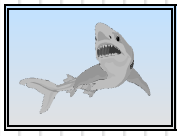

The IBM

TotalStorage Enterprise Storage Server

◘ Guide to Value and Capabilities for Business ◘

The Enterprise Storage Server (ESS) model 800, with an optional Turbo performance accelerator, is IBM's latest leadership disk system.

The ESS represents the integration of multiple advanced technologies for safely storing, quickly accessing, and easily managing data. The ESS can be described as an all-in-one disk system for the entire enterprise that stands apart from the competition in significant ways.

Delivering innovations in design, performance, and function, the ESS meets - and generally exceeds - the stringent requirements that demanding customers place on large-scale disk storage systems. Satisfying needs ranging from easier management of storage resources, to 24x365 access, to higher levels of performance and lower cost of ownership, the ESS provides demonstrable business value. With concurrent support for a broad range of server families - such as Windows servers, UNIX-based servers, and IBM zSeries (S/390) and iSeries (AS/400) servers - the ESS is truly a storage solution for the entire enterprise. And the ESS is backed by IBM quality and support.

With the ESS, online data storage is not just a peripheral anymore.

This Guide identifies specific benefits the ESS provides today that make a positive contribution to your business. Together, these comprehensive benefits make a compelling case for the ESS.

Aligning storage technology with business needs

Notices

Published June 11, 2003

Please send comments via e-mail to: djsacks@us.ibm.com

“Enterprise storage solutions such as IBM’s Enterprise Storage Server create an information delivery platform that provide[s] the needed flexibility to adapt in times of rapid change...Through use of an enterprise storage solution, companies can more effectively utilize their existing staff and — through the intelligence of the new solutions — free their staff to focus on the strategic issues of application value rather than the mundane and repetitive issues of data availability, capacity planning, and performance management.”

-- John McArthur, International Data Corporation (IDC)

"With IBM's new Shark, we have significantly increased performance and tripled our storage capacity, while reducing our cost to manage information," said Joe Giacometti, senior vice president of Information Technology at Ahold, which operates some of the largest retail grocery store chains in the world. "We selected Shark because of its superior technology to that of the competition. This technology has allowed us to better manage our information and provides us the flexibility to address our changing business needs. We've been able to reduce the time to market for new applications because of Shark."

Contents

This Guide discusses each value area from two perspectives. *Value to you* identifies specific business/ownership benefits. *How the ESS does it* identifies associated ESS attributes and facilities that deliver those benefits through technology.

Notes:

- Not every identifiable *Value to you* item or associated *How the ESS does it* item are necessarily cited in this guide.
- All items apply to all supported processors and operating systems unless specifically qualified.
- The term "midrange" is used throughout this guide to refer to non-mainframe platforms such as Novell NetWare, UNIX-based servers, Windows, and IBM iSeries (AS/400) servers.

For additional information on IBM storage please contact your IBM storage sales representative. Or visit the IBM storage Internet site at www.ibm.com/storage. The specific home page for the Enterprise Storage Server is currently: www.ibm.com/storage/ess.

Enterprise Storage Server - Value Highlights

This table summarizes some of the most important benefits and How-the-ESS-does-it items on one page. Of course, one page cannot identify every benefit of this sophisticated disk system. Please refer to the detailed itemized lists in this paper for a more comprehensive guide to ESS business and technical values.

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 is standard

Strategy / Vision

The ESS combines and integrates multiple proven technologies into a single world-class disk system. By integrating multiple advanced technologies together, the ESS delivers capabilities not otherwise available in a single disk storage system. In this way, the ESS reduces barriers to efficient IT operations by reducing technology constraints between users and the data they need.

With its powerful architecture and IBM's plans for future enhancements, the ESS is poised to deliver value today - and over the long-term.

The ESS brings you closer than ever to the vision: any data / at any time / from any place / really fast / easily managed.

Potential value to you:

Integrated, proven technologies
State-of-the-art design
Conformance to industry standards
Positioned for future enhancements

How the ESS does it:

- Third generation Enterprise Storage Server builds on and enhances previous generation market-proven systems
- Seascape Architecture is based on the principles of universal data access, operating-system-based storage servers, and snap-in building blocks. Seascape's value is also reflected in the successful Virtual Tape Server solution from IBM. Seascape's incorporation of industry-standard server building blocks reduces the need for custom components that can raise the cost and complexity of competitive systems.
- Second generation ANSI-standard Serial Storage Architecture provides internal pipelined, unarbitrated internal paths to disk
- Integrated popular industry standards include SCSI-3, Serial Storage Architecture, Fibre Channel, ESCON, FICON, PPRC, XRC, Java, Ethernet, SNMP, TCP/IP, HTML , Secure Sockets Layer (SSL), and compliance with SNIA's SMI-S management API.
- A history of functional compatibility as new models have been delivered, providing common support for many functions including interoperability for remote copy configurations.

