

# IBM TotalStorage Enterprise Storage Server (ESS)

ESS provides:

- ✓ *Ultra High Availability*
- ✓ *Massive Scalability 420GB to 22.4TB*
- ✓ *High Performance with I/O Rates Exceeding 45,700 Per Second*
- ✓ *SAN Ready*
- ✓ *Advanced Copy Features*
- ✓ *Heterogeneous Support for NT, Novell NetWare, Win2000, UNIX, and IBM eServer(AS/400, RS/6000, System/390, Intel) Data Storage*
- ✓ *Business Continuation Options*
- ✓ *Investment Protection*
- ✓ *Flexible IBM Lease Terms*
- ✓ *Price Competitive*
- ✓ *Non Disruptive Upgrade Options*
- ✓ *RAID-5 or Non RAID*
- ✓ *Automatic I/O Load Balancing*
- ✓ *Synergy with IBM Tape Devices*
- ✓ *SCSI, Fibre Channel, ESCON, FICON*
- ✓ *FICON channel performance can run over two times faster than ESCON*
- ✓ *Full Hardware Redundancy*
- ✓ *Web Management Tools*

## The Storage Server Standard For The New Millennium

E-business is driving a data explosion, generating exponential growth in the need for immediately accessible, always available and highly functional storage capacity. In the increasingly e-centric world, information demand is like electricity: you plug in, you get it. Appliance-like intelligent storage served by its own network is required to provide "information as a utility."

The IBM Enterprise Storage Server(ESS) is the ultimate Storage Area Network (SAN) utility, providing the information "fuel" that runs the e-business "engine". Extensive heterogeneous server connectivity makes the ESS a natural fit for server consolidation requirements. The ESS supports rapid universal access to vast quantities of data through many advanced functions and features, making it a workhorse for support of business intelligence and other business-critical applications. Because of its enterprise-wide support and management scope, the ESS is tailor made to help provide consistent, efficient and effective enterprise resource planning. Using the ESS to address any or all of your strategic and tactical business initiatives will give your organization the business advantage needed to survive and thrive in the e-world.

The ESS is the third generation of the IBM Seascope (TM) architecture for disk systems. It is a solution that provides the outboard intelligence required by SAN solutions, offloading key functions from host servers which frees up valuable processing power for applications. As a comprehensive SAN-based storage solution, the ESS provides considerable management flexibility to meet the fast-paced requirements of today's businesses.



## Universal Data Access and Future Enhancements

ESS allows you to consolidate your data from different platforms on a single high performance, high availability storage server.

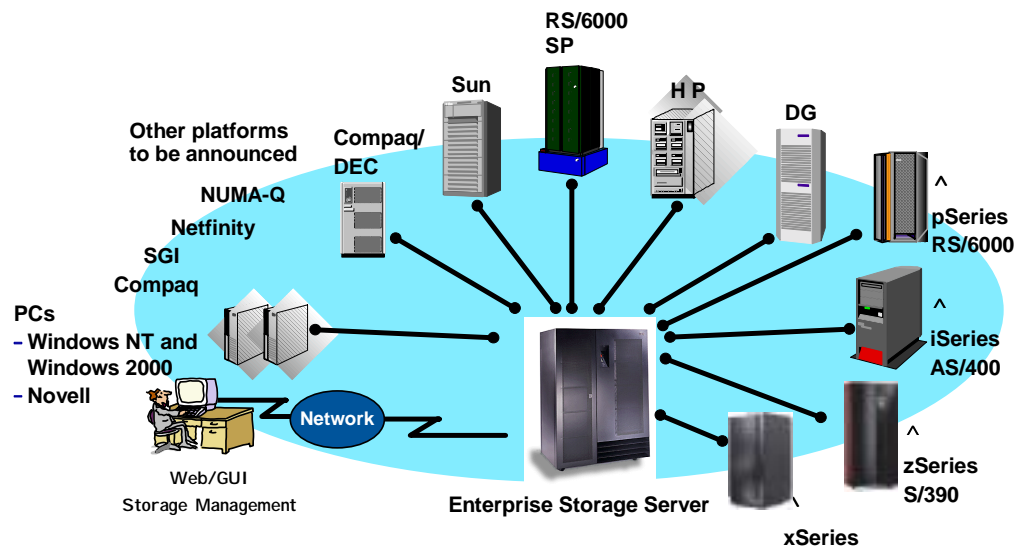
For example, 50 NT, and 5 UNIX servers could use Fibre Channel connected to the ESS.

