First In
Storage Solutions

IBM 7133 Serial Disk Systems

Presentation
Guide

Twelfth Edition (May, 1999)

References in this document to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM program product in this document is not intended to state or imply that only IBM's program product may be used. Any functionally equivalent program may be used instead.

While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this document was obtained in a controlled environment based on the use of specific data. The results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data in their specific environment.

The following terms are trademarks of the IBM Corporation in the United States or other countries or both: IBM, Seascape, Versatile Storage Server, AS/400, VM/ESA, ESCON, AIX, OS/400, Magstar, DFSMS/MVS, RS/6000, RS/6000 SP, ES/9000, ADSTAR, MVS/ESA, VSE/ESA, S/390, MVS, ES/3090, RISC System/6000, and S/370.

Other company, product, and service names may be trademarks or registered trademarks of their respective companies.

About This Document

This high-performance storage system provides powerful, responsive, flexible, cost-effective storage for RS/6000, RS/6000 SP, Sun, Hewlett-Packard, DEC, Compaq, and Intel-based server systems.

This presentation covers the IBM 7133 Serial Disk System and the attachment options to IBM and non-IBM Unix platforms. **Please refer to the IBM Serial Storage for Windows NT Servers (SSA4NT) presentation for information on our SSA products for Intel-based servers.**

The presentation also describes many of the Serial Storage Architecture concepts. SSA is ideal for open systems environments and provides the only industry standard interface for high-performance serial data transfer in the marketplace today. The SSA architecture provides a cost-effective way to overcome the performance bottleneck faced by computer users who currently rely on parallel interfaces like the widely used SCSI architecture.

Foil 1 - Title Page
IBM 7133 Serial Disk Systems
A Seascape Building Block

In today's "open systems" marketplace a revolution is occurring in the world of information storage. Higher levels of availability, fault tolerance, performance and connectivity are fundamental business requirements.

Imagine a world where these fundamental business requirements are addressed with cost-effective solutions - not a year or two in the future, but today!

SSA, an innovative interface technology for storage, is available today to provide unparalleled connection for a wide spectrum of storage requirements. With visionary insight, SSA was designed to meet today's customer needs while supporting tomorrow's expanding requirements.

Foil 2
IBM Serial Storage Takes the Market by Storm!

In September of 1995, the Storage Systems Division (SSD) of IBM introduced a disk storage system for the open systems marketplace based on a second generation serial interface. This product, the IBM 7133 Serial Disk System, has enjoyed extraordinary market acceptance. IBM has shipped 5 petabytes (5000 terabytes) of serial storage through March 1999 and the demand continues to be strong.

What makes the 7133 so popular with customers and business partners is the serial interface called Serial Storage Architecture (SSA). SSA is a second-generation serial architecture invented by IBM and made available to the industry. Leading vendors in the industry formed the SSA Industry Association which made SSA available as an open industry standard under the American National Standards Institute (ANSI) X 3T10.1 committee. An SSA infrastructure has been well established since 1994, with component suppliers embracing the technology and producing SSA chips, cables, connectors, adapter boards, peripherals, enclosures, and subsystems. The SSA design source was offered to the industry by IBM free of charge and has been distributed to nearly 50 companies worldwide.

The high volume of serial storage disk systems shipped to date is a clear indication that customers and business partners are enthusiastic about IBM's serial storage and the 7133. There are over 100 active customer references for serial disk systems in the Worldwide Customer Reference Database at w3.ibm.com/e-business/ibmref.nsf/. Application briefs for customers like Cornell and Phillips Petroleum are on the World Wide Web at www.ibm.com/storage/disk under Solutions.

Foil 3
Serial Storage - A Seascape Building Block

The 7133 Serial Disk System is a Seascape(TM) building block. Seascape is IBM's Storage Enterprise Architecture, a blueprint for comprehensive storage solutions, optimized for the networked world. Seascape is a comprehensive enterprise storage architecture enabling rich software function, global management, and a full range of disk, tape, and optical solutions.

The Seascape architecture has three basic principles:

1st, Universal Data Access. Seascape is focused on cutting across the boundaries and complexities of a mixed platform world, to enable the sharing of storage resources and data within the enterprise.

2nd, Powerful Storage Server. To tackle the issues of complex data movement and enable revolutionary storage solutions that are tuned to specific applications, Seascape architecture provides a powerful
integrated processor.

And 3rd, Snap-in Building Blocks. Software and hardware building blocks are the key to keeping pace with the future both in function and technology. This key principle enables superior investment protection at an affordable cost.

The 7133 is your entry into the Seascape family of products.

Now, let's look at SSA in greater detail.

Foil 4
SCSI Architecture

Before the introduction of SSA and the new, exciting IBM 7133, 7131-405, and 3527 Serial Disk Systems, customers had basically one fundamental option, SCSI, which stands for Small Computer System Interface.

SCSI refers to the original SCSI standard approved in 1986, that became the primary way in the "open systems" world to attach IBM and non-IBM storage products. SCSI involves the parallel transmission of data across a parallel set of wires. The wires also carry the addressing (limited to eight total addresses on 8-bit SCSI and 16 on 16-bit SCSI). As SCSI evolved, there have been a number of SCSI implementations (Fast, Wide, Differential) which require different SCSI adapters, cables and connectors. Increasing the speed of SCSI has involved the doubling of the data width to 16 bits and increasing the clock speed. This has increased the size of the bulky cables and connectors. In order to reliably send 16 bits of data across a distance of 25 meters, 68 wires are required for a Fast/Wide Differential ended cable. An increase in clock speed is limited by electro-magnetic interference (EMI) and data skew complications, which limits cable lengths.

The performance of SCSI disks is impacted by other devices on the string. If slower devices such as tape or CD-ROM are operating on the same string as disk, the disk performance will be decreased to match the slower devices.

SCSI offered generally lower price with reasonably acceptable performance, but has serious architectural and performance limitations due to an arbitrated, shared bus. The shared, arbitrated bus architecture limits the number of devices and distances. It also requires complex arbitration, cabling, terminators and switches.

SCSI solved yesterday's problems. Let's take a closer look at the limits of the SCSI bus architecture.

Foil 5
SCSI Bus - Limited Architecture
Shared and Arbitrated

With a SCSI bus every disk drive is connected to the same cable, therefore data can travel in only one direction at a time. This architectural limitation is common to all SCSI implementations (i.e., SCSI-2 Fast, SCSI-2 Differential Fast/Wide, Ultra SCSI, etc.).

SCSI has two serious architectural limitations:

- A shared bus
- Mandatory arbitration

In concept, SCSI is analogous to a two-lane country roadway with a one lane bridge. Two lanes must share a one lane bridge, allowing only one-way traffic on the bridge! Clearly, constant coordination and communications are required for traffic control.

Approaching traffic has to determine if the bridge is in use before attempting to cross it. So, traffic arbitrators must constantly communicate to avoid "head-on collisions", or lost data in the case of information processing.

For traffic to change direction, NO TRAFFIC can be on the bridge! This means there is a lot of waiting, the enemy of higher performance. Even when traffic is moving across the bridge in one direction, it is stopped and queuing in the other direction. For moving lots of data, a one way traffic architecture is clearly not efficient.

How does SSA's serial link architecture address these limitations?

