfaces (e.g., 8 Gbps FC) would enable IBM customers to consolidate ever increasing numbers of virtual machines onto a single pool of storage. IBM has advised ESG that future releases of the DS5300 will support 8 Gbps FC and up to 448 drives.
- ☑ For applications requiring extreme performance beyond that provided by FC and SATA drives, ESG Lab believes that the DS5000 would be an ideal architecture for the addition of solid state disk drives. In particular, solid state disk drives could be used to improve the performance of highly referenced database indexes and temp files. IBM has advised ESG that future releases of the DS5300 will support solid state disk drive technology.
- ☑ 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 do capacity planning and testing in your own environment to validate a storage system configuration.

ESG Lab's View

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 of these benefits 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, LSI and VMware approached ESG Lab in 2008 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 I/O intensive mix of virtualized business applications. Instead of designing a storage benchmark based on a single application workload with the goal of producing a big number, the benchmark was designed to simulate a mix of application workloads including e-mail and on-line database applications. Instead of designing a test to determine the scalability of a single server, the benchmark was designed to stress a storage system shared by multiple servers. In other words, the benchmark was designed to see what happens when a mix of virtualized business applications are deployed on multiple servers sharing a consolidated storage platform.

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 x3850 servers, with an excellent published VMmark score of 13.5 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 x3850 VMmark test.

IBM and LSI Logic have more than a decade of experience delivering modular FC-attached storage systems that were designed to meet the cost-optimized performance demands of medium-sized organizations, mid-tier applications, remote departments and near-line applications. This ESG Lab report focused on the latest product of the LSI and IBM partnership—the IBM DS5300. 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 DS5100 and DS5300 has been turbo charged to meet the real-world performance demands of virtualized applications. With twice the host connectivity and up to four times the performance of the previous generation DS4800, the DS5000 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 DS5300 test bed can deliver up to 6.2 GB/sec of raw aggregate throughput—with only half of the available FC drives. This outstanding result was an early indicator that the IBM DS5300 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 DS5300 simultaneously supported 17,512 simulated Exchange users *and* 9,164 small database I/Os per second *and* 4,551 simulated web server IOPs *and* 425 MB/sec of throughput for bandwidth intensive streams of read traffic. Predictably fast response times were maintained as the number of virtual servers was increased.

With horsepower to spare for extended storage functions including copy services and remote replication, ESG Lab is pleased to report that the DS5000 delivers the real-world performance needed for a mix of I/O intensive business applications running on virtual machines sharing a common pool of consolidated storage.

Appendix

TABLE 5. TEST BED OVERVIEW

Storage	
IBM DS5300, Firmware: 07.30.18.00	16 drive trays, 256 15K RPM FC drives
Server	
Two x IBM xSeries 3850 servers 128 GB RAM	CPU type: Quad processor, quad core CPU speed: 3 GHz
SAN Components	
FC Switch: Cisco MDS 9513 FC switch, SAN-OS version 3.2	FC Host Bus adapters: Four QLE2462, dual port 4 Gbps HBAs per server
Virtualization Software and Guest Operating Systems	
Server Virtualization	VMWare ESX Server 3.5, update 2
Guest OS	Windows 2003 Enterprise x64 Edition, SP 2

TABLE 6. BENCHMARK UTILITIES/WORKLOAD GENERATORS

E-Mail	Microsoft Jetstress, version 08.02.0060.000 Jetstress parameters: <ul style="list-style-type: none"> • Thread – 32 (per storage group) • 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 Orion parameters: <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 • 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 ¹¹
Scanner/Reader	Iometer, version 2006.07.27 One worker, one outstanding I/O per physical drive 100% Sequential Reads, 256 KB per I/O request

¹¹ See Figure 13 for workload details!

FIGURE 11. E-MAIL RESULTS

This is an example of the output created by the JetStress utility. This example 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.

