The IBM TotalStorage Enterprise Storage Server

Proven High Performance

Author:
David J. Sacks

A **Technology Insights** White Paper from IBM

The Enterprise Storage Server (ESS) is IBM's industry-leading disk storage system.

The ESS is designed to meet - and generally exceed - the stringent requirements that demanding customers place on high-end enterprise-class disk storage systems. With its sophisticated server/OS-based architecture, the ESS is intended to provide demonstrable business values such as 24x365 access to information, easier management of storage resources, extensive function, low total cost of ownership -- and outstanding performance. With concurrent support for a broad range of processors - from small PC servers to UNIX-based servers to the largest mainframes - ESS is truly a storage solution for the entire enterprise. And ESS is backed by IBM quality and support.

That said, a key trait of the ESS that has been satisfying customers since the first generation ESS became available just a few years ago is its ***raw speed***. ***Low*** response times. ***High*** I/Os per second. Specialized performance ***accelerators***. And ***optimized*** copy functions.

Anyone can claim their product delivers high performance. This paper documents customer tests and industry-standard benchmarks that demonstrate it for the ESS.

Notices

Published August 12, 2003

Please send comments via e-mail to:
David J. Sacks, Senior Storage Systems Consultant
djsacks@us.ibm.com

IBM TotalStorage Enterprise Storage Server

Note: This ESS was artificially slowed down so it would not appear blurred in this photograph.

Contents

Introduction

Since its introduction in 1999, IBM's Enterprise Storage Server (ESS), often referred to by its code name "Shark", has gathered accolades from customers and industry analysts alike.

This paper provides multiple perspectives on ESS performance, including descriptions of ESS facilities that contribute to its high performance, and numerous comments from customers and industry analysts. Of particular interest, the paper includes multiple "Proof Points" that document the results of ESS performance benchmarks. One proof point shows the results of running the vendor-neutral industry-standard SPC-1 benchmark, while other proof points document benchmarks run by customers to compare the ESS to other disk systems. These proof points focus on performance metrics most important to applications: response time and throughput.

For additional information on the IBM Enterprise Storage Server and other IBM storage products please contact your IBM storage sales representative. Or visit the IBM storage Web site - the home page for the Enterprise Storage Server is currently: www.ibm.com/totalstorage/ess.

Enterprise Storage Server - Value Highlights

This table summarizes some of the most important attributes of the ESS on one page. Of course, one page can't identify every benefit of this advanced disk system. Please ask your IBM storage sales representative or visit www.ibm.com/totalstorage/ess for information on other significant benefits the ESS can bring to your business.

ESS - Value Highlights
Comprehensive Storage Consolidation Platform
Scalable to 55.9TB physical capacity
Configurable as RAID-5 and RAID-10, intermixable under customer control
Fibre Channel, FICON, UltraSCSI, and ESCON host connections
Up to 32 direct host connections; hundreds of hosts can be connected via a SAN
Capacity can be reassigned among servers under customer control while the ESS is online
Fibre Channel port sharing and SAN LUN masking support built-in
Outstanding Performance
SPC-1 standard benchmark results published (www.storageperformance.org)
Up to 64GB of system cache
Up to 64 unarbitrated and pipelined internal disk paths
Up to 32 direct host connections / all can be transferring data concurrently
2Gbit/second Fibre Channel and FICON
All data is automatically striped for balanced disk utilization and high parallelism
Intermixable disk capacities at 10K and 15K RPM
Specialized IBM zSeries performance accelerators
Round-the-Clock Access to Data
Hardware fault tolerance, including dual write caches
Concurrent repairs ("hot swap")
Concurrent hardware upgrades
Concurrent software/microcode changes
Ease-of-Management / Operations / Maintenance
Web-based, CLI, and SMI-S compliant API management interfaces
Logical volumes presented to hosts in customized capacities
Auto-call home for service / scheduled proactive call home to confirm status
Remote problem analysis / diagnostics
Multiple event customer notification facilities: SNMP, e-mail, pager
Storage Networking Support
Flexible Fibre Channel SAN support including directors, hubs, and switches
SAN LUN masking (security) support standard
Support for CIFS, FTP, HTTP, NetWare, and NFS file-I/O (NAS) protocols
Support for iSCSI via a CISCO Storage Router
Advanced Software-based Value-add Capabilities
Internal copy - high-speed volume copy with specialized performance optimizations for short-term and long-term copies; consistency group support to avoid application downtime for multi-volume copies; data set (file) copy for z/OS
Remote Copy - synchronous and asynchronous protocols; cascading configurations provide additional function such as data currency at long distances
Host multipathing software (Subsystem Device Driver for load balancing and failover)
Attractive Cost of Ownership
Many popular features are standard (e.g., system management, host multipath software)
High performance space-efficient RAID-5 option lowers cost and increases scalability
Fibre Channel / FICON adapters can be used for either protocol
Three year hardware and software feature warranty are standard

The Value of Performance

.

The time it takes to perform I/O operations is often the single largest contributor to response time, particularly in commercial environments. Thus, a faster disk system will often have a larger impact on overall application performance than faster host processors. This shows up in two measures: I/O response time and I/O throughput - in simple terms, "how fast" and "how many."