Foil 6
SCSI Storage Evolution

The limitations of SCSI-2 are well recognized. The introduction of new technologies like high performance disk, and new application requirements like audio/video and data mining are bumping up against this limit. The ANSI standards committee is attempting to address this with a new definition of SCSI-3. SCSI-3 can be implemented through a parallel interface, sometimes known as Ultra SCSI, and three Serial interfaces. Ultra SCSI is essentially an extension of SCSI-2 parallel where the clocking speed has been doubled to provide a maximum of 40 MB/sec data rate. Its advantages are that it can attach SCSI-2 devices, provide a faster data rate, and support the enhanced SCSI-3 command set. The disadvantages are shorter cable lengths than SCSI-2 (with the exception of differential cables) and the same problems of an arbitrated, non-duplexed, bus architecture. The three serial interfaces that support SCSI-3 are FCS (Fibre Channel standard)/FC-AL (Fibre Channel-Arbitrated Loop), IEEE P1394, and SSA. FCS was invented by IBM and is designed for high speed (100 MB/s) transfers over switching networks. It requires expensive adapters and switches. FC-AL helps offset some of the expense issues, but introduces performance limitations due to arbitration. IEEE P1394 was developed by Apple to facilitate the attachment of external disk, tape, CD-ROM, printers, etc., to a desktop processor without the need for bulky SCSI cables and power cords. The IEEE P1394 is also known as "Firewire".

Some of the advantages of SSA over FC-AL are parallel data transfers, no arbitration, spatial re-use, and fault tolerance in a single loop without the need for bypass circuits. (This is not a complete list.)

SSA combines low cost with high performance and availability. This makes it an ideal interface for both the high end and low end computing environments.

Foil 7
SSA's Opportunity

SSA is a robust high performance serial storage interface, which has been designed to overcome the bottlenecks of SCSI bus architectures, while leveraging the strengths of SCSI protocols. Like SCSI, SSA is designed to be a low cost open connection for IO devices such as disk, tape, CD-ROM, optical, printers and scanners to workstations, servers and storage subsystems. SSA supports the SCSI command set, for compatibility with existing application software. At the same time SSA provides superior performance, which scales, or increases, as we add additional devices to an SSA configuration. With today's high performance disk, the MB/sec performance on a SCSI bus begins to flatten out after adding six disks, while performance on an SSA loop continues to increase.

Connections to SSA devices are simpler, allowing multi-host attachments, hot swap drives, increased distances, and non-disruptive reconfiguration. SSA is a second generation serial interface for IBM. The IBM 9333 High performance Serial Disk Subsystem was introduced in 1991 and was very popular with high-end clustered servers, that required high performance and multiple host attachments. However, the 9333 was more expensive than comparable SCSI products like the 9334 and was considered a "closed" IBM only interface. With SSA, IBM has improved on the 9333 by adding a loop architecture, increasing the link speed from 8 MB/sec to 20 MB/sec, improving packaging of the disk subsystem and adapter card, and significantly decreasing the costs. And then in 1999, IBM again improved on this technology by improving the link speed to 40 MB/s, and continued to provide more function in the disk systems and adapter cards.

Although invented by IBM, SSA was made available to be developed as an "Open System" architecture by the American National Standards Institute (ANSI), ANSI X3T10.1. An industry standard which describes the protocol and electrical specifications of the SSA interface were approved in October of 1996. An industry association of leading storage subsystems companies was formed in 1995 and has been instrumental in bringing SSA to standardization.

Foil 8
SSA Architecture
Enhanced Non-shared and Non-arbitrated

SSA is best characterized as a multi-lane highway compared to a SCSI one-lane bridge. With SCSI, two lanes must share a one lane bridge, allowing only one-way traffic across the bridge. Constant coordination and communications are required for traffic control - arbitration.

In comparison, SSA is a multi-lane highway. Reads and writes have their own 40 MB/sec (with SSA 80 it is 20 MB/s) fast lane. When configured into a loop, as we will see later, SSA provides multiple read/write paths. You will notice from our simple picture that there are multiple cars in each lane. This represents another important feature of SSA known as frame multiplex. With SSA, data transfers are broken up into frames or packets that can be multiplexed with frames from other data transfers on the same read or write path. This prevents a large data transfer from dominating the SSA path at the expenses of other data transfer requests. Frame multiplexing provides greater throughput and more consistent response times. This facilitates the attachment of devices with different data transfer speeds and blocksizes (disk, tape, printers) on the same SSA loop

Foil 9
SSA 160 - Architecture

SSA consists of intelligent nodes connected by a full duplex (40 MB/s in both directions for SSA 160, 20 MB/s for SSA 80) serial link. A node can be a target like an IO device or annotator like a host adapter. Each node has two ports which allow them to be connected into a loop. Loops provide two paths to every node for availability and performance. An SSA 160 Serial Interface Chip (SIC) on an adapter which is configured into a loop, has a maximum bandwidth of 160 MB/sec, and 40 MB/sec in and out in both directions around the loop.

The SSA loop architecture supports up to 128 nodes or a maximum of 127 devices with one adapter node, versus 15 devices for SCSI-2. The links between nodes can be 25 meters apart versus 25 meters for the entire SCSI-2 bus from adapter to the terminator. Since SSA is a serial interface, it only requires four wires for duplex reads and writes, versus 68 for non-duplex SCSI-2 Differential Fast/Wide, which is the best that SCSI-2 can offer today. With an Advanced SSA Optical Extender, distances between nodes can be up to a maximum of 10 km for flexible configuration.

SSA addresses are set dynamically, while SCSI-2 addresses must be set manually. Routing of data between an initiator and target is optimized. When an initiator (adapter) is powered up or when the loop configuration is changed, the adapter will "examine" the loop and determine the optimum (shortest) path to each target device and automatically determine the addressing. In the case of four devices, it will access two from one side and two from the other side.

Foil 10
Frame Multiplexing for Reliability and Throughput

Note to Presenter:

This foil contains technical detail which might best be used as a backup foil.

Each SSA node contains a serial interface chip, consisting of two ports and a three-way switch. Data comes in one port and is examined for an address. If the address is for that node, the data is directed to the node function which may be a device or an adapter. If it is not for that node, the data is directed to the other port for transmission to the next node in the loop. Data on an SSA loop is divided into packets or frames.

A frame of information is initiated by a host adapter and passed from one node to the next until it reaches the target node. A frame consists of a header containing Flag, Control, and Address bytes, a maximum data area of 128 bytes, and a trailer containing Cyclical Redundancy Check (CRC) bytes.

The nodes on either side of a link are responsible for passing the frame across a link and ensuring that the data is passed correctly, based on the CRC bytes in the trailer of an SSA frame. Since only the nodes on either side of a link are involved in the transfer of an SSA frame, each link can operate independently of other links, and several transfers can occur concurrently at full bandwidth as long as they use different links. This is known as spatial reuse and is different from bus architectures which require the dedication of the entire bus for the duration of each data transfer.

The adapter counts the number of nodes it traverses to get to a device (sometimes called a "hop" count) and uses this count to address the device. Each device thinks its address is zero. This means that when a frame arrives at a node containing an address of zero, the node accepts the frame into its node function. If the address is not zero, the node decrements the address by one and passes the frame to the next node. The overhead for routing frames is very low since it is done in the node hardware. SSA is a forward and store interface, which means that frames are routed immediately after the first few bytes are analyzed and the address is not zero.