Comprehensive Storage Consolidation Platform

Distributed computing has often led to islands of disks and data locked inside isolated processors. This not only makes enterprise storage management more complex, but also raises costs due to the inability to share storage resources.

The ESS enables installations to replace islands of isolated disks with a smaller number of shared disk systems. Some immediate benefits include reduced operational complexity, and capacity that can be easily distributed to processors that need it. Moreover, while many midrange systems' isolated disk storage is historically low-function, nearly all of the ESS's sophisticated capabilities are accessible to all attached servers, improving the overall quality of service the I/S organization delivers to its customers.

Drawing on IBM's years of experience in attaching storage systems to both IBM and non-IBM platforms, a single ESS supports concurrent attachment to a wide range of hosts. Multiple hosts can be connected to a single ESS at the same time - with the ability to grow this number into the hundreds in large-scale SANs.

Potential value to you:

Improved service to internal users and external customers
Ability to shift disk capacity to the host that needs it
Better control of assets/resources
Centralized management
Reduced operational complexity
Reduced staff/skills to manage storage
Enhanced data security
Shared (pooled) resources
Shared functions / common procedures

How the ESS does it:

- A wide variety of servers and operating systems can concurrently connect to the same ESS, and even share the same ESS Fibre Channel host port
- Flexible combinations of SCSI, Fibre Channel, ESCON, and FICON processor attachments
- Support for file sharing protocols (CIFS, FTP, HTTP, NetWare, NFS) via the IBM TotalStorage Network Attached Storage 300G appliance
- iSCSI support via a Cisco Storage Router
- Scalable growth to double-digit terabytes of RAID-protected, usable space, in one system / one footprint
- Availability and performance attributes make consolidation practical
- Capacity is a pool that is flexibly partitioned for distribution to attached hosts
- Capacity can be reassigned by the customer among attached hosts while the ESS is online
- Capacity is customizable as logical volumes of flexible sizes to fit the needs of different servers
- Logical volumes can be assigned to multiple hosts/paths for server clusters and data sharing
- Internal data replication (FlashCopy) supports all host environments
- Remote mirroring (Peer-to-Peer Remote Copy) supports all processor environments
- Dual-active internal clusters, multiple RAID arrays, and separate internal paths to groups of disks allow workload separation to help control internal resource contention
- Single (Web-based) management interface to one or multiple ESSs
- Multiple SCSI-attached servers can optionally be daisy-chained to a single ESS SCSI port
- Fibre Channel, ESCON, FICON, the 300G, and iSCSI support flexible distances between hosts and the ESS
- Fibre Channel and FICON allow a reduced number of host connections compared to SCSI and ESCON

Outstanding Performance

It is external performance that ultimately delivers value, not "under-the-covers" measures such as the bandwidths of various components in isolation. The two primary dimensions of external performance are: response time and throughput. In simple terms, *how fast* and *how many*. The time it takes to perform I/O operations is often the single largest contributor to end-user response time, particularly in commercial environments. Thus reducing I/O response time can make a significant difference to the ways users experience "the system".

The ESS incorporates sophisticated self-optimizing technology throughout that is designed to deliver high levels of external performance to attached host platforms.

Potential value to you:

Faster and more consistent 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 bottlenecks
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 does it:

- 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 down 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) allows optimized 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

(Item list continued on next page)

Outstanding Performance (page 2 of 2)

- Data (file) transfer between IBM z/OS (OS/390) and selected midrange hosts can be performed at channel speeds and offloaded from both the network and disk system, using the Parallel Data Mover product from Alebra (www.alebra.com)
- 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 paths to cache
- Remote mirroring performance optimizations and choice 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.

Round-the-clock Access to Data

Stored data can be converted to useful information only when applications can access it. Therefore, some of the most significant values of a disk system are its facilities that contribute to continuous data availability.

The ESS incorporates a fault-tolerant design that avoids single points of failure, meaning that the failure of an individual hardware component should not prevent applications from accessing data. ESS fault tolerance begins with support for redundant connections to attached processors, extends to two copies of data written to the system maintained in two separate caches, includes storing data on RAID-protected disks, and more.