Overall Test Result Pass

Machine Name JE-04
Test Description
Test Start Time 8/15/2008 10:09:51 AM
Test End Time 8/15/2008 12:14:08 PM
Jetstress Version 08.02.0060.000
Ese Version 08.01.0240.005
Operating System Microsoft Windows Server 2003 Service Pack 2 (5.2.3790.131072)
Performance Log [C:\Program Files\Exchange Jetstress\Performance_2008_8_15_10_9_58.blg](#)
[C:\Program Files\Exchange Jetstress\DBChecksum_2008_8_15_12_14_8.blg](#)

Database Sizing and Throughput

Achieved I/O per Second 2231.58
Capacity Percentage 100%
Throughput Percentage 100%
Initial database size 1951889391616
Final database size 1958430408704
Database files (count) 2

Jetstress System Parameters

Thread count 32 (per-storage group)
Log buffers 9000
Minimum database cache 64.0 MB
Maximum database cache 512.0 MB
Insert operations 40%
Delete operations 30%
Replace operations 5%
Read operations 25%
Lazy commits 55%

Disk Subsystem Performance

Logical Disk	Avg. Disk sec/Read	Avg. Disk sec/Write	Disk Reads/sec	Disk Writes/sec	Avg. Disk Bytes/Write
Database (E:)	0.017	0.003	588.175	533.223	(n/a)
Database (F:)	0.016	0.004	585.058	525.125	(n/a)
Log (G:)	0.001	0.002	6.881	201.775	6286.999
Log (H:)	0.001	0.002	6.877	201.086	6281.840

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	11.071	3.021	16.719
Available MBytes	2616.813	2615.000	2661.000
Free System Page Table Entries	3918422.000	3918422.000	3918422.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	25990365.867	25223168.000	26284032.000
Pool Paged Bytes	97061563.733	96985088.000	97251328.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

8/15/2008 10:09:51 AM -- Jetstress testing begins ...
 8/15/2008 10:09:51 AM -- Prepare testing begins ...
 8/15/2008 10:09:55 AM -- Attaching databases ...
 8/15/2008 10:09:55 AM -- Dispatching transactions begins ...
 8/15/2008 10:09:55 AM -- Database cache settings: (minimum: 64.0 MB, maximum: 512.0 MB)
 8/15/2008 10:09:55 AM -- Database flush thresholds: (start: 5.1 MB, stop: 10.2 MB)
 8/15/2008 10:09:58 AM -- Database read latency thresholds: (avg: 0.02 sec/read, max: 0.05 sec/read).
 8/15/2008 10:09:58 AM -- Log write latency thresholds: (avg: 0.01 sec/write, max: 0.05 sec/write).
 8/15/2008 10:09:59 AM -- Operation mix: Sessions 32, Inserts 40%, Deletes 30%, Replaces 5%, Reads 25%, Lazy Commits 55%.
 8/15/2008 10:09:59 AM -- Performance logging begins (interval: 15000 ms).
 8/15/2008 10:09:59 AM -- Attaining prerequisites:
 8/15/2008 10:14:03 AM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 483606500.0 (lower bound: 483183800.0, upper bound: none)
 8/15/2008 12:14:05 PM -- Performance logging ends.
 8/15/2008 12:14:05 PM -- JetInterop batch transaction stats: 100849, and 100734.
 8/15/2008 12:14:05 PM -- Dispatching transactions ends.
 8/15/2008 12:14:05 PM -- Shutting down databases ...
 8/15/2008 12:14:08 PM -- Instance2560.1 (complete), and Instance2560.2 (complete)
 8/15/2008 12:14:08 PM -- Performance logging begins (interval: 30000 ms).
 8/15/2008 12:14:08 PM -- Verifying database checksums ...
 8/15/2008 1:46:06 PM -- E: (100% processed), and F: (100% processed)
 8/15/2008 1:46:09 PM -- Verifying log checksums ...
 8/15/2008 1:46:15 PM -- Volume E: has 0.0171 for Avg. Disk sec/Read.
 8/15/2008 1:46:15 PM -- Volume F: has 0.0159 for Avg. Disk sec/Read.
 8/15/2008 1:46:15 PM -- Volume G: has 0.0018 for Avg. Disk sec/Write.
 8/15/2008 1:46:15 PM -- Volume G: has 0.0010 for Avg. Disk sec/Read.
 8/15/2008 1:46:15 PM -- Volume H: has 0.0018 for Avg. Disk sec/Write.
 8/15/2008 1:46:15 PM -- Volume H: has 0.0010 for Avg. Disk sec/Read.
 8/15/2008 1:46:15 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
 8/15/2008 1:46:15 PM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.