Foil 11
SSA - Superior Availability

With the SSA loop design there is no single point of path failure. There are two paths to each device. If there is a failure in one path, the SSA adapter will automatically reconfigure and use the second path to access a device. The SSA loop then operates as two strings. Once the path failure is corrected, the SSA adapter will recognize this and reconfigure back into a loop and restore the normal mode of operation.

If there is a disk failure, the disk can be removed while the other disks in the loop continue operations. Since SSA disks are "hot swappable", a new disk can be inserted in place of the old disk without stopping the system, or powering down the subsystem. SSA will recognize the new disk and reconfigure the loop dynamically.

Growing an SSA configuration can also be done dynamically. The SSA loop can be disconnected and a new disk inserted or a new subsystem connected into the loop while activity continues to all the other disks. SSA will automatically add the new disk into the configuration of the loop.

In addition to these availability features which are inherent in the SSA architecture. IBM SSA products are implemented with additional availability features such as redundant power and cooling, RAID, and host bypass features which will be covered in the product section.

Foil 12
IBM 4-Port SSA 80 Adapters - Available Today!

IBM has implemented SSA architecture in a family of products which are available today. IBM SSA products provide a cost-effective, high performance serial connection that meets a broad range of interface requirements from PCs to massively parallel processors like the RS/6000 SP. The fundamental SSA performance building block is a single SSA Serial Interface Chip (SIC) that supports two SSA ports which can be connected into an SSA loop. Each port is capable of carrying on two conversations at once, one outbound at 20 MB/sec and another inbound at 20 MB/sec, for a total of 80 MB/sec for an SSA 80 SIC chip. If a link in the loop is disconnected at some point, the two ports will operate as two separate strings. IBM SSA adapters support two SSA SIC modules, for two loops or four ports, which have a combined maximum data rate of 160 MB/sec. The four ports can operate independently as separate strings. The optimum configuration for availability and performance is to attach SSA loops to each SSA SIC module. Note that the SSA architecture doesn't dictate the number of loops supported per adapter - different adapters may support one or more loops depending on the manufacturer. Loops cannot be cross connected between the two SIC modules on the same adapter. Up to 48 disks and eight SSA adapters (feature code 6216, 6215, 6219, 6225) can be configured in one loop today. One SSA adapter could access up to 48 disks on each loop or 96 disks across its two loops. Up to 3,000 I/O operations per second have been demonstrated on these SSA 80 Adapters.

Foil 13
Scalable Performance with SSA Disks

This chart represents measured data for a testcase running 70% reads and 30% writes with 4 KB blocks and short seeks, comparing SCSI-2 Fast/Wide to SSA 80. Operations per second are measured against the number of disk drives. The SCSI ops/sec performance curve begins to flatten after six disks. At this time, the IO response time begins to increase exponentially due to increasing queuing delays. With today's high performance disks, most users limit the number of disks per SCSI bus to six to avoid such bottlenecks.

SSA enables the performance of each disk drive to show through up to about 32 disks. A point to note is that the difference between SSA and SCSI performance is not seen in very small configurations. When benchmarking against other subsystems, be sure to have enough disks to showcase SSA's bandwidth performance. Another point to consider is that this benchmark caused 80% of the RS/6000 59H processor to be used. If you benchmark, be sure you do not have a bottleneck in the processor or memory.

Foil 14
Data Warehouse Performance with 7133

This is the latest TPC-D result published on 2/7/97 and available on the Internet at http://www.vfecteau.com/tpcd. The Transaction Processing Performance Council (TPC) defines standard benchmarks which are widely used by UNIX vendors to rate their systems against competition. TPC-D is a benchmark designed to evaluate Decision Support applications that require complex, long-running queries against complex data structures. Price performance is measured in dollars per query per hour.

The winning IBM solution used DB2 running on an RS/6000 SP with 12 high nodes (8-way SMP). The SSA configuration consisted of 12 SSA adapters per SP wide node, and 72 x 4.5 GB 7133 disks. This solution provided the best price/performance with $/Q/hr of 6,935. Although the IO configurations for Tandem and NCR are not available in this report, they were the best that they could offer. At equivalent system cost, IBM beat Tandem on load time and $/Q/hr. At equivalent load time, IBM beat NCR/Teradata in system cost and $/Q/hr. SSA clearly had an impact on total system cost and load times which helped to provide the best $/query/hour. Note: this benchmark was done with SSA 80 disk systems and adapters. Even greater performance would be seen with IBM's 160 MB/s SSA technology.

Foil 15
New Paradigm for Disk Systems

IBM's implementation of SSA has introduced a paradigm shift in the traditional way we build disk systems. SCSI disk subsystems are usually built into a cabinet or drawer with a controller, which consists of a processor, an internal bus, and several SCSI adapters which attach the disks. This controller is then attached through a SCSI interface to a SCSI adapter in the host. Some controllers had multiple processors, multiple internal busses, cache, as well as multiple SCSI interfaces. The purpose of the controller was to overcome some of the limitations of SCSI, such as addressing and cabling, and to add availability or performance enhancing functions like RAID. This is costly to build, and when new technologies or functions were required, it often meant a costly box replacement. With SSA, the controller function is moved from the disk system to the SSA adapter card which defines the personality of the disk system. New functions like RAID or PCI support can be added to existing SSA disk systems by changing a low cost adapter, and utilizing the same SSA 7131-405 or 7133 disk system. This provides lower costs, investment protection, and ease in implementing new technologies and functions for both the manufacturer, business partners, ISVs and customers.

Foil 16
RS/6000 SSA Adapters and Disk Systems

IBM's Family of high performance SSA adapters for the RS/6000 can attach any of the 7133 or 7131-405 SSA disk systems. The basic features of each RS/6000 adapter and SSA disk systems is summarized on this chart.

Foil 17
7133 Serial Disk System - Advanced Models

The 7133 Serial Disk System models T40 and D40 are the latest models of IBM Serial Disk Systems to be announced. And, they are the only models that support SSA 160 bandwidth. The 7133-D40 is a drawer model. The 7133-T40 is a deskside tower model. Both models have bays for 4 to 16 drives. Unlike the 020 and the 600 models, the T40 and D40 support 36.4 GB and 18.2 GB disk drives as well as 4.5GB and 9.1 GB drives. The Advanced models also support new 10k RPM drives in 9.1 and 18.2 GB capacities. Drawer capacity is 582 GB, yielding a total 6-drawer rack with a storage capacity of 3.5 TB.

These models also provide other advancements over the 020 and 600 models with a new carrier design and enclosures which yield greater overall disk system reliability and serviceability. Like the 020 and 600 models of the 7133, the D40 and T40 have hot-swappable drives and host bypass circuits. They have 3 fans where 2 were standard with the older models. And, two power supplies instead of three as were provided in the 020/600 models, are now provided with the T40 and D40. Each of these power supplies is capable of supporting a fully loaded drawer or tower. This makes it easy to connect a separate power source to each power supply to create an even more highly available power configuration. Power supply options include AC as well as -48V DC power in support of the Telecommunications industry customers.

Like other 7133 models, these advanced models can be attached to a wide array of UNIX and Intel-based server platforms using IBM SSA adapters, or the IBM 2108-S20 SAN Data Gateway for Serial Disk. Additionally, an Advanced SSA Optical Extender is an optional feature of the D40 and T40 models. This optical extender allows T40 and D40 models to be connected up to 10 km apart, allowing greater distances for configuration flexibility and disaster recovery configurations.

Like the 7133-020, the 7133-D40 is supported in the Versatile Storage Server.