Improved performance is more than just speed for its own sake - it makes a significant difference to the way you run your business, and to the way users and customers perceive the quality of your business.

Potential Value to You:

Better service to users and customers

Improved productivity

Batch work completed in less time

Better handling of peak or unpredictable I/O workloads

Reduced I/O bottleneck

Reduced need for performance tuning

Minimized application delays waiting for backup tasks

Increased flexibility to consolidate distributed storage into a single disk system

Increased flexibility to mix applications together in one disk system

Reduced need to replicate data for performance reasons

Reduced application development time

Extended processor life

How the ESS Delivers the Value of Performance

Anyone can assert that a given disk system provides high levels of performance. But performance leadership has to be earned, not just claimed. It is earned by incorporating state-of-the-art technology that provides a platform for superior performance, and then by demonstrating superior performance repeatedly and consistently in different environments.

Consider the following list of ESS facilities that make the promise of outstanding performance a reality:

- Balanced high performance design can benefit both cache-friendly and cache-unfriendly workloads
- Large read/write cache can satisfy many I/O requests at electronic speeds. Multiple algorithms optimize cache efficiency. (Examples: 1) a read miss initiates staging into cache the optimum amount of data based on ESS analysis of current I/O patterns. 2) sequential read-ahead plus accelerated discarding of such data from cache after it is read by applications makes more cache available to other users.)
- All host connections can transfer data concurrently
- Up to 64 data transfers to/from disk can concurrently be in process across internal disk adapters. Further, every single disk can be in the process of transferring data in an interleaved manner over pipelined paths.
- *All* logical volumes are automatically striped across multiple physical disks providing: parallel cache-miss processing even for a single logical volume, reduction of disk "hot spots" by naturally balancing utilization across disks, faster sequential throughput due to efficient ESS sequential processing algorithms, and, as a result of these benefits, the additional benefit of reduced manual tuning. The use of a relatively small "strip" size further contributes to ESS striping efficiency.
- Disks have high-performance characteristics, including 10,000 RPM and 15,000 RPM options
- Disk capacity selection (18.2GB, 36.4GB, 72.8GB, 145.6GB) helps optimize price/performance
- Dual active n-way SMP RISC processor clusters manage many system activities, offloading that work from the ESS internal host and disk adapters
- Floating spare design eliminates the overhead of moving data back from a spare disk to a replaced disk
- Near-instantaneous internal data replication (FlashCopy). Users can either make a full byte-for-byte physical copy, or can select one of two performance optimizations for a given copy operation. "Copy-on-first-write" (NOCOPY) is optimized for short-term copies, minimizing internal data movement and freeing target volumes for other uses. "Incremental" update (a.k.a. resync or refresh) is optimized for periodic updates of long-term copies.
- RAID-5 and RAID-10 options which can be intermixed. The high-performance RAID-5 implementation satisfies the performance requirements of most workloads while lowering costs and increasing scalability. RAID-10 may improve performance for very high random-write workloads.
- RAID rebuild, invoked in the case of a physical disk failure, favors preserving performance of production I/O
- RAID management is localized to the internal disk adapters, offloading that overhead from cache and the internal paths to cache
- Remote mirroring performance optimizations and selection of protocols
- Dual internal clusters, multiple RAID arrays, and separate internal paths to groups of disks allow optional workload separation to help control resource contention
- Fast I/O performance can offset application delays caused by higher processor utilization
- The ESS can provide a high-speed replacement for traditional tape: no manual or robotic delays.
- An ESS Model 800 *Turbo* feature can provide an additional performance boost for the most demanding workloads
- Fibre Channel paths can be used for inter-system remote copy links. (Publicly previewed by IBM and planned for fourth quarter 2003 availability.)

Platform-specific items:

- Midrange platforms:
  - 2Gbit/s Fibre Channel host connections (also supports 1Gbit/s speeds)

- Command Tag Queuing for both host connections and internal disk paths supports increased I/O parallelism.
- Host path I/O load balancing: the Subsystem Device Driver, a standard host-based software facility that ships with the ESS, balances host I/O traffic over multiple paths to the same volumes (as well as supporting path fault tolerance). Supported host platforms includes AIX, HP-UX, Linux, Windows, and Sun Solaris. (Note that some operating systems and volume managers include their own multi-path support.)

- zSeries (S/390):
  - 2Gbit/s FICON channel support (also supports 1Gbit/s speeds)
  - Significantly increased I/O parallelism for z/OS (OS/390) via the Parallel Access Volume (PAV) feature, a major advance over conventional disk systems that only support one I/O to one logical volume at a time from a given system image. Up to 256 concurrent I/Os per volume are supported.
  - Multiple system increased I/O parallelism via the Multiple Allegiance feature, a major advance over conventional disk systems that only support one I/O to one logical volume at a time from multiple system images. This feature also supports a fair access facility, designed so that a faster processor cannot delay I/Os from a slower processor.
  - Application priority support for z/OS (OS/390) via the Priority I/O Queuing feature, a significant advance over conventional disk systems that process I/Os only in first-come first-served order
  - Performance improvements through ESS support for improved channel commands
  - Asynchronous remote mirroring to minimize application delays (z/OS Extended Remote Copy, (XRC))
  - FlashCopy can optionally replicate individual data sets
  - Intra-job FlashCopy backups can eliminate tape usage/delays
  - Support for the sequential-notification bit set in channel programs by operating systems (in addition to ESS sequential-detect algorithms); this is more efficient than disk systems that rely only on sequential-detect.

