Lab Validation
Report

The IBM DS8000

Enterprise Class Performance and Functionality

By Brian Garrett
With Tony Palmer

February, 2008

Table of Contents

ESG Lab Reports

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

Introduction

This Enterprise Strategy Group (ESG) Lab Validation report examines the enhanced performance and virtualization capabilities of the IBM Systems Storage DS8000 Series. The report presents the results of hands-on testing and analysis with a focus on innovative performance and virtualization capabilities including space efficient flash copies (FlashCopy SE), Storage Pool Striping, Dynamic Volume Expansion and Adaptive Multi-stream pre-fetch (AMP),

Background

ESG Research indicates that large organizations have a number of IT priorities in common.[1] As seen in Figure 1, compliance, risk reduction, security concerns and disaster avoidance are driving a wave of storage and server consolidation. In addition to these high level corporate concerns, consolidation is also being used to reduce the costs of equipment, manpower and infrastructure including power and cooling. And of course, this all this has to be done while improving application performance for end users.

FIGURE 1. ENTERPRISE IT PRIORITIES

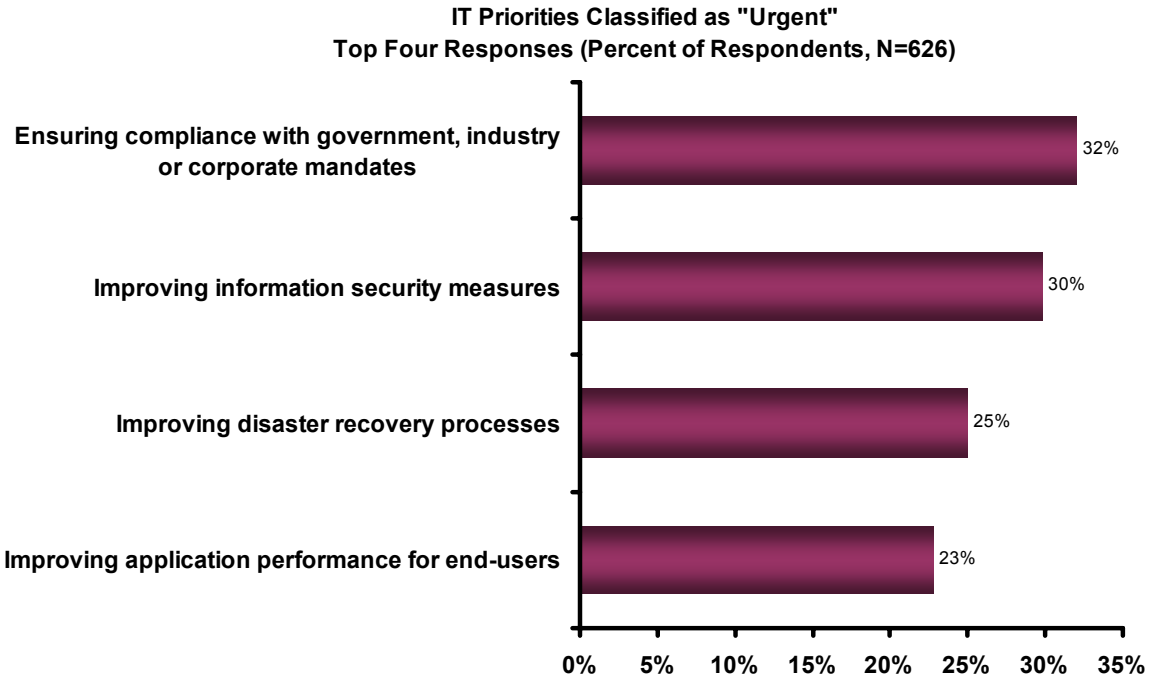

With IT priorities and budget constraints driving IT managers toward virtualization and consolidation, multiple challenges present themselves when it comes to providing storage services. ESG Research also looked at the challenges associated with information assets that are growing at the fastest rate within most organizations—disk-based data protection and digital archives. Once again, compliance, risk mitigation and performance were identified as key concerns, but the challenges that stand out the most are data growth and complexity. Never-ending storage capacity growth and the complexity that goes with that growth is clearly a problem within organizations of all sizes.

ESG Research estimates that organizations will retain over 89,000 petabytes of unstructured data in digital archives over the next five years[2]. A petabyte is a thousand terabytes. Until recently it was rare to

[1] ESG Research Report: *Branch Office Optimization,* January, 2007.
[2] ESG Research Report: *Digital Archiving: End-User Survey & Market Forecast 2006-2010,* January, 2006.

hear of a single organization storing more than a petabyte of data. Now, many of the CIOs within the world's largest organizations are dealing with a petabyte or more of digital capacity. Unstructured, or file, capacity is growing at a much faster rate compared to e-mail and database. Unstructured files in the form of documents, images, video and audio are driving the need for greater capacity AND performance.

The IBM System Storage DS8000 is a high-performance, highly scalable, enterprise-class storage solution. The DS8000 is designed to support continuous operation for mission critical applications. With it's high capacity, high performance, flexible scalability, broad server support and virtualization features, the DS8000 is targeted at clients who aim to simplify the storage environment by consolidating data from multiple storage systems into one. The DS8000, shown in Figure 2, is a modular system that scales up to 512 TB of capacity in four racks.