The integrated Web based StorWatch Specialist is used to configure the ESS subsystem.

The optional StorWatch ESS Expert product provides asset, capacity and performance management functions which are especially useful for open systems. The Expert provides functions such as number of I/O, bytes transferred, read and write response times, etc.

With a capacity of up to 22TB and up to 32 direct host connections (more with daisy-chaining and switches), an ESS can meet both your high capacity requirements and your performance expectations.

Most upgrades such as additional capacity are non disruptive.



The ESS supports many diverse platforms including the IBM's pSeries (RS/6000), many leading UNIX variants (Sun Solaris, HP-UX), IBM xSeries (NetFinity®) and other Intel-based PC servers (Compaq, etc.) running Windows NT/2000, Novell Netware, iSeries (AS/400®) running OS/400® and zSeries (System/390). In addition, the ESS supports System/390® servers running OS/390®, VM/ESA®, VSE/ESA® and TPF®. This rich support of server platforms is not limited in any way; any combination of these heterogeneous platforms may be used with the ESS. Storage capacity is partitioned among the attached servers using the flexible, web based StorWatch Enterprise Storage Specialist management tool. The ESSNet feature even provides a browser for the StorWatch interface!

Different platform types use different connection protocols, and the ESS is equipped to handle many different types concurrently including connectivity for UNIX, Intel-based PC and iSeries servers, Fiber Channel, and ESCON or FICON for zSeries servers. The ESS has up to 16 Fibre Channel or 100 MB/sec full duplex FICON ports (up to 100KM without data rate "droop") offered in pairs. Direct Fibre Channel attachment is available for connections to Windows NT/2000, Novell NetWare, and UNIX servers from IBM, Sun Microsystems and Hewlett-Packard.

The ESS also supports Storage Capacity Upgrade on Demand which allows you to get the capacity your e-business needs when it needs it without disrupting the flow of work. Storage Capacity Upgrade on Demand ESS, A Step Ahead, allows you to have extra disk storage capacity available when you need it. It gives you the option of receiving an extra amount of disk storage which is shipped and installed in the ESS, but which will be activated and paid for at a later time. This should be particularly attractive for customers with dynamic or unpredictable growth, or those who just want the assurance that the storage will be there when needed. Technical details about the ESS can be found at the IBM ESS web site: <http://www.storage.ibm.com/hardsoft/products/ess/ess.htm>

## Business Value of the ESS

In today's information-dependent world of e-business, the ability to gather and store information from every imaginable source and deliver it to users across the extended enterprise is the key to business success. The ESS provides:

- Extensive Connectivity - The ESS provides simultaneous attachment to a wide variety of host servers with storage capacity partitioned among the attached servers using flexible web-based StorWatch management tools.
- Copy Services Functions - The ESS supports several hardware copy functions for two purposes: Mirroring for disaster recovery solutions, and copy functions that provide a near instant copy of data.
- Comprehensive Availability for Applications and e-Business 24x7x365 - Everything is designed for redundancy: power, cooling, adapters, buses and microprocessor complexes. Redundancy of all components means that ESS has been designed with virtually no single points of failure as to hardware components.
- Breakthrough Performance - The ESS hardware can be configured to provide the capacity and performance required for applications supporting over 41,000 I/Os per second. New OS/390 options allow multiple simultaneous I/O operations to the same logical volume, which significantly reduces IOSQ time.
- Low Total Cost of Ownership - Snap-in-building blocks keeps IBM development costs low -- with the savings passed on to you. IBM Global Financing can develop a lease to fit every requirement. Three-year warranty is included.
- FICON Performance - FICON System/390 channels, can reduce tape backup elapsed times by up to half, provide 4:1 reduction in ESCON channels, and can run over two times faster than ESCON channels
- Extreme Scalability - The ESS adapts to growing storage requirements created by e-business applications scaling from 420GB to over 22TB with non-disruptive upgrades. Also, storage consolidation from multiple platforms reduces costs.
- E-Business - When you're an e-business, the store is always open, so you need data storage that's available 24xforever. Your customers drive your workload, so you need high performance data storage that can meet unexpected demands. Customers can flock to your door without warning, so you need storage that is scaleable. Suddenly, everyone in the world is your user, so you need data storage on which you can bet your business!
- Storage Area Networks (SAN) - The ESS offers customers the leading storage subsystem, supporting the key servers in the industry, and enabling customers to realize the promise of nonproprietary SANs.
- Remote Hardware Support - The IBM hardware support center call home connection results in virtually no user intervention required on hardware failures.