FIGURE 12. DATABASE RESULTS

This is an example of the output created by the Oracle Orion utility. This example 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 VMWareTwo -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: VMWareTwo

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: 38

Name: \\.\e: Size: 981449728

1 FILEs found.

Maximum Large MBPS=140.21 @ Small=0 and Large=16

Maximum Small IOPS=1622 @ Small=38 and Large=0

Minimum Small Latency=4.55 @ Small=1 and Large=0

FIGURE 13. WEB SERVER RESULTS

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

'Test Type Test Description
 0 ESG Lab Characterization

'Version

2006.07.27

'Time Stamp

2008-08-15 10:15:49:094

'Access specifications

'Access specific default assignment

Web Server

'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

'End access specifications

'Results

'Target Type	Target Name	Name	# Managers	# Workers	# Disks	IOps	Read IOps	Write IOps	MBps
ALL	All	Web Server	1	4	4	1137.477518	1137.477518	0	17.39527
MANAGER	IO-04	Web Server		4	4	1137.477518	1137.477518	0	17.39527
PROCESSOR	CPU 0								
PROCESSOR	CPU 1								
WORKER	Worker 1	Web Server			1	284.36155	284.36155	0	4.350689
DISK	PHYSICALDRIVE:2					284.36155	284.36155	0	4.350689
WORKER	Worker 2	Web Server			1	284.44322	284.44322	0	4.33213
DISK	PHYSICALDRIVE:2					284.44322	284.44322	0	4.33213
WORKER	Worker 3	Web Server			1	284.273164	284.273164	0	4.359669
DISK	PHYSICALDRIVE:2					284.273164	284.273164	0	4.359669
WORKER	Worker 4	Web Server			1	284.399583	284.399583	0	4.352781
DISK	PHYSICALDRIVE:2					284.399583	284.399583	0	4.352781

'Time Stamp

2008-08-15 14:00:34:141

FIGURE 14. SCAN/READ RESULTS

This is an example of the output created by the lometer utility after a scan/read test run. This example shows the performance for one of four scan/read tests running in parallel. Specifically, this report was created by the lometer scan/read job running on a virtual machine within the fourth tile of the four tile test.

'Test Type Test Description
 0 ESG Lab Characterization

'Version

2006.07.27

'Time Stamp

2008-08-15 10:14:16:535

'Access specifications

'Access specificati default assignment

Backup reader

'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	Specificatio	# Managers	# Workers	# Disks	IOps	Read IOps	Write IOps	MBps	Read MBps
ALL	All	ackup read	1	1	1	450.9572	450.957212	0	112.7393	112.739303
MANAGER	AD-04	ackup reader		1	1	450.9572	450.957212	0	112.7393	112.739303

PROCESSOR CPU 0

PROCESSOR CPU 1

WORKER	Worker 1	Backup reader			1	450.9572	450.957212	0	112.7393	112.739303
DISK	PHYSICALDRIVE:2					450.9572	450.957212	0	112.7393	112.739303