StorWatch Serial Storage Expert can also be used very effectively with the advanced 7133 models. With the advanced 7133 models, the Serial Storage expert will be able to automatically discover and provide information concerning the configurations of the attached disk storage. This information includes which disks are connected to which adapter and where those disks are located in the drawer or tower. This automatic discovery of detailed drive location information will only be available with the advanced D40 or T40 models of the 7133 when used with the StorWatch Serial Storage Expert.

Foil 18
IBM 7133 Models D40/T40 Flexibility

Most currently available SSA adapters and attachments are supported on the T40 and D40. The 7133 D40 and T40 models can be attached to RS/6000 and RS/6000 SP servers using the RS/6000 adapters 6214, 6216, 6215, 6219, 6225. These are both PCI and MCA adapters. (Note: The 6217 and 6218 adapters are not supported on the T40 and D40 models.) The 6225 is the Advanced SerialRAID Adapter - the only RS/6000 adapter that supports SSA 160 MB/s bandwidth. The Advanced SerialRAID/X Adapter (7133 Feature 4040) supports SSA 160 throughput for Intel-based servers. These models are also supported by the IBM 2108-S20 SAN Data Gateway for Serial Disk which allows Ultra SCSI attachment to many HP, Sun, and DEC UNIX and Windows NT platforms. Also, Sun systems can be attached using the IBM SSA to Sun SBus Interface Controller. Microcode on all of these attachments needs to be updated to supported levels for T40 and D40 SSA 160 connections. This can be down-loaded off the web.

The advanced models can interoperate on the same loop with other models of the 7133 as well as with the 7131-405. So, the advanced models should be compatible with most previously installed SSA disk.

The Advanced Optical Extender can be used to connect T40 and/or D40 models between the models or between SSA 160 adapters (Advanced SerialRAID adapter for RS/6000 and Advanced SerialRAID/X adapter for Intel-based servers) and the T40/D40 models up to 10,000 meters apart. This enables the easier creation of disaster recovery configurations at this extended distance. Up to 2400 meters can be reached between the advanced T40/D40 models and other SSA devices such as SSA adapters, other 7133 models. Note that the 7131-405 does not support attachment to the Advanced Optical Extender..

Foil 19
7133 T40/D40 Advanced Models - Details

The 7133 T40 and D40 models have some significant differences from the 020 and 600 models. To make the ordering process easier and more understandable for our customers and our business partners, the T40 and D40 models do not have standard drives or select features. A minimum number of 4 drives must be selected, but these drives will be priced at the same price as the rest of the 16 drives that can be placed in the enclosure. A choice of black or white drawer covers is also provided as an option so that the drawers can match whichever rack the customer chooses.

The operator panel of the 7133 has also been redesigned in these models. The LEDs and information concerning the status of power, cooling, and disk drives is displayed on a consolidated panel at the front of the enclosure. Additionally, a separate SES controller inside these units monitors the status of the components in the system and provides detailed status information to the host system..

The 7200 RPM disk drives in the T40 and D40 are new faster drives, and they will provide some performance enhancements over the older drives which had been used in the 7133 020 and 600. Although we also began to ship these drives in 020 and 600 models, the carriers used in the 020 and 600 are different from the carriers that are shipped with the drives for the T40 and D40. The drives for the T40/D40 are, therefore, not interchangeable with the drives we ship for the 020/600.

We are also now offering 10K RPM disk drives for the T40 and D40 in 9.1 GB and 18.2 GB capacities.

New SSA cables are required for use with the T40/D40. These cables have been colored blue for easy recognition. Also, new power supplies and power cords are required for the T40 and D40 models. Two outlets are now needed per D40, so a rack with more than 3 drawers will require you to order an additional rack PDU (Power Distribution Unit). This additional PDU is a chargeable feature on the rack.

Note that updated for the 7133 D40/T40 is required in order for SSA 160 MB/s to be enabled on these models.

Foil 20
7133 More Performance and Capacity

We have new additions to our SSA disk drive family, including
a 36.4 GB disk drive. This new 7200 rpm drive doubles the capacity of our 7133 enclosures. Using this 36.4 GB drive, our 7133 enclosures now provide a capacity of 582 GB and a 6-drawer rack can hold up to 3.5 TB. That's a lot of storage packed into a small footprint.

We also now have 10K RPM disk drives in 9.1 and 18.2 GB capacities. These drives are the same form factor as our 7200 RPM drives but run at a faster rotation speed. They provide approximately a 20% random throughput performance improvement and a 35-40% sequential throughput improvement over our 7200 rpm drives.

These 10K RPM drives are available in addition to our 7200 rpm drives, so you now have a choice of drive sizes as well as rotation speeds. You can order 7200 RPM drives in 4.5, 9.1, 18.2 or 36.4 GB capacities, or 10K RPM drives in 9.1 or 18.2 GB Capacities. The 7200 RPM drives are priced at a lower price point than the 10K RPM drives. So, the choice is yours.. 10K RPM drives for better performance, or 7200 RPM drives for cost effectiveness..

Additionally, , included in the drive announcement, we made a statement of direction announcing our intention to provide Novell Netware Version 4.11 and 4.2 as well as Netware Version 5.0 support for 7133s attached to supported Intel-based servers. This support will be provided using our Advanced SerialRAID/X Adapter - the new SSA 160 adapter for Intel-based servers that we announced in March. We have also stated plans to support RAID 0+1 on that adapter. I know that many of our customers have Novell and RAID 1 requirements This statement of our plans should provide assurance to our customers that this support is in our plan and is coming soon..

Foil 21
7133 10K vs 7200 RPM Drive Throughput
4-K Random -Read/Write 70/30

This chart shows a comparison of disk drive throughput using random 4K writes in a 70:30 read/write ratio. As you can see from the chart, our 10K RPM drives are significantly faster than our 7200 RPM drives - around 20% for this random workload. And our newer 7200 RPM drives which became available after 10/98 are significantly faster than the 7200 RPM drives we had before that point in time. It is clear that the 10K RPM drives offer a performance advantage to our customers. It should be noted that the performance advantage is even stronger when using sequential writes. There we see a 35-40% throughput improvement using the 10K RPM drives.

The other thing to note on this chart is the improvement on throughput of even the larger 10K RPM drives over the smaller capacity 7200 RPM drives. The greatly decreases the need for using the smaller capacity drives in order to get the required performance for your customers applications..

Foil 22
7131 Model 405 SSA Multi-Storage Tower

This is an SSA version of the 7131 Model 105 SCSI Multi-Storage Tower. It is designed for low cost entry configurations and customer setup. Although it is called a tower, its small size suits it for desk top as well as desk side storage. It attaches to all SSA adapters and supports 48 disks per loop and 96 disks per adapter. 7131 Model 405 and 7133s can be intermixed in the same SSA loop. The base unit requires two 4.5 GB or 9.1 GB disks, located in the top and bottom bays, and comes with one SSA cable. Be sure to order a second cable in order to install it as a loop. While it is called a multi-storage tower, like its SCSI model, there are no SSA tape or CD-ROM modules available at this time. The two top bays which would normally contain tape or CD-ROM modules are not used. This is supported on UNIX and NT servers - HP, Sun, and DEC with the IBM 7190, on Sun Servers by using the IBM Sun SBus card, or on RS/6000 and selected Intel-based servers with IBM SSA adapters.