*SCSI, FICON, Fibre Channel, and ESCON provide robust host server connectivity.*

*PPRC, XRC, Concurrent copy and FlashCopy provide data copy functions.*

*Maintenance and upgrades normally can be performed concurrently with full operation, including the concurrent activation of microcode.*

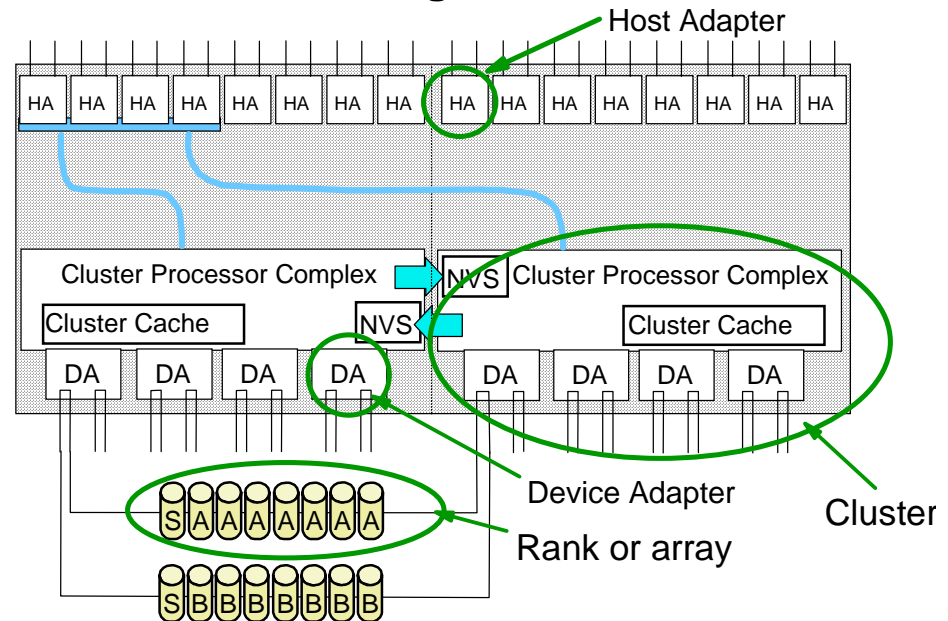
*ESS "future-proofs" users by providing a new "snap-in" design. With ESS, you can easily add performance, capacity and connectivity upgrades, providing investment protection and 24x7 productivity, saving money and time.*

*ESS has one of the smallest footprints in the industry for 22TB*

## Architecture Overview

The schematic below illustrates the major components of the ESS. Starting at the top, there are 16 Host Adapters (HAs) which are the external interfaces. Each adapter supports two ports, either SCSI or ESCON. Fibre Channel or FICON HA's will support one channel per HA. Each HA connects to *both* Cluster Processor Complexes so that either cluster can handle I/Os from any host adapter.


### Overview of ESS Logical Structure



*ESS has full redundancy of hardware components.*

*Designed to have the highest quality ever for an IBM disk storage system.*

*ESS requires two power cords for redundancy.*



*Some vendors have only one copy of data in cache -- if the cache fails data is lost!*

*If a cluster fails there is no loss of data access since the other cluster has a Disk Adapter that can access all the data. Write data is retrieved from NVS.*

There are two Cluster Processor Complexes in the ESS, each working independently. Each contains four RISC Symmetric Multi Processors (SMP), with up to 16GB of cache (total combined up to 32GB), non volatile cache (NVS), and four device adapters. All host data is written to and from the cache using adaptive record, partial track, sequential or full track staging algorithms. In addition all write data is automatically placed into two separate caches for integrity. The NVS cache has its own 7-day battery and is managed using a least recently used (LRU) algorithm. On data writes, one copy of data is placed in one cluster and the second copy of the write data is placed in the NVS of the other cluster. In the event of a either a planned or unplanned cluster outage, write data for the failed cluster is available on the other cluster and the remaining cluster takes over the functions in progress from the failed cluster. The design philosophy is that once data is in the ESS, it is intended to be-fully protected.