FIGURE 2. THE IBM DS8000: HIGH-PERFORMANCE, HIGH-CAPACITY ENTERPRISE-CLASS STORAGE

In this report, ESG Lab focused on the performance capabilities of the IBM DS8000 and a number of valuable capabilities introduced recently in Release 3 including adaptive multi-stream pre-fetch, space efficient flash copies, storage pool striping and dynamic volume expansion.

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing at IBM's performance lab in Tucson, Arizona utilizing an IBM DS8300 disk array and an IBM p595 server, as shown in Figure 3.

FIGURE 3. THE ESG LAB TEST BED

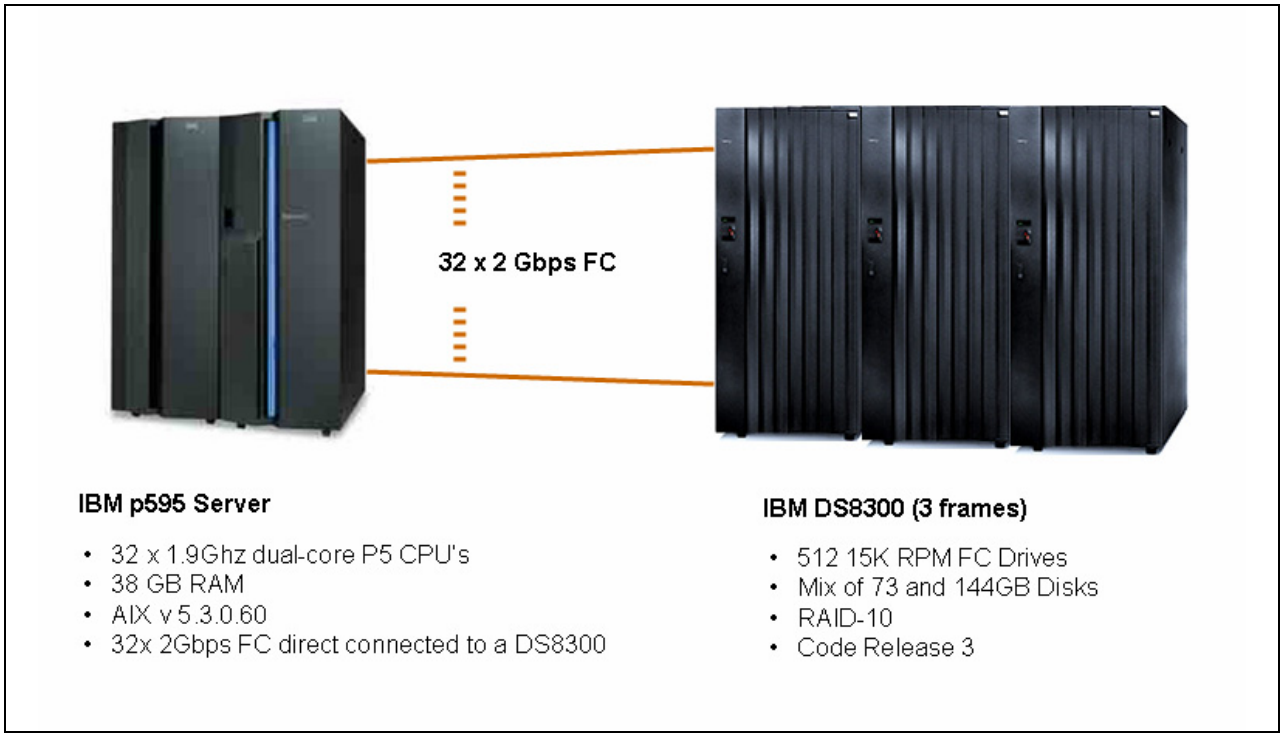

FlashCopy SE

The FlashCopy utility has been embraced by IBM customers for years. In general, a flash copy is a point-in-time copy of an application volume that is created nearly instantaneously. One of the uses of a flash copy is to decrease the time required to obtain a backup of a production database. Without a flash copy, the application must be shut down to obtain a consistent usable backup of the database application. For a very large database, it can take hours, or possibly days, for the backup to complete. The database typically remains unusable during the entire duration of the backup. In contrast, FlashCopy shrinks the time that the database is unavailable from hours to minutes. This is done by making a virtual, consistent copy of the volume that is frozen at a point in time. A space efficient variant of this technology, called FlashCopy SE, is supported in the latest release of the IBM DS8000.

As shown in Figure 4, a traditional Flash copy requires that you allocate space for a full copy whether you need it or not. In this example, 200 GB of DS8000 capacity has been reserved to create a point-in-time image of a DB2 volume. With Flash Copy SE, the IBM DS8000 provides a virtual copy of the volume without having to pre-allocate capacity. Besides a bit of space needed to maintain pointers to un-changed data, the only space consumed is a record of writes completed since the flash copy command was executed. In this example, the 20 GB of data that was written by the DB2 application since the flash copy was executed is all that is needed to maintain a point-in-time image of a 200 GB DB2 volume.