What People have been Saying

The ESS has a history of praise for its performance capabilities. The current generation, the ESS model 800, provides enhanced performance well beyond the high levels of performance of previous ESS generations.

"The Enterprise Storage Server has dramatically improved our batch and online throughput performance, increased our operating efficiency, and enabled us to process data much faster. The results are extremely positive."

-- Dave Corley, director of Information Systems Operations, Western Resources

"The entire [reservation] system relies on the availability and performance of the database. At peak, Amadeus is processing 150,000 I/Os per second. If there had been any question about the availability or performance, then we would have looked elsewhere. As it is, not only does the ESS [model F20] deliver on our demanding availability targets, the performance has impressed us in the extreme!"

-- Hamilton Baird, Vice President of Operations and Deputy General Manager of Data Processing, Amadeus

"...IBM's technology is superior [to EMC's]. It wasn't the cost. It was performance."

-- William Homa, Chief Information Officer at Hannaford Bros, quoted in Forbes Magazine article "What's Eating EMC", Daniel Lyons, November 26, 2001.

"With the implementation of both Parallel Access Volumes and Multiple Allegiance, both IOSQ time and PEND time were virtually eliminated. These components of Response Time have throttled the maximum throughput we were able to achieve through our storage subsystems. With the ESS, those bottlenecks have been substantially reduced or eliminated. This enhancement is one of the most significant performance changes to a storage subsystem since the introduction of DASD Fast Write."

-- Wayne Hartman, Systems Director of Storage Systems, Prudential Securities

"SAP is enthusiastic about the specifications of IBM's Enterprise Storage Server... With the Enterprise Storage Server's Parallel Access Volumes capability, SAP R/3 fully expects to improve the performance of its parallel disk IO intensive R/3 tasks..."

-- SAP AG

"To call upon IBM...for us and our customers meant benefiting both from the credibility of a large manufacturer and from a reliable and powerful [ESS] storage solution at a modest cost"
--- Denis Harscoat, founder and chairman of Zden

"Consolidated storage will lower the Total Cost of Ownership and protect investment. ESS looks like a well constructed machine delivering the latest SSA technology; it is fairly flexible and quite easy to manage. We are very happy with the ESS performance as well..."
-- Director of IT for an international energy company

A white paper by a respected industry consultant, Dr. H. Pat Artis of Performance Associates, Inc., provides an in-depth analysis of the OS/390 Parallel Access Volumes (PAV) and Multiple Allegiance performance accelerators supported by the ESS. The paper concludes that these facilities can provide significant benefits. The paper is available on the Internet at www.perfassoc.com.

"As part of IBM's Early Support Program, we've experienced excellent performance and availability from ESS. During our testing period, we found that ESS was easy to use for both S/390 and open storage."
-- Tom Lund, executive vice president for Pengeinstitutternes Betalings Systemer, PBS, (Danish Payment Systems, Ltd.)

"...our tests [of ESS] have shown the performance to be exceptional....Over time, we aim to simplify our storage management procedures by consolidating all disk storage onto a single disk-based architecture, the Enterprise Storage Server."
-- Safeway PLC, United Kingdom

" 'The new Sharks help improve online response times, and it cuts an hour off of our batch processing cycle...That's a huge number to any company.'...heavily-used data on the company's site now can be accessed more quickly. That includes master files used for insurance and common database files."

-- From "IBM Shark tears into rival EMC", CBS MarketWatch, March 18, 2000. The customer quoted is Greg Scroggs, a senior vice president in charge of data services for Primerica Financial Services

"It is apparent Shark is superior in design and architecture and will blow the competition away."

-- Executive at a large New York insurance company

"Having a centralized, highly available, ultra-fast storage system across our S/390, UNIX and NT platforms that is tightly coupled with our disaster recovery and business resumption plans is critical to our strategic IT vision. 'Shark' delivers on this goal."

-- David R. Wade, executive vice president and chief information officer at Primerica Financial Services

"ESS' price/performance ratio is extremely good and the management facilities are excellent."

-- Wilfried Vergeer, Systems Specialist, Delta Lloyd Zorg, financial services provider

"The FlashCopy features offered by the ESS make it possible to make a complete copy of the production environment available for a test or for the data warehouse within two minutes."

-- Michaël de Mul, IT consultancy Rodata

"We evaluated both EMC and IBM, and went with the Shark because IBM provided the best total picture -- from performance to service."

-- Gene Berger, Chief Technology Officer, Pacific Data Technology, Inc.

"... the IBM [Shark] offering 'has exceeded the competition in both speed, reliability and availability'."