But the ESS's comprehensive availability design goes well beyond such self-healing facilities. The ESS is designed to let applications continue to access data while a failed component is being replaced. Additional disk capacity can be added dynamically. Software running on the internal, redundant processor clusters can be changed while applications continue to run. Further, innovative intra-system data copy facilities can reduce application outages for making copies of data for backups or other purposes from hours to seconds -- or can even eliminate such outages altogether. And, inter-system (remote) copy facilities support continuing business operations even across extended planned or unplanned outages.

Potential value to you:

Consistent service to internal users and external customers
Stable/reliable/dependable storage environment
Greater assurance of meeting production schedules
Protection against both planned and unplanned outages
More practical to consolidate distributed storage into a single disk system
Comprehensive protection against loss of data or loss of data integrity
Support for 24x365 business operations

How the ESS does it:

- Data integrity is protected in multiple ways as it moves through an ESS. Examples: ECC (error correction) for data in cache and data on disks, Longitudinal Redundancy Checking (LRC: additional error checking) for data as it moves through the ESS, the use of dual write (mirrored) caches to protect against cache failures, and more.
- Predictive Failure Analysis can predict/prevent disk failures before they occur. Examples: internal periodic measurements of signal quality and head flying height.
- Periodic cache and disk "scrubbing" provide detection/correction of selected bit errors , helping to keep such errors from accumulating beyond the ability to be corrected.
- FlashCopy data replication minimizes or even eliminates application downtime for backups, data transfers, creating test copies, etc.
- A self-optimizing design reduces the need to take data offline for performance tuning
- The ESS design supports online upgrade of Licensed Internal Code (internal software or microcode) for host adapters, the internal Seascape servers, and internal disk drives
- Multiple, built-in spare disks support rapid, automated restoration of RAID protection after a predicted or actual disk failure

(Item list continued on next page)

Round-the-clock Access to Data (page 2 of 2)

- Internal batteries insure clean system shutdown in case of total loss of external power; if power loss is only transient, the ESS continues with normal operations.
- Additional, specialized batteries protect write data in cache, but not yet destaged to disk, for up to 72 hours. This provides additional protection such as if the primary batteries do not have sufficient time to recharge following multiple intermittent power outages during a short time interval.
- Redundant AC and redundant DC power supplies
- Redundant power cords
- Redundant cooling fans
- Redundant, dual active processor clusters with automatic failover/failback
- Redundancy for data on disk via RAID protection, ensuring data remains accessible even if a disk drive fails. (RAID-5 and RAID-10 arrays can be intermixed within the ESS.)
- Redundant internal adapters and data paths for disk modules
- Repair/replace actions ("hot swap") performed while the ESS remains online
- Upgrades performed while the ESS remains online. (Examples: adding new disks to increase capacity, or adding new host attachment ports, or increasing cache size.)
- Internal hard disks (one per cluster) hold both the current and previous level of internal software, for quick fallback if necessary
- FlashCopy consistency groups support the ability to make a multivolume copy without stopping applications
- Remote mirroring (Peer-to-Peer Remote Copy) provides for continuing business I/S operations without data loss across an extended planned or unplanned outage
- Specialized data integrity protection for remote mirroring (Peer-to-Peer Remote Copy and Extended Remote Copy)
- Asynchronous remote copy (PPRC-XD), optimized for controlled (point-in-time) data transfer operations over unlimited distance with minimal impact to application performance
- Cascading remote copy (a.k.a. multi-hop). A 3-site configuration is designed to support data currency at any distance, and a 2-site configuration is designed to support asynchronous point-in-time copies at any distance without needing to quiesce applications.

Platform specific items:

- Midrange hosts:
  - Host path I/O load balancing. The Subsystem Device Driver, a standard host-based facility shipped with the ESS, manages I/Os over multiple paths to the same volumes. In case of loss of a path, I/Os are automatically routed over other paths. 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.)

- IBM zSeries (S/390):
  - Changes to parallel I/O processing can be made while the ESS remains online; for example, Parallel Access Volume (PAV) aliases can be added dynamically
  - Concurrent Copy minimizes application down time during copy and dump operations
  - Extended Remote Copy (XRC) provides asynchronous remote mirroring for optimized application performance in a remote copy solution
  - Peer-to-Peer Remote Copy and Extended Remote Copy both fully support Geographically Dispersed Parallel Sysplex
  - GDPS/PPRC supports HyperSwap to reduce application outage time when switching to remote ESSs
  - GDPS/PPRC supports open system volumes (LUNs) as well as z/OS volumes

Ease-of-Management / Operations / Maintenance

A high-capacity high-function disk system is generally mission-critical in day-to-day business operations. A disk system needs to be easily and efficiently manageable to allow a business to fully realize the benefits of its technology.