FIGURE 4. TRADITIONAL VS. SPACE EFFICIENT FLASHCOPY

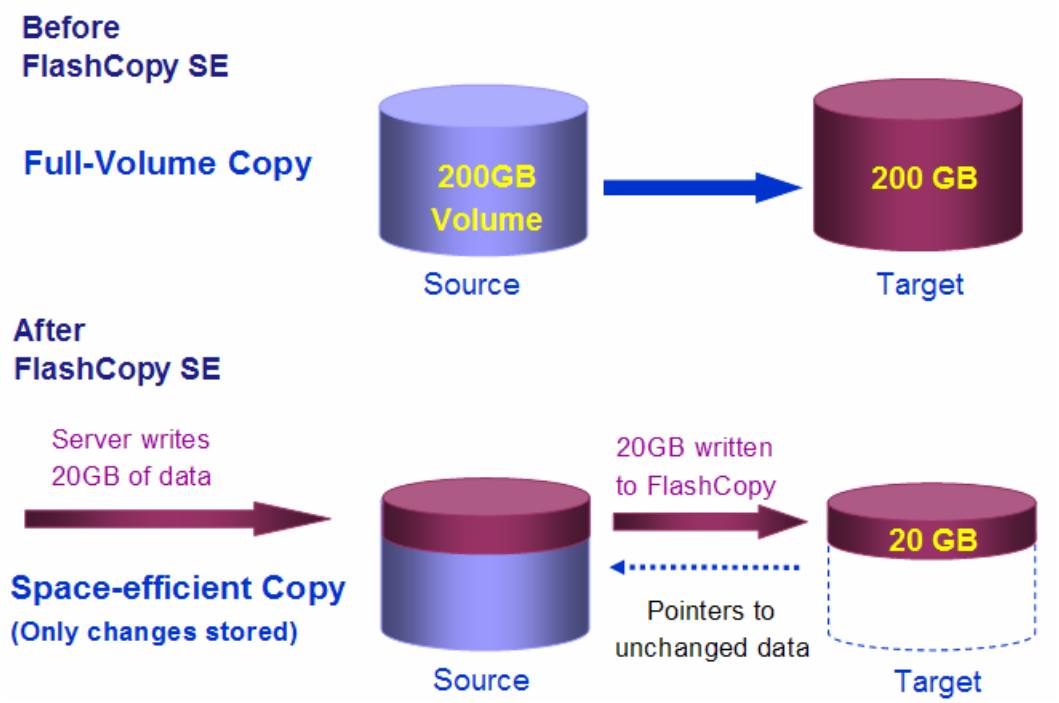

ESG Lab Testing

ESG Lab created a 10 gigabyte file system and copied 5 gigabytes of files into it including a recently created text file. A point-in-time flash copy of the file system was created and mounted via a different host adapter on the server. The creation of the copy was instantaneous. The text file was modified and an additional 132 megabyte file was copied to the live volume. At this point, the flash copy was examined and the text file was confirmed to be a frozen image of the file system at the time of the flash copy. ESG Lab also verified the space consumed by FlashCopy SE using a DSCLI command. It was observed that FlashCopy SE consumed only the capacity written since the copy was made (132 megabytes).

Why This Matters

The most common usage of point-in-time copy technology is to make backups more efficient by allowing applications to remain online as backups are run. Point-in-time images of critical application volumes can be used for a variety of other purposes as well. Test and development, remote replication and the maintenance of disk-based images of critical applications for quick and reliable recovery can also be beneficial uses of point-in-time images. When used as directed for application volumes that aren't write intensive, space efficient flash copies can significantly reduce the cost of capacity allocated for point-in-time images. More cost-effective flash copies can be used to create more frequent point-in-time images, resulting in better recovery point objectives as well as improved satisfaction and service level agreements.

Dynamic Volume Expansion

Dynamic volume expansion is used to increase the size of a storage volume while applications which rely on that volume remain on-line and active. Some time after the volume has been expanded, a short period of application downtime is needed so that operating systems and file systems can recognize and use the new capacity. For example, on a server running a Windows operating system, a disk administrator rescan operation or reboot is required to recognize expanded capacity. This capability has been added in recent years to most operating systems and file systems.