Foil 23
IBM RS/6000 SSA Adapters

This is a reference chart which summarizes all of the currently available RS/6000 SSA adapters.

Foil 24
IBM 6216 Enhanced SSA 4-port Adapter

The 6216 is an 8-way SSA adapter which means that up to eight of these adapters can be configured into the same SSA loop. The 6216 is ideal for use in multi-node RS/6000 SP configurations. 8-way 6216 and 2-way 6214 can be intermixed in the same Loop. However, when they are intermixed only two adapters are allowed on a loop. Adding additional adapters to a loop can provide increased performance and availability. The 6216 adapter cannot be intermixed on the same loop with the IBM SSA Multi-Initiator RAID EL adapters (6215 and 6219), or with the Advanced SerialRAID Adapter (6225).. The 6216 can be used to attach to the 7131-405 and all models of the 7133, including the Advanced T40 and D40 models.

Foil 25
7133 with Redundant SSA Adapters

This drawing shows four SSA adapters on one loop. The presenter can use a pointer or colored pen to trace out the single loop. This configuration has the added advantage of higher MB/s (up to 55 MB/s across the two adapters on each host). It also provides redundant adapters on each host for increased availability in the event of an adapter failure. Note the use of mirroring and Domains.

Foil 26
Sharing Two SSA Loops With 8 Hosts
8-Adapter Domain Configuration

This is a configuration that one might have on an SP2. With the Domain concept, each host can work on its own set of disks but are capable of accessing any disk in the loop. This is storage area network which is like a local area network for storage, without the overheads required for networks and with the performance of direct attachment. On an SP2 with Virtual Shared Disk, data residing on storage in one node is often shipped across the high performance switch to another node to be processed. This creates a load on a resource that was designed to handle commands and control information and not long data transfers. An SSA loop with the 6216, 6215, 6219, or 6225 offer a way to offload this data transfer from the high performance switch to an SSA storage network. One node might be designated as a backup node with a 3494 tape library. During off-peak hours, it can do direct backups of the disks on the other nodes.

Foil 27
Winning with IBM RAID

IBM has successful RAID products in every platform from the PC, AS/400, RS/6000 to RAMAC on the mainframe. The initial patent for a disk array was filed by IBM in 1978. The RS/6000 and Intel-based servers use SSA RAID adapters to provide RAID 5 protection for serial disk storage. On the RS/6000, all disks can be software mirrored with the Logical Volume Manager, a standard feature of AIX. SSA architecture facilitates the low cost implementation of multiple RAID 5 and non-RAID disks on SSA loops. The IBM Versatile Storage Server also provides RAID 5 protected storage along with dynamic storage allocation across a wide variety of UNIX, NT, and AS/400 servers.

Foil 28
RAID Architecture Summary

RAID definitions were initially developed in a paper published by the University of California in Berkeley, California. An industry association, the RAID Advisory Board, has further defined the various implementations of RAID. The three main types of RAID are RAID 1 or Mirroring, RAID 3 which is the parallel striping of data with dedicated parity, and RAID 5 which has distributed parity. This foil summarizes their advantages and disadvantages. RAID 1 and RAID 5 are most commonly used today. RAID 3 requires raw mode programming support and is not generally used. RAID 0 is a type of RAID that is not on this chart. It provides striping for performance, however RAID 0 does not provide data protection.

Foil 29
6215 & 6219 SSA Multi-Initiator/RAID 5 Adapters

IBM extended the serial storage function and performance with the introduction of the SSA Multi-Initiator/RAID Enhanced Loop adapters, available for PCI and MCA channel systems.

With these adapters, mainframe class protected Fast Write technology comes to SSA adapters. The optional 4 MB Fast Write Cache may only be enabled on 1-way RAID 5 or 1-way non-RAID configurations. The Fast Write Cache option dramatically improves RAID 5 write response time to less than 1 msec and also improves RAID 5 write throughput up to 13 MB/sec. It utilizes non-volatile RAM and has a data memory retention of over seven years.

SSA RAID adapters are very flexible. They can support up to 32 RAID 5 arrays on its two SSA loops. RAID 5 and non-RAID disks can be intermixed on the SSA loops. A RAID 5 array can be made up of between three to 16 disks. The capacity of one disk in RAID 5 arrays is required to store the parity. RAID 5 uses a parity which is rotated among the disks in the RAID 5 array. Since the parity is not tied to one disk, RAID 5 supports multiple concurrent users. A RAID 5 array of three disks is also referred to as a "2 + P", meaning two data disks and one parity disk. A "15 + P" would be an array of 16 disks. A RAID 5 array appears to AIX as one logical disk even though it may be three to 16 disks.

Hot spares can be assigned in one loop and will dynamically support any of the RAID 5 arrays in the loop. Hot spares are defined for a given adapter. When a disk in an array group fails, the hot spare joins the array group and data from the failed disk is rebuilt on the hot spare using the remaining members of the original RAID 5 array. Rebuilding of the data on the hot spare can be done in the background while normal processing continues. Multiple hot spares can be designated. On RS/6000s, RAID 5 arrays and hot spares are configured through the SMIT interface for ease of use. The use of hot spares is highly recommended.

6215 and 6219 adapters can be configured on the same loop with other 6215 and 6219 adapters and also with the Advanced SerialRAID Adapter (6225), however, they are not compatible with older SSA adapters like the 6214 or 6216.

An announcement was made April 1998 which enhances the function of these adapters to support 2-Way RAID 5 or 8-Way Non-RAID function. The availability date for this enhancement was July 31, 1998. 6215 or 6219 adapters shipped BEFORE the general availability date, were shipped with the microcode which supports 1-way RAID 5 or 2-way non-RAID function. Adapters shipped AFTER general availability have the microcode to support 2-way RAID 5 or 8-way non-RAID. This microcode upgrade is down-loadable via the world wide web, or orderable via RPQ. This microcode will allow previously purchased 6215 and 6219 adapters to support the new 2-way RAID and 8-way Non-RAID function. When used in the RAID 5 configuration, the drives can be configured in any combination of (2+P) to (15+P) RAID arrays.

Foil 30
Advanced SerialRAID Adapter

This is IBM's most recently announced SSA adapter for RS/6000 servers. The IBM Advanced Serial RAID adapter is the only RS/6000 adapter IBM announced that supports 160 MB/s bandwidth when attached to 7133-D40 or T40 models.

The Advanced SerialRAID adapter is a high performance storage adapter that supports the attachment of the 7133 Serial Disk System to RS/6000 and RS/6000 SP servers. It requires AIX 4.2.1 or 4.3.2 software. (4.1.5 customers will need to use the 6215/6219 SSA Adapters, though those don't support SSA 160.)

The Advanced SerialRAID adapter supports 4.5, 9.1, 18.2, and 36.4 GB SSA disk drives. Just like other SSA adapters, it supports up to 96 disks, 48 disks per loop. It is possible to create RAID 5 arrays of up to a maximum capacity of 546 GB per array (15+p) with a total storage capacity of up to 3.2 TB RAID-5 or 3.5 TB JBOD per adapter.

You can have up to 8-way JBOD configuration, up to 2-way RAID-5 configurations, or 1-way RAID 0 configurations..

Unlike other SSA adapters for RS/6000 servers, 64 MB of read cache is standard. Also, a 32 MB Fast Write Cache is an optional feature of the adapter for single host configurations that enables improved write response times. The Fast Write Cache card is operational only in standalone configurations, but it is possible to have the adapter support multi-way (up to 8-way) operations on one loop and the Fast Write Cache in a single mode on a second loop.