Within each cluster the Device Adapters (DAs) are used to connect disks to the Cluster Processor Complexes. DAs are always installed in pairs, one in each cluster. Disks arrays, or ranks (IBM 9, 18, 32 or 72 GB disks) are attached to the two DAs via 160 MB Serial Storage Architecture (SSA) loops. The ranks can be configured as RAID 5 or non-RAID arrays (Just a Bunch of Disks -- JBOD).

ESS supports AS/400, RS/6000, NUMA-Q, Netfinity, Intel, HP, Sun and other host servers. For the complete list see:

<http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm>

ESS offers open systems users 24x7 availability with redundant design and non disruptive maintenance.

The StorWatch ESS Specialist enables centralized physical storage management which reduces cost and management complexity.

IBM 7133 storage systems capacity can be used with the ESS providing investment protection for assets already paid for.

Linux is now supported on Intel server environments using either Red Hat or SuSE Linux distributions.

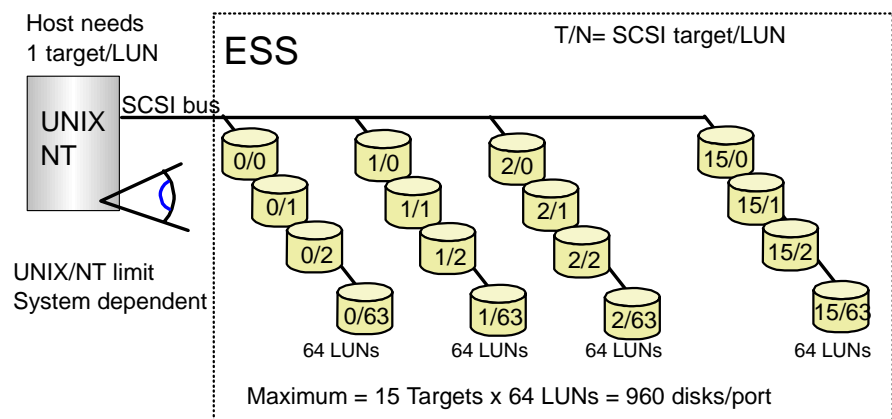
IBM Subsystem Device Driver (SDD) provides:

- More than one path from the host to the ESS.
- A single LUN can appear as 2 - 16 LUNs
- Host path failover
- Load distribution across paths
- Supports AIX 5L and Windows NT and 2000 with other UNIX server types planned for the future.

## ESS Open System Option

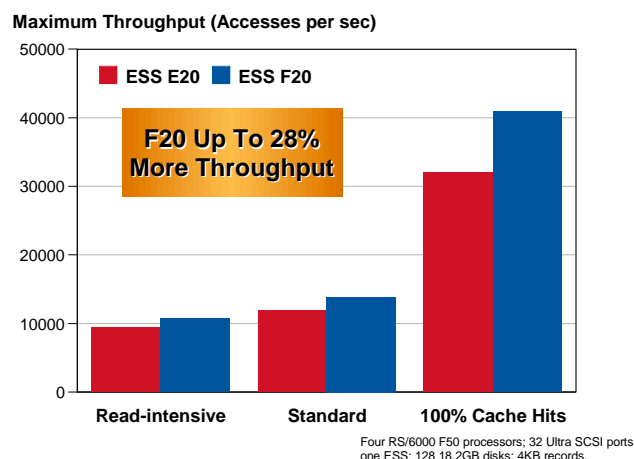
The ESS provides universal data access including: pSeries and SP2 running AIX, many leading UNIX variants, IBM xSeries and other Intel-based PC servers running Windows NT and Novell Netware, and iSeries running OS/400. This rich support of server platforms is not limited in any way; any combination of these heterogeneous platforms may be used with the ESS (including S/390). The storage capacity is partitioned among the attached servers using flexible Web-based StorWatch Enterprise Storage Specialist management tools. If you imagine a SCSI host's view of an ESS, it looks like a group of SCSI disks attached to a SCSI bus. The actual number that any UNIX or NT system can support is considerably less than the maximum shown below:

### A UNIX or NT View of an ESS



IBM performance tests have demonstrated that the ESS has significantly exceeded IBM's design goals. The example below used four pSeries processors, 32 Ultra SCSI ports, 128 18.2GB drives with 4K records achieving over 40,000 I/Os per second. No other competitor comes close.