The ESS pursues this goal from multiple perspectives. As a single, large-scale storage system, the ESS provides a single point of management control. The ESS supports both interactive and programmable management interfaces. A thoughtful maintenance philosophy minimizes customer awareness of maintenance activity. The ESS can be customized for the installation environment; for example, ESS capacity is defined by the customer as logical volumes of potentially different sizes, independent of the capacities of the physical disks installed.

Potential value to you:

Simplified operations / reduced staff
Any-time / any-location view and control
Secure access to management facilities
Support of enterprise-wide systems management solutions
System maintenance performed without customer intervention
Logical configuration tailorable to the I/S environment
Easy migration of data from conventional disk systems to an ESS
Centralized backup of midrange data to z/OS (OS/390) tape
Indirect improvement of access to other I/S resources

How the ESS does it:

- The ESS Specialist provides a convenient, interactive Web-based management interface. Access is secured by password, SSL (Secure Sockets Layer) protocol, and/or private network
- A Command Line Interface (CLI) can be used to manage system configuration and Copy Services. This can aid in automating management of the ESS.
- A SNIA SMI-S compliant Application Programming Interface (API) can manage many ESS facilities
- The TotalStorage Expert reports on the ESS internal (system) view of performance and capacity
- Through IBM's TotalStorage Proven program, the ESS has been pretested with multiple leading industry servers and applications (see www.ibm.com/totalstorage/proven). The ESS has also been certified by some vendor programs; for example, ESS Copy Services (FlashCopy) is identified by Oracle as an “Oracle Compatible Snapshot Technology”, and ESS Copy Services (PPRC) is identified by Oracle as an "Oracle Compatible Remote Mirroring Technology".
- E-mail and pager notifications of system events can be automatically sent to designated personnel
- SNMP messages of system events can be sent to designated addresses such as systems management applications
- Spare disks support automated restoration of RAID protection after a disk failure
- Data Copy Services can be combined (example: you can make a FlashCopy of a remote-copied volume for input to backup processes)
- FlashCopy: using the NOCOPY (copy-on-first-write) optimization, target volumes can be reused for other purposes without losing “move changed data only” performance optimizations for the next copy cycle
- FlashCopy: up to 12 copies can be maintained concurrently for a given source volume
- FlashCopy: automatically invoked by selected Tivoli Storage Management products
- FlashCopy: target volumes are specified dynamically by the customer; dedicated volumes are not required
- FlashCopy: tape drive use for disk backups can be reduced or rescheduled. (Once the FlashCopy is made, a process that may take no more than a second or so, production applications can continue without waiting for disk data to first be written to tape. Further, multiple backups can potentially be batched onto one tape reducing tape cartridge requirements.)

(Item list continued on next page)

Ease-of-Management / Operations / Maintenance (page 2 of 2)

- Reduced manual tuning effort due to designed-in self-optimizing facilities such as system-wide striping
- Flexible intermix of RAID-5 and RAID-10 arrays, under customer control
- Logical volumes: up to 8192 can be created in one ESS (4096 for midrange servers plus 4096 more for IBM zSeries (S/390) servers)
- Logical volumes are configurable over a wide range of capacities (from as small as 100MB to the capacity of an entire RAID array, in 100MB increments).
- "Phone home" automatically notifies the IBM Support Center if the ESS needs service
- Proactive, periodic "phone home" on schedule confirms operational status of the ESS to the IBM Support Center
- IBM product specialists can run system diagnostics from a remote support center and download any needed software changes; access to the ESS is password-protected.
- High performance, high scalability, high availability, and added-value function attributes make it practical to consolidate multiple disk systems into a single ESS, simplifying storage operations and management
- The ESS may be a potential tape replacement due to its large capacities and other benefits (such as RAID protection and elimination of manual and robotic delays)
- The IBM SAN Volume Controller can add storage virtualization capabilities to a SAN

Platform Specific Items:

- Midrange hosts: Data migration to an ESS can be via standard operating system commands or host-based mirroring facilities. Optional IBM services are also available.