When combined with the SSA 160 Advanced 7133 D40/T40 models, this adapter provides exceptional performance. Most of the benchmark numbers show 2-3x the performance of SSA 80 which was already a high performing disk system technology. And unlike arbitrated technologies, with SSA, performance scales. Adding more drives and more servers yields more performance. An 8-way cluster configuration has shown 600 MB/s and 60,000 I/Os per second performance. Those numbers are pretty hard to top..

This adapter can be attached to the 7133 D40 and T40 and will operate with SSA 160 MB/s bandwidth in these configurations. Attachment to other 7133 models as well as to the 7131-405 and 3527 is also supported, however, when attached to those disk systems, the adapter will auto negotiate down to their link speed and will therefore only provide 80 MB/s bandwidth.

Foil 31
Advanced SerialRAID - Scalability

The Advanced SerialRAID Adapter provides excellent scalability . It allows up to 3.5 TB of disk storage to be attached per Advanced SerialRAID adapter. Up to eight servers are supported in a clustered configuration which allows the Advanced SerialRAID adapters installed in each server to share the same loop of disks. HACMP is also supported with this adapter so servers can be configured for back-up and availability, or can be clustered for performance. The RAID levels that are supported are RAID 0 and RAID 5, as well as non-RAID. Please note that RAID-0 is only supported in 1-way single server configurations.

The RAID 5 arrays range from 3 (2+p) to 16 disks (15+p) which gives a maximum partition size of 546 GB. The maximum number of arrays that is supported per adapter is 32 with each array being 2+p size.

Also, the adapter scales in connection distance. If used with the Advanced Optical Extender and the 7133 D40/T40, up to 10,000 meters can be spanned between the host server and the disk system.

Foil 32
SSA 160 RS/6000 Advanced Scalability

The Advanced SerialRAID adapter supports clustered configurations with up to eight servers running AIX.. When Advanced SerialRAID adapters are configured in a single loop, the disk drives can be pooled with some servers owning some drives and other servers owning other drives. HACMP can be installed and configured so that if either the host or the adapter fails, the remaining adapter can access data in the arrays. Additionally HACMP and selected databases can be used to create clusters of servers that can have concurrent access to the disks for additional performance..

The Advanced SerialRAID Adapter is SSA160 capable and provides up to three times the performance of the previous adapters that run at 80 MB/sec bandwidth.

One of the main differentiators of SSA when compared to other technologies is that the performance really scales with added drives. You can not only expand the capacity but also scale the performance as drives are added. Because SSA is a full duplex, point to point architecture, it allows for simultaneous read and write operations in the same loop. Higher performances can be achieved with SSA's spatial reuse capability which can be exploited heavily with large configurations or clusters. Up to 60,000 I/Os per second or 600 MB/s have been seen in an 8-way RS/6000 cluster using this SSA 160 technology.

Foil 33
Advanced SerialRAID Adapter & 7133
D40/T40 Configuration Example

This is an example of a RAID configuration which highlights the versatility and availability of SSA RAID. You can satisfy almost every level of performance and data availability with this one SSA RAID 5 adapter and the AIX Logical Volume Manager. For write intensive, high performance, high availability data like a data base log, you can use LVM mirroring across the two loops. A catalog which has a high read ratio, and high data availability requirement might use a small RAID 5 array consisting of 3 disks. Large data base tables and indexes with 70/30 read/write ratios might use larger RAID 5 arrays. Raw data which might be loaded from tape then loaded into the data base would not need RAID protection since a copy resides on tape, and could be a candidate for LVM data striping or RAID 0 striping on the disks.

Since the capacity of one disk in a RAID 5 array is required for parity data, the more disks in an array, means less capacity is required for the parity and more capacity is available for data. However, more disks increases the probability of a disk failure in the array, and increases the time to reconstruct data in the event of a disk failure. A good rule of thumb for the most efficient use of parity is eight disks. This is known as a "7 plus P", or seven data disks and one parity disk. Although this picture allows the members of an SSA RAID 5 array to be physically adjacent to each other, this is not required. Members of a RAID 5 array may be anywhere on the loop. Hot spares can be dynamically supported on any of the arrays on the same loop.

As in any SSA loop, disks of any size can be intermixed. However, if disks of different sizes are intermixed in a RAID 5 array, the capacity of each disk becomes that of the smallest disk.

Foil 34
SSA - Superior RAID Capabilities

This foil illustrates the advantages of implementing RAID on an SSA loop versus SCSI strings. Due to SSA's node and loop architecture which supports multiplexed, concurrent data transfers with no single point of path failure, all devices in an SSA RAID array can reside on the same SSA loop. By contrast, to implement RAID using SCSI architecture, one would need multiple busses to isolate busses and drives since each bus and all drives on that bus represent single points of failure. SCSI bus separation is also needed to overcome the overhead of bus arbitration. These multiple SCSI busses are contained in an external controller which adds a great deal of cost. The mixing of different RAID arrays, and support of hot spares and hot swapping of disks is also more difficult and more expensive to implement with the SCSI architecture.

Foil 35
SSA RAID 5 Considerations

This chart summarizes the points that were made on the previous configuration example. The 6225, 6215 and 6219, should be positioned as low cost data protection. These adapters also provide write cache, multi-host attachment, and concurrent non-RAID and RAID 5, and RAID 0 (RAID 0 on the 6225 only)..

Foil 36
RAID 1 Mirroring with SSA 7133

RAID 1 mirroring with LVM provides the highest performance and availability, especially in High Availability Clustered Multi-Processor configurations. Since a 7133 can be configured into as many as four independent loops, mirroring can be done in one 7133 disk system, while SCSI require s mirroring across separate SCSI subsystems which increases the cost. With 36.4 GB disks in a 7133, we can mirror up to 291 GB at about the same cost as a SCSI RAID 5 subsystem and provide superior performance and availability to as many as eight hosts.

We can expect to see improvements to SSA RAID. However, we do not have to wait to install SSA. You can install 7133 with mirroring today, and when the appropriate RAID function is available, you can convert the same 7133 to a RAID configuration by replacing an SSA adapter card. Data would have to be unloaded from a mirrored non-RAID logical volume and reloaded to a RAID 5 logical volume. Not only will investment in the 7133 be protected, but additional capacity will be freed up by converting mirrored pairs to n + p RAID array groups.

Foil 37
SSA 160 Superior Throughput

SSA 160 offers truly exceptional serial storage performance for Unix and Windows NT servers. Combined with IBM's 7133 Serial Disk System Advanced Models D40 and T40, the Advanced SerialRAID adapters can provide up to 2-3 times the performance available with other SSA adapters. The throughput is excellent compared to other technologies as well. This SSA 160 architecture offers a bandwidth of 160 MB/s. Compare that to disk enclosures using other architectures such as SCSI-2 with a bandwidth of 10 MB/s, Ultra SCSI with a bandwidth of 40 MB/s, or even FC-AL with a bandwidth of 100 MB/s and you can see how impressive SSA 160 is.

Foil 38
Disk Storage Technology

Storage technology has come along way in the last few years. SSA 160 offers truly exceptional serial storage performance for Unix and Windows NT servers. Combined with IBM's 7133 Serial Disk System Advanced Models D40 and T40, the Advanced SerialRAID adapters can provide up to 2-3 times the performance available with SSA 80 technology, and 4 to 5 times the performance of Ultra SCSI technology..