-- Joe Dwyer, VP systems performance and services, Credit Suisse First Boston, quoted in *Security Industry Online*, November 13, 2000

"We needed to be able to ensure that our services improve productivity and bottom line. The combined IBM RS/6000 and Shark Storage servers answered the challenge -- while neither Sun nor EMC could match their price and performance. We selected the Shark because of its price and function and weren't persuaded by EMC, especially since it was unwilling to provide performance specifications for the Symmetrix."

-- Gene Deans, Database Administrator, EchoMail, October, 2000, from an Internet Wire press release

"We went with Shark for its superior performance and outstanding value. With Shark's new capabilities -- particularly Peer-to-Peer Remote Copy -- IBM has delivered the industry's premier, open storage subsystem."

-- Greg Scroggs, Senior Vice President and Data Center Manager at Primerica, a subsidiary of CitiGroup, December 14, 2000

"With IBM's Shark servers we've seen a great performance boost from our DB2 applications and the overnight processing of our clinical and financial records as well as the nightly data backup has improved by 40 percent."

-- Nelson Gray, IT Manager, Shared Medical Services, December, 2000

"Our disk storage capacity was maxed out. We needed to resolve response problems caused by slow disks. The City of Boston needed a storage solution that offered the best technology for today, not just a promise for something better in the future," said Jack Malinsky, director of Operational Technology, City of Boston, in describing why the City selected IBM over EMC to manage the data intensive applications. "In looking at a solution to these problems both EMC and IBM were seriously considered. The performance benefits of the Shark sold us."

--Jack Malinsky, Director of Operational Technology, City of Boston, January, 2001

Performance Perspectives - Words of Wisdom

Disk system internal bandwidth has received a lot of focus in the industry as vendors compete for "bragging rights". But the simple fact is that internal bandwidth is not a useful indicator of external system performance, which is the performance customers actually care about and that applications can actually attain. That is one reason why benchmarks are more important than the performance attributes of isolated internal components, no matter how impressive. For more information on the many drawbacks of disk system internal bandwidth as a useful way to compare the performance of different disk systems, please refer to "Disk Systems and the Internal Bandwidth Wars", a white paper currently available on the public IBM Web at http://www.storage.ibm.com/disk/ess/whitepaper.htm .

Many system attributes can impact disk system performance. Here is a partial list:

- system architecture/design
- installed configuration (model and options) vs. a "brochure" configuration
- interaction of multiple components in the data path of any particular I/O request
- disk system host port bandwidth (individual and aggregate)
- host-to-disk system path multipathing host-based software
- disk system internal disk adapter bandwidth (individual and aggregate)
- disk system disk path bandwidth (individual and aggregate); number of disks per path
- disk system path protocols (arbitration overhead or not, number of transfers at a time per path)
- disk drives (number and speed)
- internal data copy services implementation
- external (remote) data copy services implementation
- cache size and management algorithms
- I/O workload characteristics (I/Os per second, read:write ratio, random/sequential mix, cache hit ratios, data transfer sizes, data reference patterns)
- manual management including data placement and (ongoing) tuning
- the configuration and use of any optional performance-related system features

The number, variety, and interaction of the elements in the above list (which, again, is only a partial list) shows how very difficult it can be to compare the performance of different disk systems based on nothing more than specification sheets, vendor claims, or even analyst opinions if they are based on those kinds of factors. All of this reemphasizes the importance of measured benchmarks to help compare the performance of different vendor's systems.

Performance Leadership - "Proof Points" of the Pudding

The ESS earns its high performance credentials in three ways: by incorporating the state-of-the-art technology and design that make high performance possible, by outstanding results based on the vendor-neutral SPC-1 benchmark, and by actually delivering high performance in customer environments.

The following pages contain ESS performance results - all are actual measurements.

Rather than going through the time, effort, and expense of doing their own measurements, it appears that in practice many customers base performance assessments and comparisons on various sources including prior product experience, the experiences of other customers, the results of industry-standard vendor-neutral benchmarks, the opinion of selected industry watchers, and on some understanding of disk system designs. Some customers do choose to run their own performance tests and comparisons. The ESS measurements reported back by customers are not always documented in the detail an engineer would provide, but that is how real customers often report the results of in-house performance tests.

Because the mainframe (IBM zSeries) environment generally has better tools for performance reporting, most benchmarks tend to be run in that environment. The majority of ESS components and function - such as internal paths, cache, and disk drives - apply to I/O requests from all server platforms.

Measurements were made by customers unless otherwise noted. The customer measurements were not audited by IBM. As you'll see, the performance advantages of the ESS are consistently demonstrated across different customers and different types of workloads.

POINTS TO KEEP IN MIND WHILE VIEWING THE BENCHMARK RESULTS

- Some customer benchmarks used second-generation ESS "F" models, not current generation "800" models. The ESS model used is identified in the proof-point description. As a guideline, the newer ESS 800 model delivers up to 2 to 2.5 times the performance of the earlier ESS F model.
- A given test does not necessarily show maximum ESS capabilities, even if the ESS provided superior performance in the test.
- Performance results were reported by IBM customers or achieved under controlled conditions. Performance results that may be obtained in other environments or by other users may vary.
- When considering the meaning of benchmark results, keep in mind that what is most pertinent is whether a system can meet your performance requirements. Beyond that, higher speed is of little if any business value, and other criteria are likely more useful in a comprehensive product comparison.