- z/OS (OS/390) ⇔ Midrange file transfer: The Parallel Data Mover product from Alebra (www.alebra.com) provides high-speed data transfer between selected midrange and z/OS servers. The implementation is disk system-independent. Data transfer overhead is offloaded from both the network and the disk system. One-step transfer of database data is supported, eliminating the need for intermediate temporary flat files required by conventional data transfer facilities.
- z/OS (OS/390) ⇔ Midrange backup/restore: The Parallel Data Mover product from Alebra (www.alebra.com) provides high-speed backup of midrange data to z/OS tapes/libraries, offloading the network. The Tivoli Storage Manager (TSM) exploits this high-speed solution.

- IBM zSeries (S/390):
  - FlashCopy is detected and automatically invoked by the standard DFSMSdss copy utility - no JCL changes are needed. Both volume and data set copy are supported.
  - RMF and the IDCAMS LISTDATA command report ESS performance statistics
  - Multiple data migration methods are available: Examples: IBM services (which offers a nondisruptive migration solution), or using the ESS as an XRC target from a source disk system that supports XRC.
  - Because of the ESS's high performance design augmented by zSeries accelerators such as PAV, it is practical to define logical volumes as large capacity 3390 devices with up to 32K cylinders, reducing the total number of volumes to manage, reducing processor memory requirements, and conserving disk addresses.
  - Parallel Access Volume (PAV) tuning can be automated by the z/OS (OS/390) Workload Manager; support for up to 255 PAV aliases per volume can help address peak I/O loads and reduce manual tuning
  - P/DAS (Peer-to-Peer/Dynamic Address Switching) provides nondisruptive volume movement
  - GDPS/PPRC HyperSwap provides ease-of-operations benefits beyond P/DAS
  - Control Unit Initiated Reconfiguration (CUIR) reduces operator interactions when the ESS is serviced

Storage Area Network (SAN) Support

Storage Area Networks (SANs) hold the promise of significantly improving the overall storage environment - particularly in those installations with multiple heterogeneous midrange processors and storage systems.

IBM's Enterprise SAN initiatives include connectivity/components, management, applications (e.g., capacity pooling, LAN-free and server-free backup), and service offerings. The ESS, as a disk system, participates as a SAN component that connects to Fibre Channel-based SANs which can include hubs, directors, and switches. This ability offers immediate benefits such as increased distance, performance, and wider access to storage. As SAN applications evolve, additional benefits will likely accrue.

ESS storage networking protocol support extends beyond SANs alone. The ESS can also attach to iSCSI (Ethernet)-based SANs. And, the ESS can provide true file sharing through support of NFS, CIFS, and other file-I/O protocols.

Please visit the IBM SAN Web site at www.ibm.com/san for details on the IBM Enterprise SAN.

Potential value to you:

Separates storage from processor dependencies
Offloads I/O traffic from the LAN/WAN network to a storage-optimized network
Supports increased connectivity of many servers to one storage system
Storage capacity and function shareable by multiple processors
Positions installations to support emerging SAN-based applications
Supports file sharing and block-I/O in one disk system

How the ESS does it:

- Flexible intermix of Fibre Channel , UltraSCSI, FICON, and ESCON host connections. Fibre Channel and FICON support 1Gbit/s and 2Gbit/s speeds.
- The ESS supports a selection of popular hubs, switches, and directors for open interoperability, and to provide for extended distances between the ESS and attached servers.
- The ESS can attach to an iSCSI-based SAN via a Cisco Storage Router.
- The IBM TotalStorage Network Attached Storage 300G appliance (i.e., a NAS "head" or gateway) can be connected to an ESS to provide true file sharing for applications using standard CIFS and NFS file sharing protocols. FTP, HTTP, and NetWare protocols are also supported.
- A single ESS can be used at the same time for both file-I/O (e.g., NAS) and block-I/O (e.g., Fibre Channel and iSCSI SAN) applications.
- IBM's commitment to open SANs, demonstrated by its SAN Interoperability Lab
- IBM actively supports SNIA, the Storage Network Industry Association at www.snia.org). IBM participates in all registered multi-vendor "Supported Solutions" as of June 2003.
- The ESS is positioned to support SAN applications such as LAN-less and Server-less data movement.
- The ESS provides built-in LUN masking support
- The Tivoli SANergy product provides a software-based true file sharing solution that is supported by the ESS
- IBM Global Services can assist in SAN planning, design, implementation, and testing
- The IBM SAN Volume Controller can add storage virtualization capabilities to an ESS in a SAN

Attractive Cost of Ownership

In the dynamic storage marketplace, customers are increasingly concerned that they may unnecessarily pay more than they need to for an enterprise-class disk system and added-value features. They are also often concerned that any investment they make may too soon be obsolete. The ESS is designed to address these concerns in multiple ways.