Foil 39
SSA 160 Adapter - Performance

The SSA 160 adapters - the Advanced SerialRAID Adapter for RS/6000 servers and the Advanced SerialRAID/X Adapter for Intel-based servers - offer up to 2 to 3 times the performance of other SSA adapters. The performance increase is dependent on the SSA 160 MB/s technology that is a capability of these adapters when attached to the SSA 160 disk technology of the 7133 Advanced models D40 and T40. This new technology takes advantage of the new, powerful SSA processors that are in the 7133 advanced models as well as the firmware that has been developed to support this function in order to create overall high performance. However, there is more to it than that.

Up to 10,000 I/Os per second have been measured on sequential single sector reads using one adapter. This high performance is also possible because 64 MB of read cache is now a standard part of these adapters. Other SSA adapters only had up to 32 MB of read cache. Also, 32 MB of non-volatile fast write cache is now an option for single server configurations. (Single server or, single adapter per loop configurations.) Previous adapters only had up to 4 MB of write cache, so this is an 8X increase of write cache. The write cache enables response time to improve down to the 1 ms time frame. Also, the cache allows sequential write streams to be saved up and written all at once which can then decrease the number of writes that are necessary to improve the throughput. This can be especially helpful sequential data like that of database log files. Additionally the new adapters allow full stride writes that increases performance on RAID 5 configurations by the writing whole array width - available on up to 4+p configurations..

Foil 40
SSA 160 Configurations

This chart was taken from the SSA 160 T3 and is intended to be used more for field education than as a customer presentation. It lists the IBM SSA products that currently support the 160 MB/s capability.

The 7133 D40 and T40 are both SSA 160 capable, however, updated microcode is required in order for this to be implemented. There are two SSA 160 adapters currently available, the Advanced SerialRAID Adapter for RS/6000 servers and the Advanced SerialRAID/X Adapter for Intel-based servers.

When the Advanced 7133s and Advanced SSA 160 adapters are used, the new SSA cables that are features codes on the 7133 D40/T40 must also be used in order to get SSA 160 capability. Also, optical extender configurations for extended distance must use the Advanced SSA Optical Extender with the advanced 7133s and SSA 160 adapters in order for SSA 160 performance and also in order to get the improved distance capability of the Advanced Optical extender - up to 10,000 meters. The advanced 7133 models with the updated microcode and blue SSA cables will operate at SSA 160 MB/s capability within the drawer. If they are connected to SSA 160 adapters, the 160 MB/s bandwidth will also occur up to the adapter. However, if they are connected to SSA 80 host adapters, they will still operated at 160 MB/s speed within the drawer/tower, but the speed up to the host will only be 80 MB/s. So, the SSA devices will negotiate down to the lowest common link speed. Advanced disk systems, adapters, microcode, adapters and extenders must be used in order to ensure SSA 160 MB/s bandwidth. Errors on the SSA 160 links can also drop speeds down to 80 MB/s

Foil 41
SSA 160 Superior Performance

IBM Serial Disk Storage provides superior performance - up to 2 to 3 times that of SSA 80 which already provided excellent performance. These are measured performance numbers that compare favorably to almost all other offerings. (You may see some bigger numbers quoted for some FC-AL configurations, but many of those numbers represent multi-server, 100% Cache hits only. - Make sure you compare apples to apples) Most of the benchmarks on this chart are for random read/write performance...closer to what a customer would actually see. It is important to note that the performance a customer actually sees will vary based on application work-load, environment, and server configuration. This 7133 Serial Disk System performance should provide more than enough bandwidth than most customer environments need..

Foil 42
SSA 160 Performance Scalability

IBM Serial Disk Storage provides advanced scalabiity in performance. SCSI typically ranges from 10 to 40 MB/s bandwidth. Best of breed implementations with multiple strings and adapters are sometimes up to 150 MB/s maximum. FC-AL only has 100 MB/s bandwidth unless 2 hubs and loops are implemented and then it goes to 200 MB/s. SSA 160 has a bandwidth of 160 MB/s for just one loop and one adapter. Unlike the other technologies it also scales well past that. Up to 8-way clusters can be implemented (with an RS/6000), yielding performance of up to 600 MB/s and 60,000 I/Os per second. (This is non-RAID localized traffic 8-way performance.) This is unparalleled scalability in performance.

Foil 43
SSA 160 RAID 5 I/Os per Second

This chart shows the RAID 5 I/O operations per second performance that has been measured with the SSA 160 adapters and disk systems. It is important to point out that the RAID 5 performance is excellent on these new adapters and that measured benchmarks show the SSA 160 RAID 5 performance close to that of the SSA 80 non-RAID or RAID 1 performance (RAID 1 was an option on some of our SSA RAID adapters for Intel-based servers.) Benchmarks do tend to push the upper limits with the number of drives used and data pumped through, and SSA scalability shines through in those configurations. So small numbers of drives, or more cpu intensive workloads may not see this kind of advantage. However, in general, most configurations will see an advantage with the new SSA 160 bandwidth and the RAID 5 performance. These measurements do not include the use of fast write cache.

Foil 44
SSA 160 RAID 5 MBs per Second

This is similar to the previous chart except that it shows Mbs per second instead of I/Os per second. Also, 64K block sizes were used instead of 4k.

Foil 45
Advanced SSA Optical Extender

The Advanced SSA Optical Extender (Feature Code 8851) for the 7133 models T40 and D40 can extend the distance between advanced 7133 models T40 or D40, or between Advanced 7133 Models D40 or T40 and SSA 160 adapters (like the IBM Advanced SerialRAID adapter) up to 10,000 meters. It can extend the distance between SSA 80 nodes (7133 020/600, SSA 80 adapters, 7190-200, VICOM, 3527, etc.) up to 2400 meters. (Note the 7131-405 does not support attachment to the advanced optical extender.)

One end of the adapter plugs into the SSA port on the 7133 or other SSA device, the other connects to single-mode optical cable. The Optical extender converts the electrical signals into two optical signals, one inbound and the other outbound. So, each link requires two fiber optic cables to make a link connection between the optical extenders. The Advanced SSA Optical Extender is available as feature code 8851 on the 7133 D40 or T40 models.

Foil 46
Advanced SSA Optical Extender Performance

Response time increases over distance due to the requirement to receive positive acknowledgment that the data transmitted to a node has been received correctly. With the Advanced SSA Optical Extender, when connected to advanced SSA 160 7133 models D40 or T40 or SSA 160 adapters, 10,000 meters is the maximum distance that is supported before a time-out occurs. The maximum throughput in ops/sec or MB/sec depends on the application workload, blocksize and fiber cable length. Many applications with balanced read/write ratios see little, if any, performance degradation. In many cases the limitations of the processor or the adapter card will have a greater effect on performance than the optical cable lengths.

Foil 47
SSA Optical Extender Configurations

Optical extensions of up to 10,000 meters makes it feasible to place storage in a protected area for disaster recovery. Remote backup to a processor in a secured area, or remote access to mirrored data in case of a local site outage is a common requirement. Many customers have a need for a central storage pool that can be accessed by many different processors in the building or campus. In a campus storage area network, each node (7133, SSA adapter, or other SSA device), can be located up to 2400 meters from each other, using the SSA Fibre Optic Extender or the Advanced SSA Optical Extender. Additionally, 7133 D40 or T40 models can be separated from other D40 or T40 models up to 10,000 meters using the Advanced SSA Optical Extender, and host servers with SSA 160 adapters, like the Advanced SerialRAID adapter, can be separated from D40 or T40 models up to 10,000 meters using the Advanced SSA Optical Extender. Many customers have decided on Serial Disk System technology because of this capability which is not readily available on SCSI systems..