Performance Proof Point #1

Bottom Line:

The ESS 800 model delivers very high levels of performance in an open systems environment as measured by the vendor-neutral SPC-1 standard benchmark, an OLTP-like benchmark.

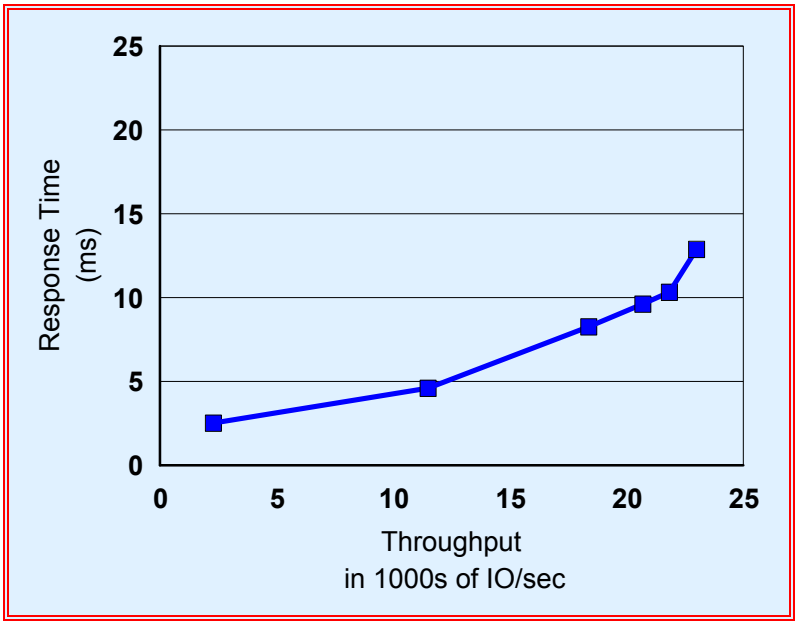

The SPC-1 benchmark is created and owned by the Storage Performance Council, www.storageperformance.org. It is an open systems-based benchmark with OLTP-like characteristics. The SPC is a multi-vendor organization that creates and promotes open, vendor-neutral disk system performance benchmarks. Vendors can acquire the SPC-1 benchmark from SPC and publish audited measurement results on the SPC web site. More details about SPC-1 and about the IBM ESS 800 benchmark are available at the SPC web site.

IBM for the ESS, and many other vendors for their disk systems have published SPC-1 results. EMC has not been willing to publish SPC-1 results for Symmetrix or CLARiiON. A noted industry analyst has this to say:

> "Performance 'benchmarketing' is a long-practiced art in this business," says Randy Kerns, senior analyst at Evaluator Group. "Don't believe anything regarding performance benchmarks from vendors. Skepticism is a good thing. Use the SPC benchmarks, and if a company doesn't have any, they probably don't compare favorably with the competition."
>
> - *Is NetApp SANbagging?,* Byte and Switch, 3/19/2003, By Todd Spangler (www.byteandswitch.com)

Performance Proof Point #2

Bottom Line:

Three ESS 800 systems outperformed three Symmetrix 8830 systems in an IBM mainframe environment.

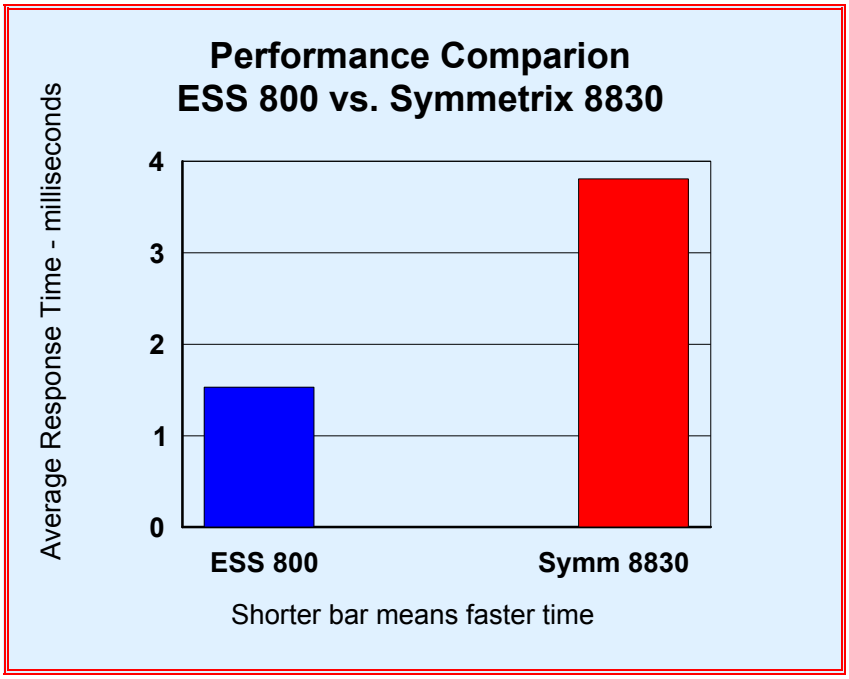

Background:

The customer, a large financial institution in Asia, migrated nearly all volumes from three Symmetrix 8830 systems to three ESS 800 systems. The systems were attached to six z/OS servers.