For example, the newer ESS design can lower costs compared to conventional disk systems with older designs. Several ESS advanced features are optional with pricing based on system capacity. And, ESS capacity can be configured in multiple ways to meet individual price/performance requirements.

The ESS is also designed to incorporate "releases" of enhancements over time, so when possible newer software-based function can be installed on current models and selected previous ESS models. Similarly, hardware enhancements such as newer disk drives are sometimes made available on both current and selected previous ESS models.

Potential value to you:

Attractive upfront costs
Attractive ongoing costs
Attractive upgrade costs
Support for budget constraints
Ability to affordably acquire new technology
Investment justification
Investment protection
Long asset life
Potential elongation of processor life
Minimized technical/business risk

How the ESS does it:

- The TotalStorage ESS Specialist and associated Command Line Interface and Application Programming Interface facilities to control the ESS configuration and manage many of its functions, are standard at no extra charge
- A three year (24x7) hardware and software warranty (for priced software features) is standard
- Design efficiencies minimize upfront hardware costs. (Examples: Efficient cache algorithms and fast internal data paths mean less cache may be required to achieve higher performance than in competitive disk systems. A high-performance space-efficient RAID-5 design can lower cost while maximizing scalability, eliminating the need for a more costly RAID-1 (mirrored) configuration.)
- All supported hosts may be concurrently attached to an ESS without additional charge
- Flexible configuration and upgrade options (including disk capacities, RAID types, host connections, cache size, and optional features) allow you to buy what you need when you need it
- The ESS 800 model can be upgraded concurrently to add the Turbo performance accelerator feature
- Optional added-value functions are incrementally priced based on system capacity. (With many ESS software functions standard, there may be a lower software "delta" cost in ESS due to adding disk capacity that enters a new capacity tier, compared to competitors that have tiered-pricing for most or all system software.)
- With the its high-performance RAID-5 design, less raw capacity is likely needed in an ESS compared to mirrored systems to deliver the same usable capacity. This means that software that is priced based on raw capacity may have lower costs in an ESS than in competitive mirrored systems.
- ESS software (Licensed Internal Code) and purchased features can be transferred with the ESS to a new customer, potentially increasing residual value compared to competitive systems that do not permit such license transfers

(Item list continued on next page)

Attractive Cost of Ownership (Page 2 of 2)

- Multipath software is standard (Subsystem Device Driver). (Some competitive vendors charge for equivalent software on a per server basis so that adding a new server incurs additional disk system software charges)
- FlashCopy: target volumes are specified dynamically under user control; pre-purchased, dedicated volumes are not required
- Most ESS functions and facilities are shareable across all server platforms (the main exceptions are IBM zSeries-only features such as PAV)
- High ESS performance may extend processor life / allow processor upgrades to be delayed. (The simple reason: processor response time degrades at higher utilizations - this may be offset by the speed of the ESS.)
- Independent scalablity of the number of disks, the number of host connections, and cache size
- Quality IBM technology, service, and support reduces business risks
- An ESS is flexibly redeployable across different servers in the enterprise, increasing usable asset life
- Scalability / consolidation facilitate cost savings (compared to buying additional disk systems)
- Storage capacity is reassignable across hosts under customer control, while the ESS remains online. This maximizes the use of existing resources without vendor involvement or billable services.
- Two ESS footprints, base frame + optional expansion frame, scale floor space with capacity
- Disk capacities and RAID types can be intermixed within the same ESS, helping you to optimize price/performance/capacity
- The IBM TotalStorage NAS 300G adds NFS/CIFS/FTP/HTTP/Netware file-I/O support to the standard block-I/O access for the same ESS. A Cisco Storage Router adds iSCSI support to an ESS.
- The same host adapter supports either Fibre Channel or FICON (as specified by the customer)
- The Standby Capacity on Demand offering can allow customers to quickly access pre-installed extra capacity

Platform Specific Items:

- IBM zSeries: The following performance accelerators are standard at no extra charge (Priority I/O Queuing, Multiple Allegiance, and support for new optimized I/O commands)

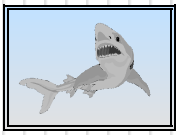

The IBM

TotalStorage Enterprise Storage Server

◘ Guide to Value and Capabilities for Business ◘

Aligning storage technology with business needs