'Time Stamp

2008-08-15 13:56:35:332

FIGURE 15. DS5300 DRIVE MAP

Volume Name	App	# Drives	RAID Level	Usable Capacity(GB)	Segment Size (KB)	Drive Pieces
Orion-01	Oracle DB	8	RAID-10	544	512	(10,1 10,2 10,3 10,4 11,1 11,2 11,3 11,4)
OS-01	vmfs - O/S	5	RAID-5	544	128	(20,9 20,10 21,9 21,10 21,11)
OS-02	Exchange Log	5	RAID-5	544	128	(20,14 20,15 20,16 21,14 21,15)
OS-03	Exchange Log	5	RAID-5	544	128	(30,9 30,10 31,9 31,10 31,11)
IO-01	Web Server	8	RAID-10	544	128	(40,1 40,2 40,3 40,4 41,1 41,2 41,3 41,4)
AD-01	Reader	8	RAID-10	544	128	(50,1 50,2 50,3 50,4 51,1 51,2 51,3 51,4)
Log-01	Unconfigured	1	RAID-0	136	128	(50,15)
OS-04	Oracle DB	5	RAID-5	544	512	(60,9 60,10 61,9 61,10 61,11)
Log-02	Unconfigured	1	RAID-0	136	128	(60,14)
Jet-01	Exchange DB	8	RAID 10	1080	128	(70,1 70,2 70,3 70,4 71,1 71,2 71,3 71,4)
Jet-02	Exchange DB	8	RAID-10	1080	128	(80,1 80,2 80,3 80,4 81,1 81,2 81,3 81,4)
OS-05	vmfs - O/S	5	RAID -5	544	128	(10,11 10,12 10,13 11,12 11,13)
OS-06	Exchange Log	5	RAID-5	544	128	(10,14 10,15 10,16 11,14 11,15)
Orion-02	Oracle DB	8	RAID-10	544	512	(20,1 20,2 20,3 20,4 21,1 21,2 21,3 21,4)
IO-02	Web Server	8	RAID-10	544	128	(30,1 30,2 30,3 30,4 31,1 31,2 31,3 31,4)
OS-07	Exchange Log	5	RAID-5	544	128	(40,9 40,10 41,9 41,10 41,11)
OS-08	Oracle DB	5	RAID-5	544	512	(50,9 50,10 51,9 51,10 51,11)
Log-03	N/A	1	RAID-0	136	128	(50,14)
AD-02	Reader	8	RAID-10	544	128	(60,1 60,2 60,3 60,4 61,1 61,2 61,3 61,4)
Log-04	N/A	1	RAID-0	136	128	(60,15)
Jet-03	Exchange DB	8	RAID-10	1080	128	(70,5 70,6 70,7 70,8 71,5 71,6 71,7 71,8)
Jet-04	Exchange DB	8	RAID-10	1808	128	(80,5 80,6 80,7 80,8 81,5 81,6 81,7 81,8)
Orion-03	Oracle DB	8	RAID-10	544	512	(10,5 10,6 10,7 10,8 11,5 11,6 11,7 11,8)
OS-09	vmfs - O/S	5	RAID-5	544	128	(20,11 20,12 20,13 21,12 21,13)
OS-10	Exchange Log	5	RAID-5	544	128	(30,11 30,12 30,13 31,12 31,13)
OS-11	Exchange Log	5	RAID-5	544	128	(30,14 30,15 30,16 31,14 31,15)
IO-03	Web Server	8	RAID-10	544	128	(40,5 40,6 40,7 40,8 41,5 41,6 41,7 41,8)
AD-03	Reader	8	RAID-10	544	128	(50,5 50,6 50,7 50,8 51,5 51,6 51,7 51,8)
Log-05	N/A	1	RAID-0	136	128	(51,15)
OS-12	Oracle DB	5	RAID-5	544	512	(60,11 60,12 60,13 61,12 61,13)
Log-06	N/A	1	RAID-0	136	128	(61,14)
Jet-05	Exchange DB	8	RAID-10	1080	128	(70,9 70,10 70,11 70,12 71,9 71,10 71,11 71,12)
Jet-06	Exchange DB	8	RAID-10	1080	128	(80,9 80,10 80,11 80,12 81,9 81,10 81,11 81,12)
OS-13	vmfs - O/S	5	RAID-5	544	128	(10,9 10,10 11,9 11,10 11,11)

ESG LAB VALIDATION
IBM DS5300 Mixed Performance Analysis

Orion-04	Oracle DB	8	RAID-10	544	512	(20,5 20,6 20,7 20,8 21,5 21,6 21,7 21,8)
IO-04	Web Server	8	RAID-10	544	128	(30,5 30,6 30,7 30,8 31,5 31,6 31,7 31,8)
OS-14	Exchange Log	5	RAID-5	544	128	(40,11 40,12 40,13 41,12 41,13)
OS-15	Exchange Log	5	RAID-5	544	128	(40,14 40,15 40,16 41,14 41,15)
OS-16	Oracle DB	5	RAID-5	544	512	(50,11 50,12 50,13 51,12 51,13)
Log-07	N/A	1	RAID-0	136	128	(51,14)
AD-04	Reader	8	RAID-10	544	128	(60,5 60,6 60,7 60,8 61,5 61,6 61,7 61,8)
Log-08	N/A	1	RAID-0	136	128	(61,15)
Jet-07	Exchange DB	8	RAID-10	1080	128	(70,13 70,14 70,15 70,16 71,13 71,14 71,15 71,16)
Jet-08	Exchange DB	8	RAID-10	1080	128	(80,13 80,14 80,15 80,16 81,13 81,14 81,15 81,16)