After the migration completed, the 1,981 logical volumes migrated from the Symmetrix 8830s were measured to run 59.8% faster on the ESS 800: 1.53 milliseconds average per I/O operation versus 3.81 milliseconds. Average "before and after" I/O rates were comparable: 4,541 on the Symmetrix 8830s and 3,863 on the ESS 800s, a difference of about 15%. Peak I/O loads were in the neighborhood of 10,000 I/Os per second in both environments. The relatively few volumes not migrated to the ESS systems had very low average activity, under 100 I/Os per second.

Configurations: The Symmetrix systems each had 16GB of cache and six FICON channels. The ESS systems each had 32GB of cache and six FICON channels. The larger cache increased the average read hit ratio, but only from about 86% to 92%. The ESSs were configured with the PAV (Parallel Access Volumes) feature while the Symmetrix systems were not, but queuing delays were low before the migration (at about 10% of average Symmetrix response time), so most of the 60% average performance improvement on the ESS systems was due to other system characteristics. Measurements were taken from RMF data.

Performance Proof Point #3

Bottom Line:

The ESS F model processed a random access workload faster than an HDS 9900 subsystem in an IBM mainframe environment.

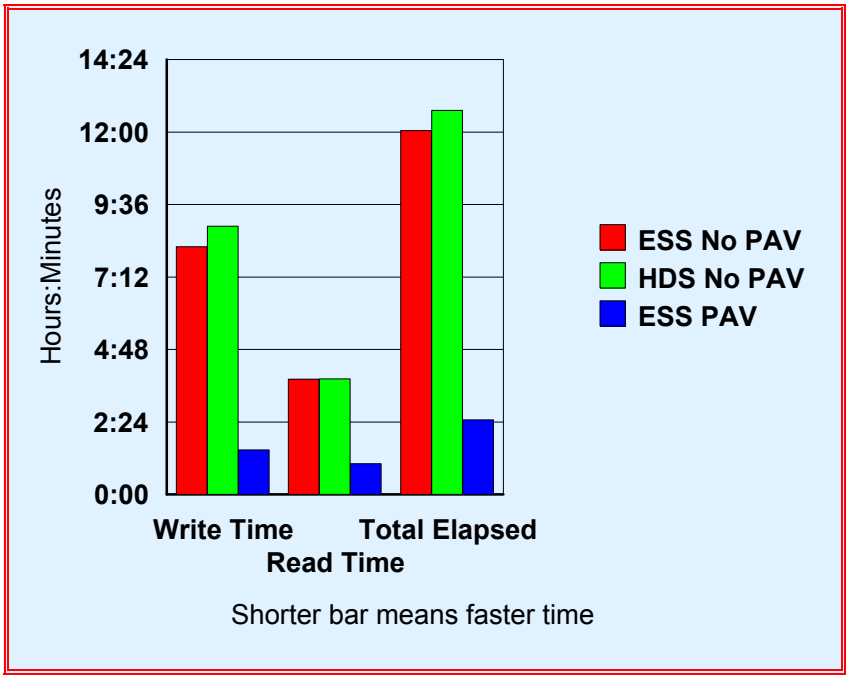

Background:

This customer, a large retailer based in the Midwest area of the United States, compared the HDS 9960 to IBM's ESS F model with a random I/O workload using their own benchmark in an IBM mainframe environment. The benchmark was representative of one type of application this customer runs. The workload consisted of ten jobs running concurrently, each job issuing a combination of 500,000 random reads and 500,000 random writes to a VSAM KSDS data set on a single volume.

Configurations: Both disk systems had 16GB cache, an equal number of ESCON channels, and 36GB disks. Both disk systems used RAID-5 protection.

The ESS was tested with and without PAV. (Parallel Access Volumes is an optional ESS feature that improves performance by adding support for parallel I/Os to the same disk address (volume), analogous to SCSI Command Tag Queuing in the midrange environment.) The HDS 9960 did not have PAV installed.

Without PAV, the ESS completed the workload in slightly less time than the HDS 9960. With PAV, the ESS ran significantly faster than the HDS 9960.

Performance Proof Point #4

Bottom Line:

The ESS F model outperformed both an EMC Symmetrix 8730 and HDS 9960 in an online transaction environment.