### Open Systems Maximum Throughput



## [zSeries \(System/390\) Overview](#)

ESS supports OS/390, z/OS, VM/ESA, z/VM TPF and VSE/ESA, and Linux operating systems.

Up to sixteen 3990 control units can be defined in the ESS.

Custom Volumes may also be configured. Custom Volumes allow a volume to have any number of cylinders you choose. One would typically define a small number of cylinders to reduce contention for high activity data sets.

In most cases, however, the PAV function would be a better choice.

The combination of StorWatch Specialist and HCD / IOCP are used to configure the ESS.

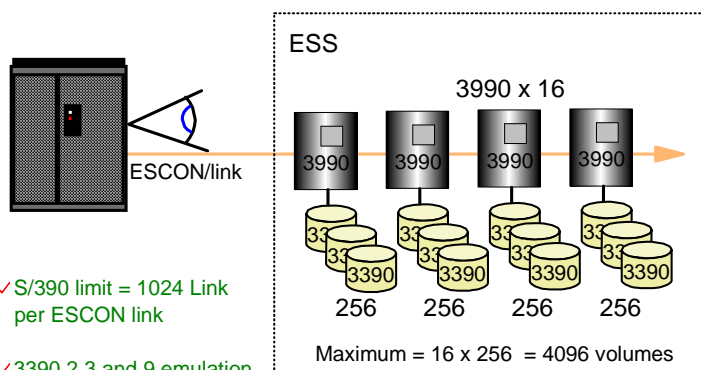
*New 3390-9 capacities are available to support up to 32,760 cylinders which relieves address constraints and allows multiple disk volumes to be consolidated.*

The chart on the right compares various storage systems from IBM to the ESS model F20.

The ESS can handle significantly more work than is shown here as more channels and capacity are added.

The ESS supports System/390 servers and inherits the S/390 functions previously provided by the 3990-6 and RAMAC families. However, the ESS far surpasses these legacy technologies in all areas such as performance, hosts supported, availability, functions, scalability and non-disruptive operation. The ESS is designed specifically to support both S/390 and open systems workloads at the same time. Typically S/390 workloads are cache friendly and take advantage of large caches, whereas the open systems workloads often are very cache unfriendly. For the S/390 workloads there are sophisticated cache management algorithms and a large cache. For workloads that do not take advantage of large cache sizes, the ESS has high performance disk arrays with 160 MB disk loops. So whatever the workload, even mixed workloads, the ESS delivers high performance. The S/390 hardware view of the ESS is shown below:

### S/390 View of Logical Subsystem



✓ S/390 limit = 1024 Link per ESCON link

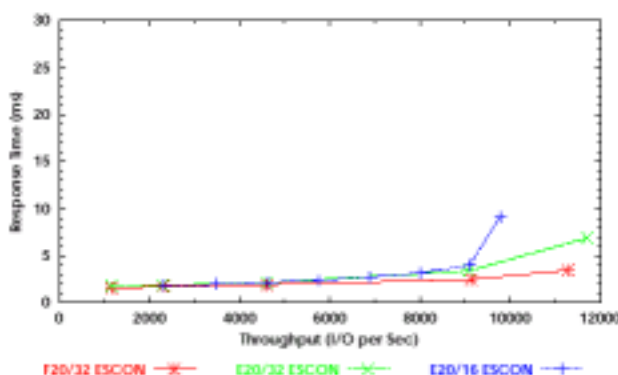
✓ 3390 2,3 and 9 emulation

✓ 3380 track format

✓ Variable size 3390 and 3380 volumes - (Custom Volumes)

The number of 3390 or 3380 devices available depends on the type of disk installed (9, 18, 36, 72 GB), how many are installed and the type of emulation you elect. With high speed microprocessors, 160 MB disk interfaces, large cache and large capacity the ESS provides extraordinarily high performance. In the example below, with a cache friendly workload, the new ESS model F20 exceeds the server's capacity to drive it.

### ESS 1.1 Cache Friendly Performance



Read hit ratio = 92%, R/W ratio = 5:1

S/390 architecture only allows one I/O operation at a time to disk devices. The new ESS changes this so that more than one I/O can be active at a time to one disk device.