FIGURE 5. DYNAMIC VOLUME EXPANSION

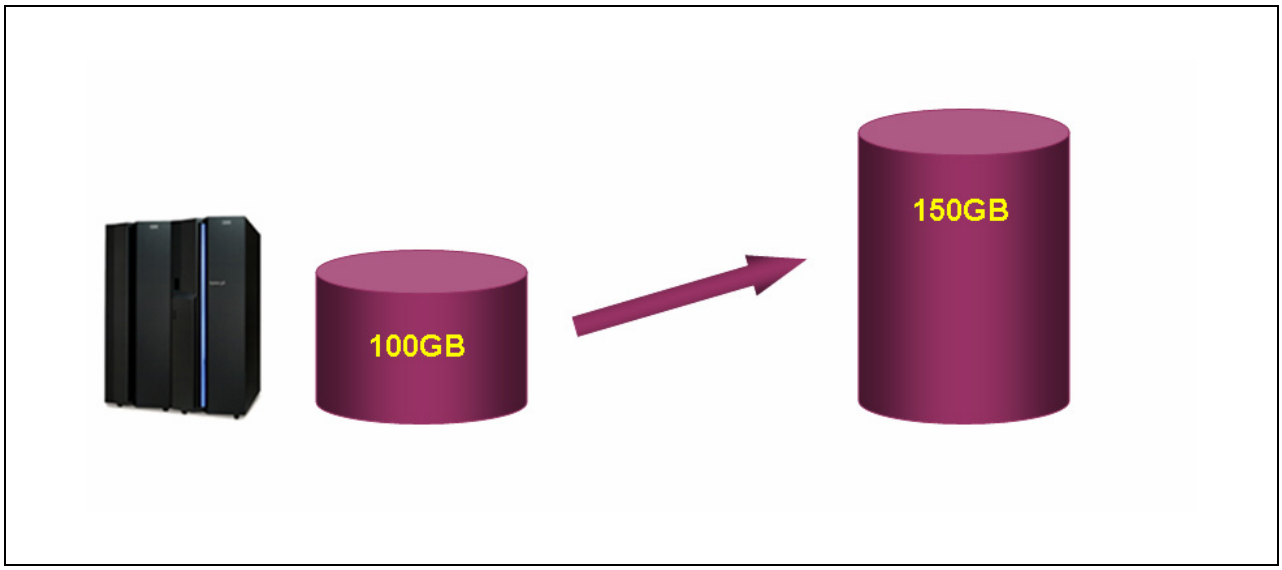

ESG Lab Testing

ESG Lab tested this feature using an AIX file system created over a 100 GB volume as shown in Figure 5. A one gigabyte file was copied to the volume to ensure that it was working properly. The size of the volume was increased by 50 GB as new files were copied to the volume. A single CLI command was used to increase the capacity of the volume to 150 GB. The new capacity was recognized from AIX using commands which expanded the volume and the file system. Files created before, during and after the expansion were verified by inspection less than three minutes after the dynamic volume expansion had begun.

Why This Matters

Traditionally, when an administrator needed to expand an open systems volume which was being accessed by an active application, she would need to back up the data, bring the volume offline and recreate the volume to expand it. While network-based copies and volume management software can be used to reduce the complexity and down-time, storage administrators often struggle with the soft and hard costs associated with allocating additional capacity for running applications. To avoid the pain associated with finding maintenance windows for application and file system expansion, administrators often tend to over-provision capacity, resulting in decreased storage utilization and increased capital equipment costs. ESG Lab has confirmed that dynamic volume expansion can be used to significantly reduce the cost and complexity of meeting growing capacity and stringent application availability requirements.

Storage Pool Striping

With release 3, the IBM DS8000 now supports the creation of volumes that are striped over more than eight disk drives. Windows, Linux and UNIX applications requiring access to very large volumes often require more than eight disk drives to meet performance and capacity needs. Until now, storage administrators have used host-based volume managers to solve this problem.

As shown in Figure 6, host-based striping was the traditional method used to combine multiple disk groups into a single volume. Storage Pool striping can now be used to create a single pool of capacity striped over dozens, and even hundreds, of disk drives. As drives are added to a pool, the performance capacity of the pool increases–particularly for common business applications like e-mail and databases that are interactive and mostly random in nature.