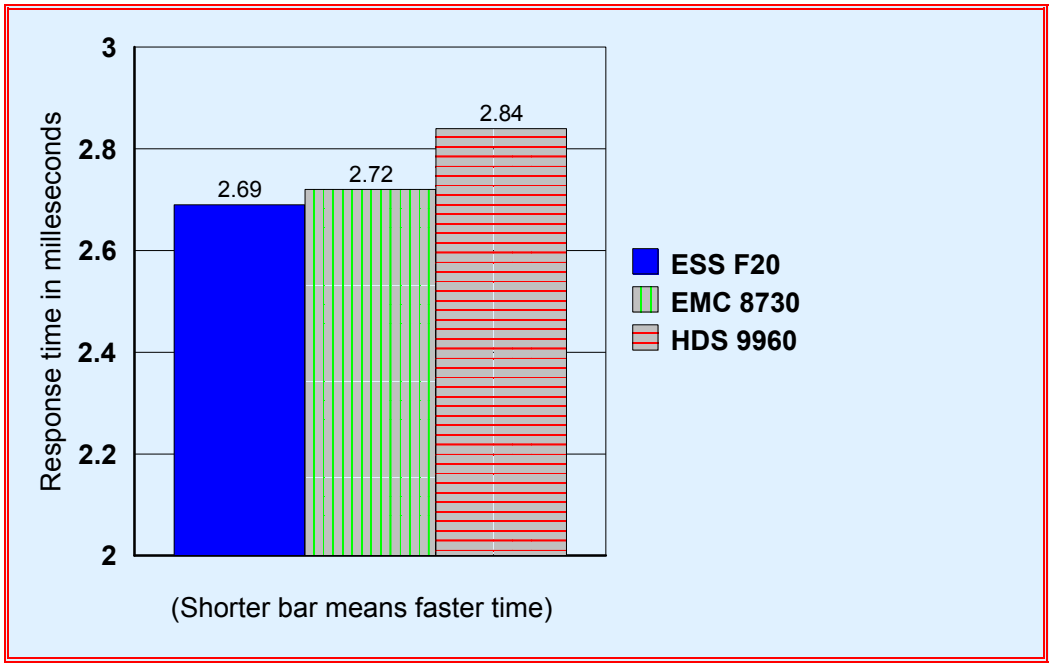

Background:

The tests were run by a North American customer in the finance industry in late 2001.

Each disk system vendor was requested to provide the following configuration: 3.3TB usable capacity, 32GB of cache, and 32 ESCON channels. Of the 32 ESCON channels, 12 were used in this test. The ESS was configured with the PAV option for z/OS servers (implemented with the minimum benefit of two concurrent I/Os per disk volume). EMC and HDS were both requested to install PAV on their systems, but neither vendor was able to do so.

The operating environment was OS/390 R2.10, IMS V6 and DB2 V6. The test database configuration consisted of approximately 400 IMS databases and 500 DB2 databases. I/O rates were approximately 4000 I/Os per second. The read:write ratio was approximately 2.3:1 (i.e., 30% writes).

Performance Proof Point #5

Bottom Line:

The ESS F model outperformed both an EMC Symmetrix 8730 and HDS 9960 in a remote copy (remote mirroring for disaster recovery) online transaction environment.

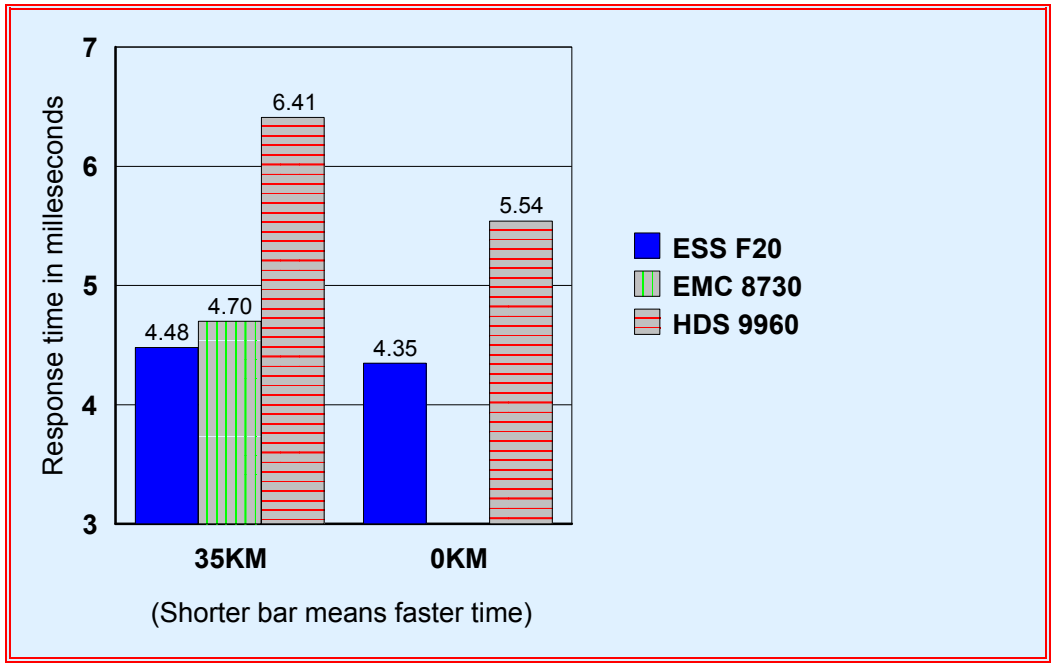

Background:

The tests were run by a North American customer in the finance industry in late 2001.

Each disk system vendor was requested to provide the following configuration: 3.3TB usable capacity, 32GB of cache, and 32 ESCON channels. Of the 32 ESCON channels, 12 were used in this test. The ESS was configured with the PAV option for z/OS servers (implemented with the minimum benefit of two concurrent I/Os per disk volume). EMC and HDS were both requested to install PAV on their systems, but neither vendor was able to do so. For the remote copy tests, each vendor configured two identical disk systems separated first by approximately 35 kilometers of cable, and then by approximately 0 kilometers (actually 20 feet) of cable.