OS/390 ESS PAV Support:

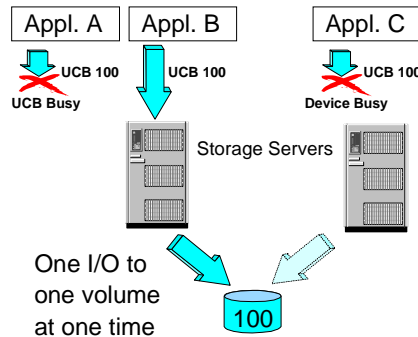
- Transparency support has PTFs to use the ESS as a 3990 only
- Exploitation static support has PTFs that recognize the ESS with PAV support and can use alias and base addresses for I/O
- The Workload Manager will automatically tune the OS/390 or z/OS I/O system using alias address by the use of policies

PAV's:

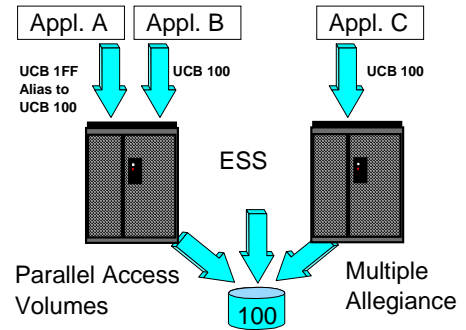
- Sharply reduces IOSQ time
- Provide multiple UCBs per logical volume
- Aliases are visible only to IOS in OS/390
- PAV support is an ESS option

Since it was developed, System/390 architecture allows only one channel program to be active to a disk volume at a time to ensure that data being accessed by one channel program cannot be altered by the activities of some other channel program. OS/390 does not try to issue another I/O to a disk volume when a previous I/O is already active for that volume. Unit control blocks (UCBs) and OS/390 queues keep track of I/Os so they are processed one after another. The ESS and OS/390 software change all this.

### Traditional OS/390 Behavior



### Enterprise Storage Server (ESS) Behavior




The ESS introduces the capability to do more than one I/O to a S/390 logical volume. The ESS introduces the concept of alias addresses. Instead of one UCB per logical volume, an OS/390 host can now use several UCBs for the same logical volume. The ESS also accepts I/Os to a shared volume coming from different hosts in parallel. This capability is called Multiple Allegiance. In the table below, the first row shows hosts A and B running separately to a DB2 volume. The second row illustrates how host B can significantly slow down host A I/O. With the ESS and Multiple Allegiance, there is no slowdown, as shown on the third row!

## Multiple Allegiance

### Concurrent access from multiple Path Groups (system images) to a volume

- Incompatible I/Os are queued in the ESS
- Compatible I/O (no extent conflict) can run in parallel
- ESS guarantees data integrity
- No special host software required, however:
- Host software changes can improve global parallelism (limit extents)



Example

### Improved system throughput

- Different Workloads have less impact on each other

	Host 1 (4K read hits)	Host 2 (32 record 4K read chains)
Max ops/sec - Isolated	767 SIOs/SEC	55.1 SIOs/sec
Max ops/sec - 100% Extent Conflicts	59.3 SIOs/SEC	54.5 SIOs/sec
Max ops/sec - full Multiple Allegiance	756 SIOs/SEC	54.5 SIOs/sec

## ESS Copy Features

Copy functions are supported and can be initiated by OS/390. UNIX systems use the StorWatch Specialist software to initiate and control copy services.

PPRC is the ideal way to provide for business continuance in the event of a major outage at one location.

With PPRC the ESS ensures that write data is in two separate ESS's before letting the application continue.

PPRC is an industry standard..

PPRC is established on a volume level and can be used by midrange and S/390 systems.

PPRC uses direct connections between ESS systems with no host involvement.

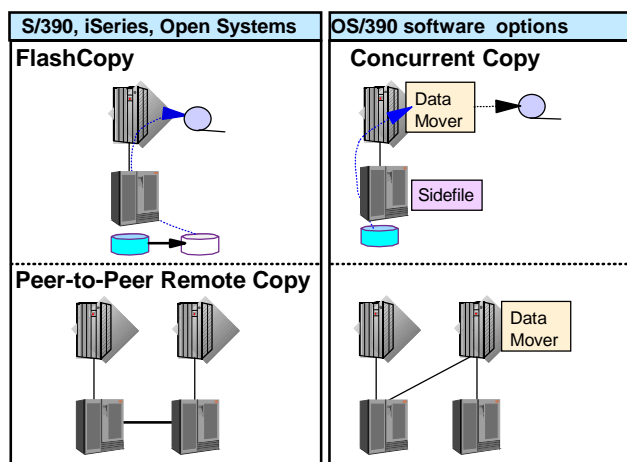
*FlashCopy is designed to make copies of volumes in less than 2 seconds!*

PPRC and FlashCopy are separate options on ESS systems.