FIGURE 6. STORAGE POOL STRIPING

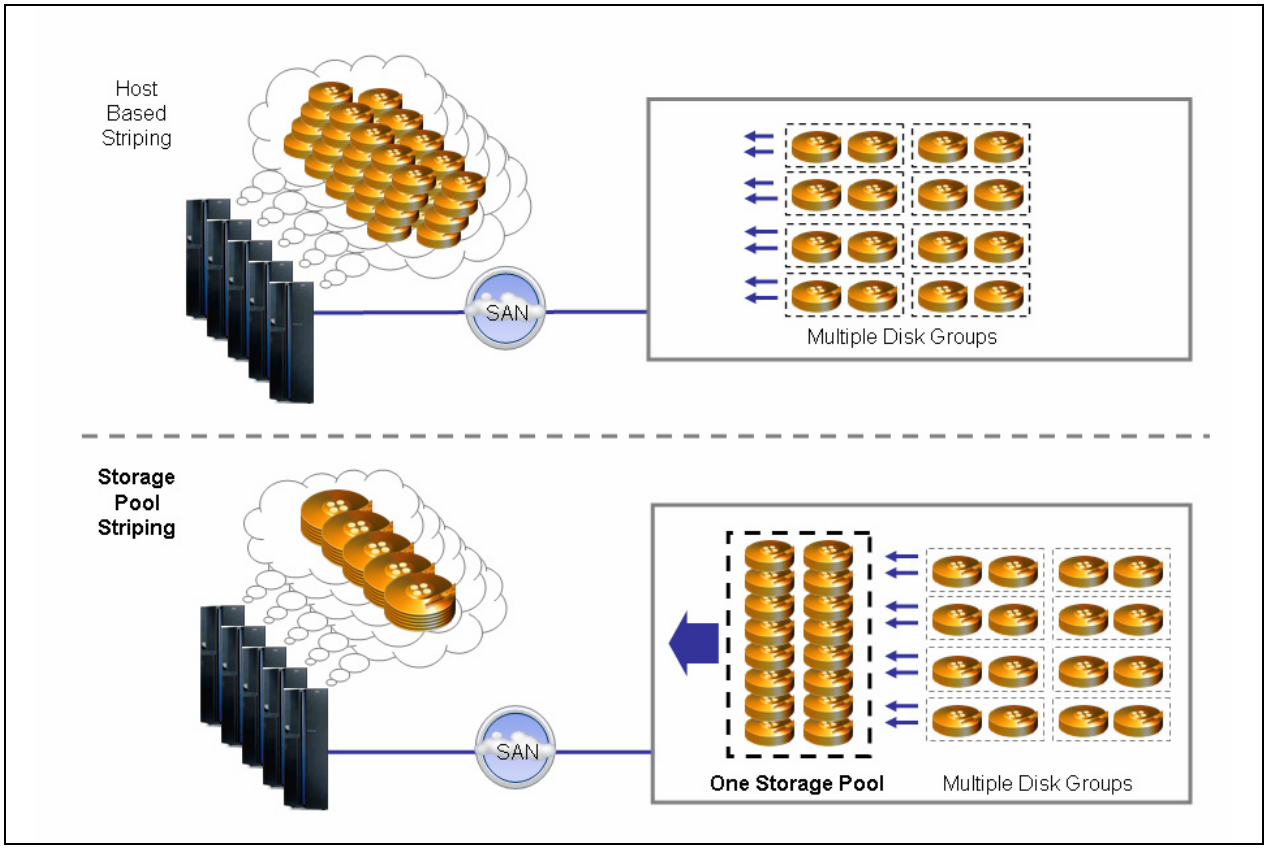

ESG Lab Testing

ESG Lab configured and exercised the storage pool striping feature and audited results of previously conducted tests. The performance of an eight drive pool was compared to a 16 drive pool via hands-on testing. A random 4KB workload composed of 70% reads and 30% writes was used during this phase of testing. The workload was designed to be serviced from disk and cache with a goal of a 50% cache hit rate. The workload was driven using 128 threads executing in parallel.

Tests audited by ESG Lab compared the same numbers of drives running a workload comprised of I/O operations designed to demonstrate the performance of a storage system running a business critical application. The workload was comprised of predominately random I/O operations requiring both queries as well as update operations (for example OLTP systems, databases systems, or mail server

applications). The majority of the I/O requests, meant to simulate update and query operations, were random 4KB I/O requests with a read/write mix of 50%. The balance of the I/O activity (28%) was composed of sequential writes with a variable requests size between 4 KB and 64 KB. The log activity was designed to run at 100% utilization on a pair of DS8000 volumes in the eight drive test case. This test was designed to mimic a skewed I/O workload distribution that often occurs with applications over time. The audited benefit of storage pool striping over 16 drives was compared to the benefits measured during ESG Lab testing. In both cases, performance nearly doubled with storage pool striping, as shown in Figure 7.

FIGURE 7. PERFORMANCE IMPROVEMENT WITH STORAGE POOL STRIPING

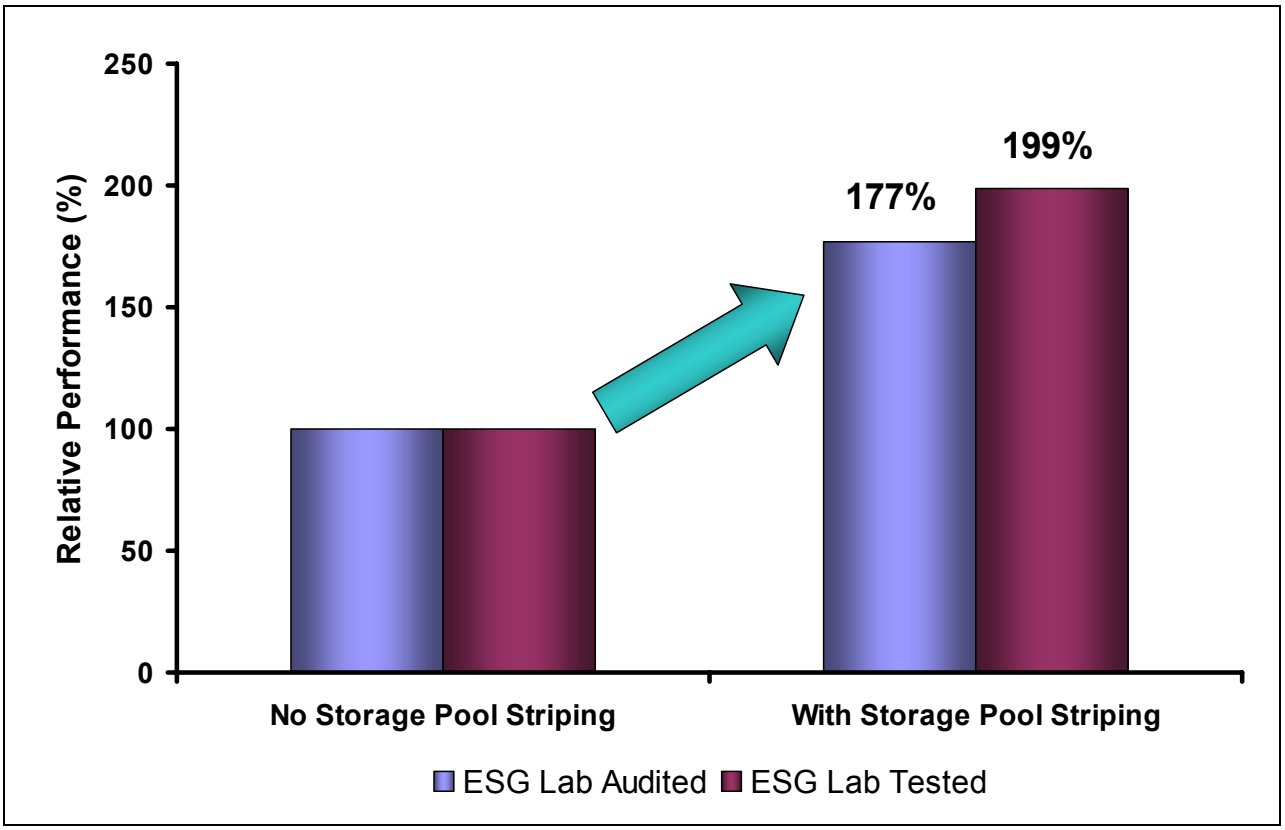

Why This Matters

Large multi-user application including e-mail and order entry systems often require shared access to a very large pool of high-performance storage capacity. Managing many storage volumes glued together using host-based volume managers running on many servers can be complex and expensive.