FIGURE 16. DS5300 CONFIGURATION DETAILS

The following excerpts were extracted from an IBM DS5300 Storage System Profile Summary taken after the last test run had completed.

```

-----
PROFILE FOR STORAGE SUBSYSTEM: ATSLAB_DS5300

SUMMARY

Number of controllers: 2
High performance tier controllers: Enabled
Number of arrays: 44
RAID 6: Enabled
Total number of logical drives used:      45
  Number of standard logical drives:      44
  Number of access logical drives:        1
Total number of logical drives allowed: 2048

Number of drives:          256
Mixed drive types:         Enabled
Current drive type(s):     Fibre (256)
Total hot spare drives:    8
  Standby:                  8
Number of drive enclosures: 16

Storage Partitioning:      Enabled
  Number of partitions used: 4 out of 256
Number of logical drives allowed per partition: 256
Default host OS:           LNXCLVMWARE (Host OS index 13)

Current configuration
  Firmware version:         07.30.18.00
  NVSRAM version:          NIBMXBB2R1030V03
  EMW version:              10.30.G5.04
  AMW version:              10.30.G5.04
NVSRAM configured for batteries: Yes
Start cache flushing at (in percentage): 80
Stop cache flushing at (in percentage): 80
Cache block size (in KB): 8
...
CONTROLLERS-----
Number of controllers: 2
  Controller in Enclosure 85, Slot A
  Current configuration
    Firmware version:      07.30.18.00
    Appware version:       07.30.18.00
    Bootware version:      07.30.18.00
    NVSRAM version:       NIBMXBB2R1030V03
  Data Cache
    Total present:         8192 MB
    Total used:            8192 MB

  Cache Backup Device
    Status:                Optimal
    Type:                  USB flash drive
...
ARRAYS-----
Number of arrays: 44
  Capacity                 544.922 GB
  RAID level:              10
  Drive type:              Fibre Channel
...
STANDARD LOGICAL DRIVES-----
Number of standard logical drives: 44
  NAME      STATUS  CAPACITY  RAID LEVEL  ARRAY  DRIVE TYPE
  
```

ESG LAB VALIDATION
IBM DS5300 Mixed Performance Analysis

AD-01	Optimal	544.922 GB	10	VG06	Fibre
AD-02	Optimal	544.922 GB	10	VG19	Fibre
AD-03	Optimal	544.922 GB	10	VG28	Fibre
AD-04	Optimal	544.922 GB	10	VG41	Fibre
IO-01	Optimal	544.922 GB	10	VG05	Fibre
IO-02	Optimal	544.922 GB	10	VG15	Fibre
IO-03	Optimal	544.922 GB	10	VG27	Fibre
IO-04	Optimal	544.922 GB	10	VG36	Fibre
Jet-01	Optimal	1.089 TB	10	VG10	Fibre
Jet-02	Optimal	1.089 TB	10	VG11	Fibre
Jet-03	Optimal	1.089 TB	10	VG21	Fibre
Jet-04	Optimal	1.089 TB	10	VG22	Fibre
Jet-05	Optimal	1.089 TB	10	VG32	Fibre
Jet-06	Optimal	1.089 TB	10	VG33	Fibre
Jet-07	Optimal	1.089 TB	10	VG43	Fibre
Jet-08	Optimal	1.089 TB	10	VG44	Fibre
Log-01	Optimal	136.23 GB	0	VG07	Fibre
Log-02	Optimal	136.23 GB	0	VG09	Fibre
Log-03	Optimal	136.23 GB	0	VG18	Fibre
Log-04	Optimal	136.23 GB	0	VG20	Fibre
Log-05	Optimal	136.23 GB	0	VG29	Fibre
Log-06	Optimal	136.23 GB	0	VG31	Fibre
Log-07	Optimal	136.23 GB	0	VG40	Fibre
Log-08	Optimal	136.23 GB	0	VG42	Fibre
Orion_1	Optimal	544.922 GB	10	VG01	Fibre
Orion_2	Optimal	544.922 GB	10	VG14	Fibre
Orion_3	Optimal	544.922 GB	10	VG23	Fibre
Orion_4	Optimal	544.922 GB	10	VG35	Fibre
OS-01	Optimal	544.922 GB	5	VG02	Fibre
OS-02	Optimal	544.922 GB	5	VG03	Fibre
OS-03	Optimal	544.922 GB	5	VG04	Fibre
OS-05	Optimal	544.922 GB	5	VG12	Fibre
OS-06	Optimal	544.922 GB	5	VG13	Fibre
OS-07	Optimal	544.922 GB	5	VG16	Fibre
OS-09	Optimal	544.922 GB	5	VG24	Fibre
OS-10	Optimal	544.922 GB	5	VG25	Fibre
OS-11	Optimal	544.922 GB	5	VG26	Fibre
OS-13	Optimal	544.922 GB	5	VG34	Fibre
OS-14	Optimal	544.922 GB	5	VG37	Fibre
OS-15	Optimal	544.922 GB	5	VG38	Fibre
OS_08	Optimal	544.922 GB	5	VG17	Fibre
OS_12	Optimal	544.922 GB	5	VG30	Fibre
OS_16	Optimal	544.922 GB	5	VG39	Fibre
OS_4	Optimal	544.922 GB	5	VG08	Fibre