Extended Remote Copy an asynchronous solution, was previously available on IBM 3990 control units and is for OS/390 systems only.

The ESS supports several hardware copy functions for two purposes: mirroring for disaster recovery solutions, and copy functions that provide an instant copy of data.

- FlashCopy provides an instant copy of a volume for SCSI, S/390, iSeries
- Concurrent Copy for S/390, of a volume or data set
- Peer to Peer Copy(PPRC) is synchronous mirroring for SCSI, S/390, iSeries
- Extended Remote Copy (XRC) provides asynchronous mirroring for S/390



ESS's capability to share logical volumes between like systems allows concurrent access to data bases by different hosts (both IBM DB2 Universal Database and Oracle Parallel Edition support this feature). The data base products provides locking mechanisms to guarantee data integrity.

With ESS's Peer to Peer Remote Copy (PPRC) function, S/390 and midrange systems (NT, NetWare, UNIX, iSeries) systems can now realize a disaster recovery solution to protect data should one location be destroyed. PPRC provides a synchronous copy of selected data up to a distance of 103 km. PPRC is administered from the web browser interface of IBM's StorWatch ESS Specialist for midrange systems and ICKDSF for S/390. No special software is required to use PPRC. ESCON cables connect the ESS systems. Write operations are completed on both copies (local and remote) before they are reported complete to the operating system.

FlashCopy provides the ability to create a copy of your data immediately, whenever you need it. FlashCopy provides an instant copy of a logical disk. For Example FlashCopy could be used to enhance asynchronous backup, creating test data from production data, checkpoints, and for non-disruptive copying for business intelligence applications. The copy of the data is available immediately after the command is issued.



## Business Solutions

### E-Business

*The ESS helps solve today's key business problems such as:*



*Completing a transaction on the Web with unpredictable utilization.*


e-business implements electronic end-to-end supply chain management, including on-line sales and order entry, billing, inventory control and shipping. It is truly a better way of doing business since it shortens product cycles; enriches customer relationships; broadens collaboration among employees, partners and vendors as well as customers; allows rapid response to market shifts and generally improves business results. Using e-business, commerce is no longer a function of time or location. The ESS is designed to help you grow and prosper in the e-business world. Significant capabilities in support of e-business include:

- ✓ Centralized management (using an Internet browser) of multiple storage servers supporting configuration, performance, asset and capacity management
- ✓ Exceptional scalability for fast growth
- ✓ Heterogeneous connectivity
- ✓ Very high data availability for real-time transactions
- ✓ Performance to support low response times during high periods transaction traffic
- ✓ Advanced replication services to support ancillary processes as well as backup and/or business continuance

### Enterprise Resource Planning

With enterprise resource planning, businesses are adopting a consistent approach to information throughout the entire enterprise. Included is information in all key categories, including financials, production planning, sales and distribution, office automation, human resources and many others. With consistency comes efficiency, and this is how enterprise resource planning pays big dividends. Business cycles are reduced since they are more integrated, and decision makers get the mission-critical information they need in a very timely manner. This integrated and consistent approach provides a key competitive advantage. In the enterprise resource planning arena, the ESS is the perfect information repository in areas such as:

- ✓ Facilitates the sharing of information throughout the organization
- ✓ Provides robust disaster recovery / business continuance capabilities
- ✓ Supports dynamic allocation of storage capacity to different servers
- ✓ Performs excellently in diverse environments



*Optimization for core business operations.*

- ✓ Provides rock-solid, dependable availability of mission-critical data
- ✓ Supports rapid database replication, as well as fast backup and restore

### **Business Intelligence**



*Customization, collection, analysis, storage, and management and distribution of corporate data for business intelligence.*

The harnessing of vast quantities (terabytes) of data to be used to achieve business solutions that provide a competitive advantage characterizes the business intelligence (BI) arena. The data must be aggregated from across the enterprise for analysis and interpretation so it can support business decisions. Businesses are looking for any advantage in today's fast-paced world, and they appreciate the way information can be transformed into action using BI disciplines. As a result, BI is growing exponentially. Ways in which the ESS can help you accomplish your business intelligence objectives include:

- ✓ Supports a variety of environments with heterogeneous multi-server attachments
- ✓ Provides a rich variety of replication services to support fast movement of bulk operational data between servers
- ✓ Provides very high performance for efficient data mining and to help maintain data currency in the warehouse
- ✓ Scales to your growing BI application needs
- ✓ Point in time copies for incremental and volume level copies

### **Server Consolidation**



*Data integration and centralization of storage for many different environments.*

Consolidating servers means recentralizing your IT workload from across the enterprise. Consolidating both applications and data reduces costs and complexity while increasing management efficiency. The following storage requirements for supporting server consolidation are perfectly matched to ESS's capabilities:

- ✓ Increased capacity that can be easily managed to achieve better resource utilization
- ✓ Capacity that is sharable between heterogeneous server platforms
- ✓ The ability to optimize performance for multiple mixed storage environments
- ✓ Continuous (7 x 24) data availability
- ✓ Scalability
- ✓ Investment value that extends far into the future

*ESS is optimized for all these workloads and more!*

## Enterprise Storage Server Summary


The Enterprise Storage Server is the storage of choice for those wanting the ultimate in enterprise disk storage consolidation and data sharing on multiple or heterogeneous server platforms with combined storage requirements of more than 500 GB.

The Enterprise Storage Server is the longer term strategic enterprise disk storage solution for those that want to manage all of their storage from a single location. Enterprise Storage Server is the right choice for those requiring high-performance RAID-5, read and write cache, or the flexibility of common storage for multiple servers with ESCON, FICON, SCSI, and Fibre Channel attachments.

The Enterprise Storage Server supports rapid universal access to vast quantities of data through many advanced functions and features, making it a workhorse for support of business intelligence and other business-critical applications. Because of its enterprise-wide support and management scope, ESS is tailor made to help provide consistent, efficient and effective enterprise resource planning. Using ESS to address any or all of your strategic and tactical business initiatives will give your organization the business advantage needed to survive and thrive in the e-world.

*Your Data  
is the  
Business*

**The Enterprise Storage Server is THE server for the next Millennium.**





11/2001 prepared by Kim Frank, IBM Architect 908-931-4588 KCFRANK@US.IBM.COM

Technical editing by Chris Saul, Storage Systems Division OEM Marketing Programs Manager

References in this publication to IBM products, programs, or services do not imply that IBM intends to make them available in all countries in which IBM operates.

This publication contains insights into IBM's future product development plans. These development plans are subject to change or withdrawal at any time without prior notice.

IBM reserves the right to change specifications or other product information without prior notice. All information contained in this publication is believed to be accurate as of the date of initial publication. IBM has no obligation to update any such information.

Copyright 2001, International Business Machines Corporation. All Rights Reserved.

Published performance information was modeled or obtained in a controlled environment based on the use of specific data. The results that may be obtained in other circumstances may vary significantly. IBM makes no representation and provides no assurances the same or similar performance results may be obtained. Users of performance information should verify the applicable data in their specific environments.

System/390, S/390, OS/390, CICS, IBM Parallel Sysplex DB2, MQseries, IMS, FICON, ESCON, VTAM, RS/6000, AS/400, OS/400, VM/ESA, TPF, VSE/ESA, PAV, AIX, ESS, PPRC, XRC, FlashCopy, RAMAC, DB2, StorWatch, SSA, P/DAS, iSeries, zSeries, xSeries, pSeries, z/OS, z/VM, and Netfinity are trademarks or registered trademarks of International Business Machines Corporation in the US and other countries.

Lotus, and Domino are trademarks or registered trademarks of International Business Machines Corporation in the US and other countries.

Windows and NT are trademarks or registered trademarks of Microsoft Corp. in the US and other countries.

Snapshot is a trademark and the property of Storage Technology Corporation for a duplication product.

UNIX is a trademark of The Open Group.

IBM home page [www.ibm.com](http://www.ibm.com)

IBM Storage [www.ibm.com/storage](http://www.ibm.com/storage)

GDPS [www.as.ibm.com/asww/offerings/mww62b1.htm](http://www.as.ibm.com/asww/offerings/mww62b1.htm)

IBM software [www.ibm.com/software](http://www.ibm.com/software)

IBM Servers [www.ibm.com/Servers](http://www.ibm.com/Servers)

ESS [www.ibm.com/storage/ess](http://www.ibm.com/storage/ess)