The operating environment was OS/390 R2.10, IMS V6 and DB2 V6. The test database configuration consisted of approximately 400 IMS databases and 500 DB2 databases. Two data sharing system images were configured. I/O rates were approximately 4000 I/Os per second. The read:write ratio was 2.3:1 (i.e., 30% writes).

The remote copy facilities compared were ESS PPRC synchronous mode, Symmetrix SRDF synchronous mode, and HDS TrueCopy synchronous mode. "Synchronous" remote copy means that writes are successfully written to both the local and remote disk system before an I/O complete signal is sent to the host processor. Synchronous mode ensures data currency at the remote site. Two sets of tests were run, one for 35KM distance and one for 0KM distance. In the 0KM case, no test was run for the EMC Symmetrix due to time constraints.

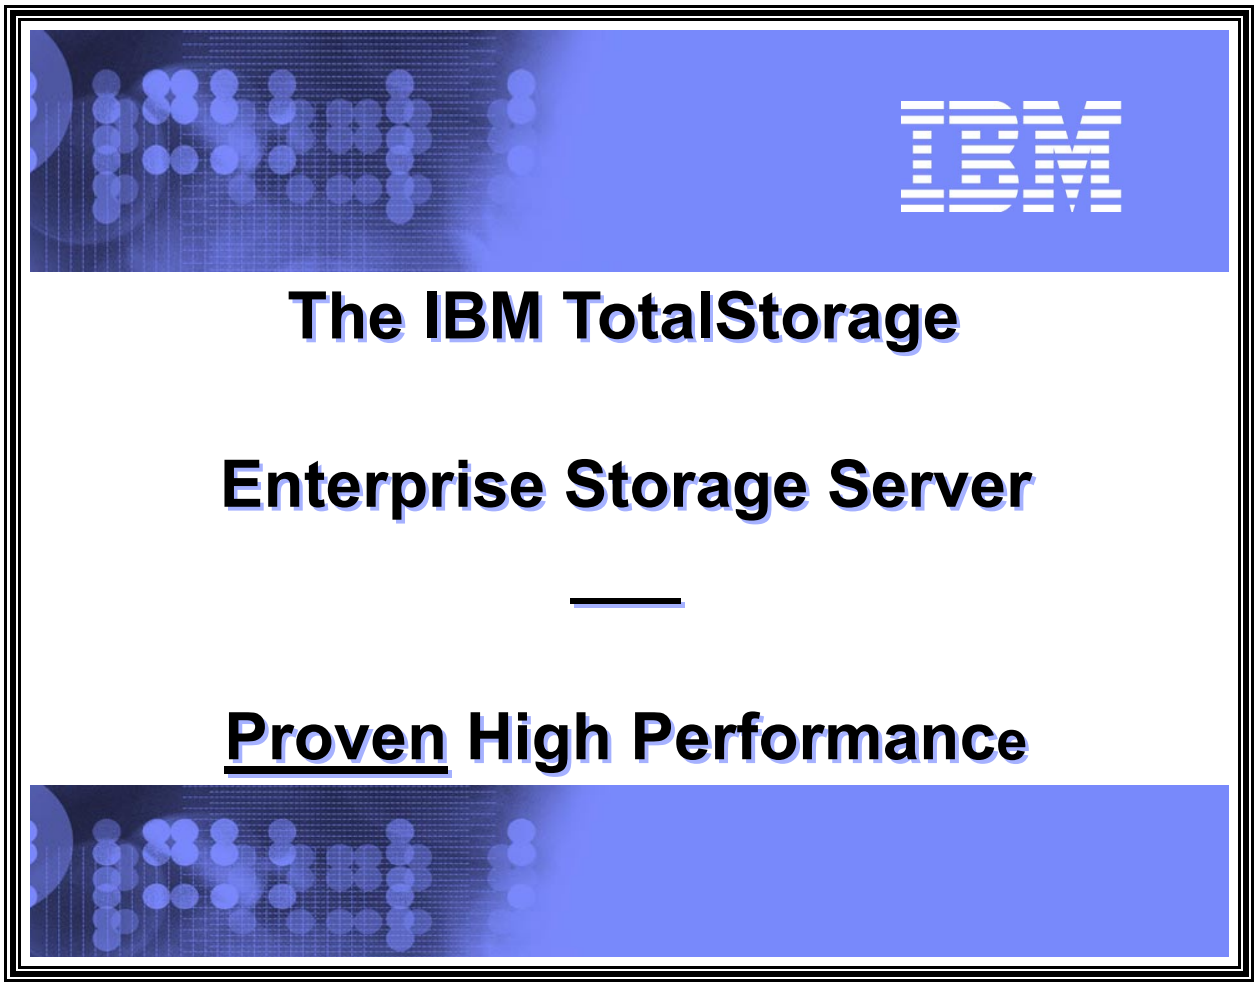
The IBM TotalStorage
Enterprise Storage Server
Proven High Performance