Storage Pool Striping gives administrators the ability to automatically stripe data across a large number of disk drives and present aggregated volumes to servers as a single virtual disk drive. This improves performance while significantly decreasing management effort—dozens, or even hundreds, of disk drives can process I/O operations simultaneously while practically eliminating the need for host-based volume management software.

Adaptive Multi-Stream Pre-Fetch (AMP)

Adaptive Multi-Stream Pre-fetch (AMP) is a recent IBM DS8000 Release 3 performance enhancement. AMP is designed to increase the performance of multiple, simultaneous sequential read workloads. The AMP algorithm detects sequential access and aggressively fetches data from disk before it is needed.

ESG Lab Testing

Testing was executed on a p595 server running AIX version as sequential read workloads were generated using the IBM SEQIO utility. The performance benefit of AMP technology is shown in Figure 8. In this test, ESG Lab performed one megabyte sequential reads to a single volume. The AIX multi-path volume manager was used to perform sequential reads through one, two and then four host bus adapters. To measure the advantage of AMP, the amount of data that could be pre-fetched was limited to the fixed Release 2 value via a CLI command. As can be seen here, performance nearly doubled with AMP. These results correlated well with a number of audited test results which compared the performance of Release 2, without AMP, to Release 3 with AMP.

FIGURE 8. THROUGHPUT WITH AND WITHOUT AMP

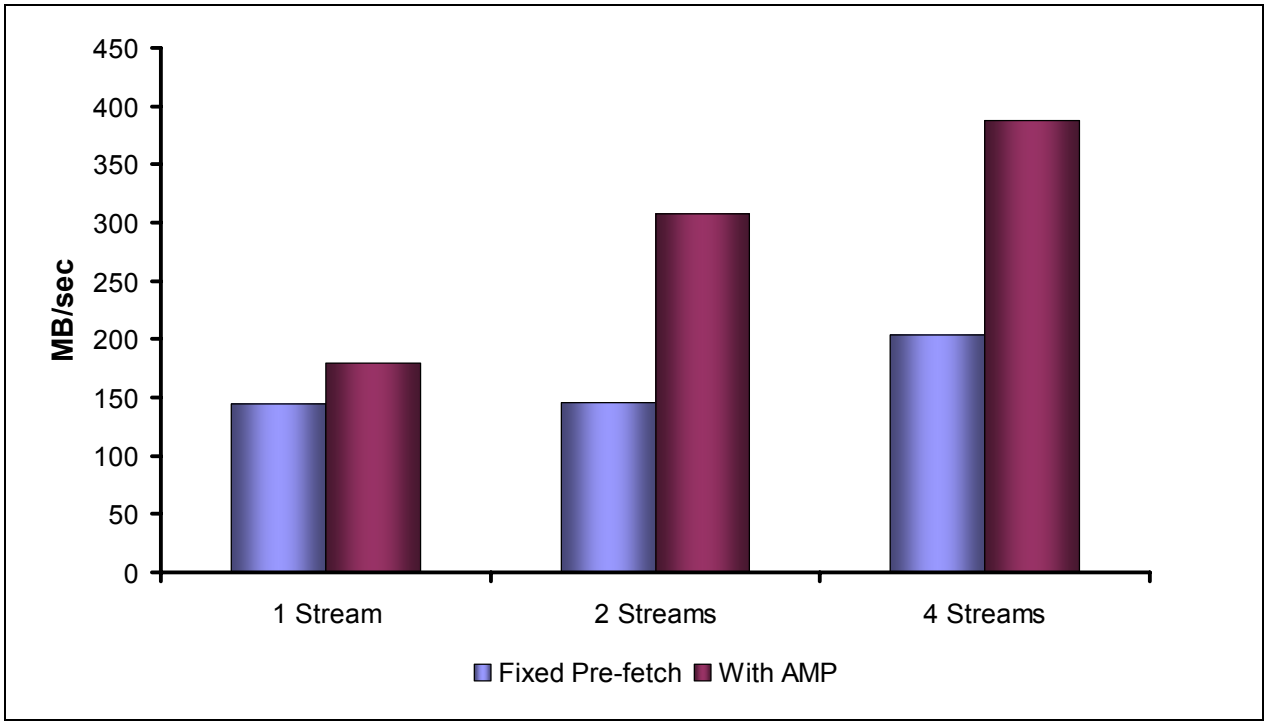

Why This Matters

Applications such as batch processing, backups, business intelligence and streaming media utilize large block, sequential I/O and require different performance characteristics than transactional workloads. Often storage systems can be tuned for one or the other but not both. This leads to higher expense (enterprises must purchase multiple arrays), or compromised performance (both types of traffic are impaired).

ESG Lab has confirmed that Adaptive Multi-stream Pre-fetching (AMP) nearly doubles throughput for applications that rely on large sequential reads. This can result in significantly shorter completion times for long running applications like batch processing, data mining, reporting and content delivery.