...
DIVE DETAILS
Drive at Enclosure 10, Slot 1

```

Status:                Optimal

Mode:                  Assigned
Raw capacity:          136.732 GB
Usable capacity:       136.232 GB
Speed:                 15,000 RPM
Current data rate:     4 Gbps
Product ID:            MAX3147FD      F
Firmware version:      S708
Serial number:         DV59P7300PY0
Vendor:                IBM-SSG
Date of manufacture:   April 18, 2007

```

...
MAPPINGS (Storage Partitioning - Enabled (4 of 256 used)

Logical Drive Name	LUN	Controller	Accessible by	Logical Drive status
Access Logical Drive	31	A,B	Host Group ESX_A_GROUP	Optimal
AD-01	5	A	Host Group ESX_A_GROUP	Optimal
IO-01	4	B	Host Group ESX_A_GROUP	Optimal
Jet-01	9	A	Host Group ESX_A_GROUP	Optimal
Jet-02	10	B	Host Group ESX_A_GROUP	Optimal
Log-01	6	A	Host Group ESX_A_GROUP	Optimal
Log-02	8	B	Host Group ESX_A_GROUP	Optimal

OS-01	1	B	Host Group	ESX_A_GROUP	Optimal
OS-02	2	B	Host Group	ESX_A_GROUP	Optimal
OS-03	3	A	Host Group	ESX_A_GROUP	Optimal
OS-09	11	B	Host Group	ESX_A_GROUP	Optimal
OS_4	7	B	Host Group	ESX_A_GROUP	Optimal
Orion_1	0	A	Host Group	ESX_A_GROUP	Optimal
Access Logical Drive	31	A,B	Host Group	ESX_B_GROUP	Optimal
AD-02	7	B	Host Group	ESX_B_GROUP	Optimal
IO-02	3	A	Host Group	ESX_B_GROUP	Optimal
Jet-03	9	A	Host Group	ESX_B_GROUP	Optimal
Jet-04	10	B	Host Group	ESX_B_GROUP	Optimal
Log-03	6	A	Host Group	ESX_B_GROUP	Optimal
Log-04	8	B	Host Group	ESX_B_GROUP	Optimal
OS-05	0	A	Host Group	ESX_B_GROUP	Optimal
OS-06	1	A	Host Group	ESX_B_GROUP	Optimal
OS-07	4	B	Host Group	ESX_B_GROUP	Optimal
OS_08	5	A	Host Group	ESX_B_GROUP	Optimal
Orion_2	2	B	Host Group	ESX_B_GROUP	Optimal
Access Logical Drive	31	A,B	Host Group	ESX_C_GROUP	Optimal
AD-03	5	A	Host Group	ESX_C_GROUP	Optimal
IO-03	4	B	Host Group	ESX_C_GROUP	Optimal
Jet-05	9	A	Host Group	ESX_C_GROUP	Optimal
Jet-06	10	B	Host Group	ESX_C_GROUP	Optimal
Log-05	6	A	Host Group	ESX_C_GROUP	Optimal
Log-06	8	B	Host Group	ESX_C_GROUP	Optimal
OS-10	2	A	Host Group	ESX_C_GROUP	Optimal
OS-11	3	A	Host Group	ESX_C_GROUP	Optimal
OS_12	7	B	Host Group	ESX_C_GROUP	Optimal
Orion_3	0	B	Host Group	ESX_C_GROUP	Optimal
Access Logical Drive	31	A,B	Host Group	ESX_D_GROUP	Optimal
AD-04	8	B	Host Group	ESX_D_GROUP	Optimal
IO-04	2	A	Host Group	ESX_D_GROUP	Optimal
Jet-07	10	A	Host Group	ESX_D_GROUP	Optimal
Jet-08	11	B	Host Group	ESX_D_GROUP	Optimal
Log-07	7	A	Host Group	ESX_D_GROUP	Optimal
Log-08	9	B	Host Group	ESX_D_GROUP	Optimal
OS-13	0	A	Host Group	ESX_D_GROUP	Optimal
OS-14	4	B	Host Group	ESX_D_GROUP	Optimal
OS-15	5	B	Host Group	ESX_D_GROUP	Optimal
OS_16	6	A	Host Group	ESX_D_GROUP	Optimal
Orion_4	1	A	Host Group	ESX_D_GROUP	Optimal
Access Logical Drive	31	A,B	Storage Subsystem		Optimal



20 Asylum Street
 Milford, MA 01757
 Tel: 508-482-0188
 Fax: 508-482-0218

www.enterprisestrategygroup.com