Performance

The IBM System Storage DS8000 utilizes a massively parallel high-speed architecture and intelligent software algorithms (e.g., AMP) to meet the high-performance needs of applications requiring parallel access to a large pool of shared storage. ESG Lab concluded the validation with a demonstration of the enterprise-class performance capabilities of the DS8300.

Performance testing began with throughput and IOPs characterization tests—often referred to as "the corners" in storage benchmark testing. "Corners" testing, which provides a good first impression of the theoretical limits of a storage system, was followed by a series of tests designed to measure the performance capabilities of the DS8000 when running real-world applications.

Corners characterization tests were performed using 112 logical units (LUNs) configured over 512 fast 15K RPM drives. The first test measured throughput, or how much data the DS8300 could move from disk in a given amount of time. Four 64 kilobyte sequential read streams were issued to each of the 112 LUNs. Cache was manually flushed between runs. As can be seen in the upper left corner of Figure 9, the DS8300 delivered 3.9 GB/sec (GB=1,000,000,000 Bytes) of aggregate throughput from disk with an average response time of 4.3 milliseconds.

FIGURE 9. DS8300 PERFORMANCE SCALABILITY

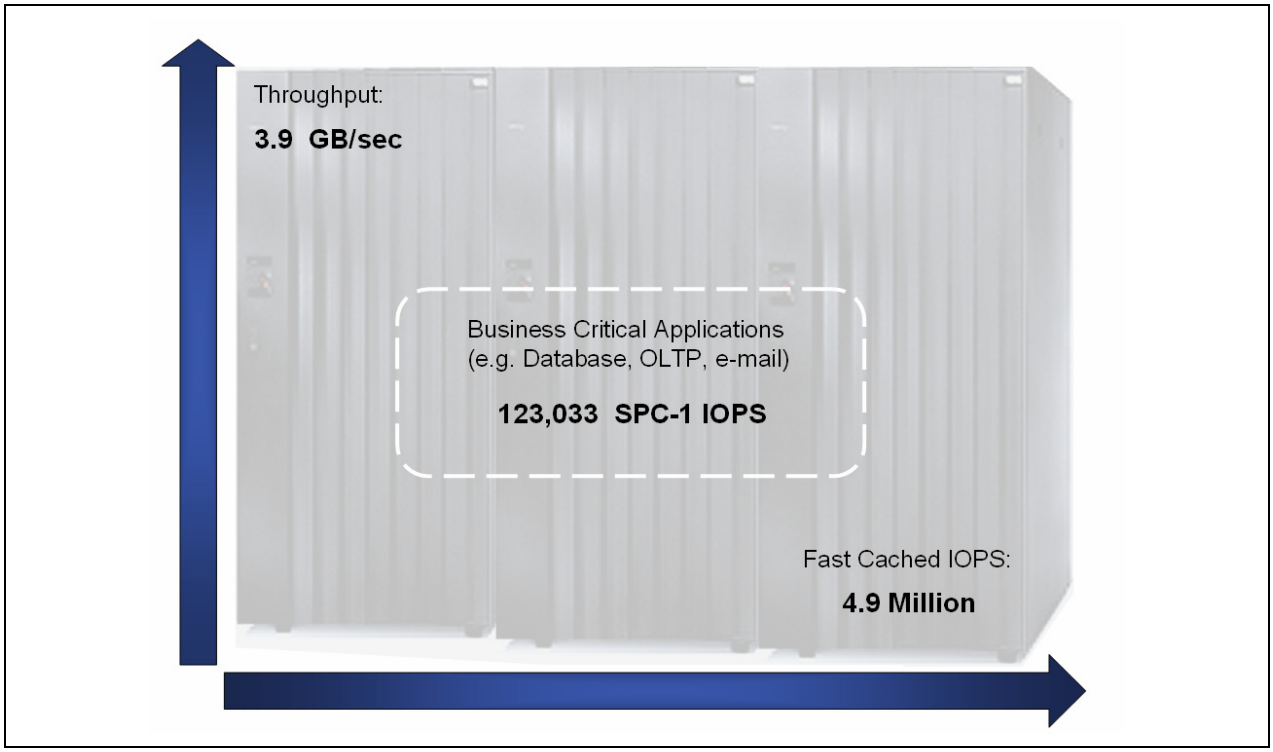

The second "corners" test measured the maximum I/Os per second (IOPS) that the DS8300 could sustain. IOPs are commonly used to demonstrate the processing power and front end efficiency of a disk array. Four 512 byte sequential read streams were run over a small region of disk that fit in the cache of each front end adapter. IBM refers to this rating as "fast cache IOPs." A series of tests were performed to learn that a 2 Gbps FC adapter can deliver 145,000 IOPS and that performance scaled linearly to 953,691 IOPS as the number of active adapters within a DS8000 was increased. It was noted during this stage of testing that the CPU utilization on the IBM p595 server was pegged at 100% which indicates that

the server had become the bottleneck. Audited 4 Gbps FC adapters results indicated that 154,000 IOPS per adapter can be used to sustain up to 4.9 million IOPs from a single IBM DS8300[3].

Moving from the corners to simulated real-world applications, ESG Lab audited the results of an application level industry standard SPC-1 benchmark posted on the Storage Performance Council (SPC) website.[4] The SPC-1 test is designed to emulate the typical functions of transaction-oriented, real-world applications. Examples of this type of application include on-line transaction systems, interactive database applications and e-mail messaging systems. As shown in Figure 9, a single IBM DS8300 delivered 123,033 SPC-1 transactions. The DS8300 also provided an excellent response time of less than 10 milliseconds at traffic levels up to 95% of its maximum transaction rate. Response time is an important metric included in an SPC report since it represents the delay that an application will experience (and pass on to users) when a storage system is stressed to the limit.

Why This Matters

Business units that rely on information assets housed within a shared storage system can suffer when a storage system is performing poorly. Productivity can suffer and, in the worst case, poor performance can lead to a loss in revenue. While some storage systems may be purpose-built to excel at one corner of performance—for example, throughput—ESG Lab has confirmed that the DS8000 performs well for a mix of applications including some of the most business-critical, multi-user applications in the world.

ESG Lab was particularly impressed with IBM's decision to publicly demonstrate DS8000 performance capabilities, including peer-reviewed industry standard SPC-1 benchmark results which prove that the DS8000 is well-suited to meet the performance needs of enterprise-class, mission-critical applications. This was a gutsy move in the often secretive world of high end, enterprise class storage. We believe that IBM's publishing of SPC-1 results has encouraged at least one other vendor in the high end space to publish and hopes that it will prompt participation by others.

[3] For a Turbo DS8300 configured with a maximum of 32 4 Gbps FC front end adapters.
[4] http://www.storageperformance.org/results/benchmark_results_spc1

ESG Lab Validation Highlights

- ☑ ESG Lab used the FlashCopy SE utility to create a point-in-time copy of a 10 GB application volume. The space efficient copy of the 10 GB volume consumed only 132 MB of disk capacity which represented the writes since the flash copy was made.

- ☑ Dynamic Volume Expansion was used to quickly and easily expand an AIX file system built over a DS8000 volume from 100 GB to 150 GB.

- ☑ Storage Pool Striping was used to create volumes striped internally within a DS8000 over 16 drives without host tuning or configuration. Application performance for simulated interactive applications that are mostly random in nature (e.g., e-mail, database) increased by 100% when the number of drives per volume was doubled.

- ☑ Performance nearly doubled for large block sequential read workloads when comparing Release 2 throughput without Adaptive Multi-stream Pre-fetch (AMP) to Release 3 with AMP.

- ☑ An excellent enterprise-class performance characterization of 3.9 GB/sec of aggregate throughput capability and 4.9M IOPs were measured. Strong industry audited SPC-1 results, which are a good indicator of real-world application performance, were audited.

Issues to Consider

- ☑ While ESG Lab testing was performed on a server running the AIX operating system with an intended audience of open systems (i.e., Windows, Linux and UNIX) administrators, it should be noted that Release 3 capabilities including Flash Copy SE, AMP and Dynamic Volume Expansion are supported on zSeries servers (i.e., mainframes).

- ☑ Flash Copy SE is well suited for applications and workflows that require a quick copy for reading (e.g., backups), but is not recommended for an update-intensive application that does a lot of writes (e.g., testing a write-intensive month-end batch update utility). Traditional full volume Flash Copies are better suited for write-intensive workflows.

- ☑ The Dynamic Volume Expansion capability requires a moment of application unavailability and an operating system specific command before applications can take advantage of additional capacity.

ESG Lab's View

Today's organizations are faced with a dizzying array of data management challenges. Accelerating data growth, increasing cost and complexity, a growing number of compliance and governance initiatives and power and cooling shortages are driving a new wave of consolidation and virtualization in the data center. Storage systems with high capacity, extraordinary performance and advanced functionality are needed to meet the needs of consolidating applications.

ESG has been tracking the success of the IBM DS8000 series for a number of years. Improvements have been made with each release in the areas which matter most in an enterprise-class storage system: availability, scalability, functionality and performance. ESG Lab has confirmed that innovative DS8000 Release 3 enhancements including space efficient flash copies, storage pool striping, adaptive multi-stream pre-fetch and dynamic volume expansion work as advertised and are well-suited to meet the needs of medium to large organizations requiring high capacity and smarter virtualization capabilities.

New capabilities in the latest release of IBM DS8000 software including AMP and storage pool striping have improved the performance capabilities and flexibility of IBM's flagship enterprise-class storage offering. While some storage systems may be purpose-built to excel at one corner of performance, ESG Lab has validated that the DS8000 can serve a mix of workloads, including some of the world's most business-critical, multi-user applications. ESG Lab was impressed with IBM's decision to publicly demonstrate the DS8000's performance capabilities, including industry standard SPC-1 benchmark results, which prove that the DS8000 is well-suited to meet the performance needs of mission-critical applications.

Based on the results of the hands-on tests presented in this report, ESG Lab believes that the IBM DS8000 is well-suited for medium to large organizations requiring high capacity, even higher performance and smarter virtualization capabilities to meet the growing needs of business units operating in an on-demand world.

Appendix

Table 1. TEST CONFIGURATION

Hardware	Software
3-Frame DS8300, 512 15K RPM FC drives, mix of 73 and 144 GB, RAID-10, 16 dual port 2 Gbps front end FC adapters	Release 3
IBM P595 server, 32x1.9Ghz P5 CPUs, 64 cores, 38 GB RAM	AIX 5.3.0.60, SEQIO, SPC-1

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

www.enterprisestrategygroup.com