Foil 48
IBM Supports SSA Open Attach

With the IBM SAN Data Gateway for Serial Disk, you can connect a loop of SSA drives to Ultra SCSI hosts. It provides attachment of serial storage subsystems (7133, 7131-405 and 3527) to various hosts including HP, Sun and DEC.

With the IBM SSA Interface Controller Card for Sun SBus, native attachment of high performance IBM serial disk storage to Sun SPARCstation, SPARCserver, SPARCcenter, and Ultra Enterprise is supported.

And of course with other SSA native host adapters, we support direct attachment of serial storage systems to RS/6000s, RS/6000 SPs, and selected Intel-based servers such as IBM Netfinity, Compaq, HP, and Dell servers.

Foil 49
IBM SSA-to-Sun Sbus Interface Controller Card

The SBus adapter further enhances the connectivity options for the IBM 7133 Serial Disk System by allowing native attachment of the 7133 to supported Sun SPARC and Ultra Enterprise servers. This adapter delivers the same breakthrough storage subsystem performance that the RS/6000 enjoys today. The features and functions of the SBus adapter are comparable to the Enhanced SSA 4-Port Adapter, RS/6000 feature 6216. Support for 4 SSA Sun Sbus adapters per Sun server as well as support for Solaris 2.6 is available starting 11/20/98.

The comparison chart shows the features and functions of the SBus adapter compared to the RS/6000 Enhanced SSA 4-Port Adapter. Note that the Sbus Adapter is a feature of the IBM 7133. Two adapters within an individual Sun server must be on separate loops. The sustained performance of up to 35 MB/s and up to 3000 IOs is an adapter card/system bus performance statement and in no way represents the maximum capabilities of the IBM 7133. The 3000 IOs performance benchmark uses 70%/30% read/write mix of 4K block size transactions.

RAID is available with the Solaris operating system and the Veritas file system and volume manager. Veritas offers RAID 0, 1, 1+0, and 5, with hot spare and Snapshot feature.

Foil 50
Four SBUS Adapters and Four Servers

The maximum SBus Adapter loop configuration with 4 adapters, 48 disk drives, and 872 GB of storage is depicted on this foil. From one to four Sun servers may be connected to this loop. A second independent loop of up to 4 adapters and 48 disk drives may be connected to the same set of servers for a total of 96 disk drives and up to 3.5 TB of storage. This total amount of storage may be connected to one Sun server or to as many as four Sun servers simultaneously.

Foil 51
IBM SAN Data Gateway for Serial Disk

Note: the following charts on the SAN Data Gateway for Serial Disk are a subset of charts taken from a bigger presentation and guide on this product called 2108-S20 Storage Area Network Gateway for Serial Disk. This presentation can be found on MKTTOOLS and the IBM Intranet at w3.ssd.ibm.com/disk.

The IBM SAN Data Gateway for Serial Disk provides high function connectivity for IBM Serial Disk Storage.

It provides Ultra SCSI and SCSI connectivity for UNIX and NT servers. Standard host SCSI adapters and drivers may be used without any change.

The intelligent SAN Data Gateway provides host server independent functions including:

- Disk concatenation of multiple disks into a single logical disk image. This is especially important for Window NT servers which have limited device addressing capability.
- RAID 1 disk mirroring provides high availability without requiring any change to host server software or operating procedures, and
- Instant Copy can create separate disk copies for tape backup processing. This improves data access because critical applications can continue to operate in parallel during the backup processing.

IBM StorWatch SAN Data Gateway S20 Specialist provides an easy to use disk system management interface. It is called the S20 because it supports IBM 2108 Model S20 Data Gateway.

Highly scalable Serial Disk loops support a mixture of up to 8 UNIX and NT servers. Each loop can support up to 64 Serial Disks. With the new 36 GB Serial Disks, over 2 TB per loop can be supported!

The Data Gateway maximizes the 40 MB/s Ultra SCSI bandwidth potential by minimizing the SCSI bus arbitration.

A final key point is that because the Data Gateway disk controller function is packaged separately from the Serial Disk, it may be replaced with new disk controllers if your Serial Disk requirement change.

Now let us take a closer look at the flexible connectivity offered by IBM SAN Data Gateway for Serial Disk.

Speaker Note:
IBM StorWatch SAN Data Gateway Specialist provides management support for the IBM 2108-G07 SAN Data Gateway.

Foil 52
Instant Copy Process

This chart depicts the instant copy process for the SAN Data Gateway for Serial Disk. In picture 1) in the upper left, we show three disk in each group with two-way disk mirroring and an Instant Copy disk. Notice that in normal operations, the host server sees only one disk image - shown as disk **(A)** in the first group. The second disk **(A')** is a second copy with the same address. The third disk **(1)** is the Instant Copy disk and is not now visible to the host server.

In picture 2) on the upper right, we have separated the Instant Copy disk from its disk mirroring group. Notice that the Instant Copy disk **(1)** now has a separate ID and LUN address so that it is usable by any host server attached to the Serial Loop. The user can now initiate a separate tape backup job, which can run concurrently with the online applications which continue to operate in parallel.

In picture 3) the backup process has completed and the Instant Copy disk has been added back into the disk mirroring group. Notice that it is now out of sync with disk **(A)** since the mainline applications has been updating disk **(A).**

In picture 4) the disk rebuild process, also used for disk sparing, is now activated to resync the Instant Copy disk.

The StorWatch SAN Data Gateway S20 Specialist provides a user interface to manage the process.

The key point is that Instant Copy can reduce the batch window and improve online service.

Foil 54
StorWatch Data Gateway S20 Specialist

The StorWatch Data Gateway S20 Specialist supports the IBM 2108 Model S20 SAN Data Gateway for Serial Disk. Please remember that the IBM SAN Data Gateway (for SCSI-attached storage) is supported by the StorWatch Data Gateway Specialist - without the S20 designation.

The S20 Specialist can operate on either a Windows NT server attached to the Serial Loop or can operate on a Window NT client which is remotely LAN attached to the Host Server. A Deamon application runs on the NT server to provide LAN connection to the remote Windows NT client.

We have discussed the role of the S20 Specialist in creating composite disks, establishing mirrored sets of disks and providing Instant Copy management.

Sample procedures and scripts are provided in the S20 Specialist documentation.
It is important to understand that the user is responsible to customize these scripts and to provide the tape backup procedures and applications to use the Instant Copy disks. This may be a great IBM or Business Partner Service Offering opportunity.

The S20 Specialist also provides monitoring, reporting and configuring support of the Serial Disk System. This eases the management and service for all supported UNIX and Windows NT servers.

Speaker Notes:
The footnote shows that currently, clustered Windows NT servers are not supported. Also, Windows NT servers require an outage to add or remove a new device such as an Instant Copy disk.

This is not as limiting as it might appear. Many customers prefer to have a separate tape library backup server dedicated to disaster recovery processing. In this case, Clustered Windows NT server Instant Copy disks can be separated from their disk mirroring group, acquired by the library server, copied to tape and returned its mirroring group without an outage.

Foil 55
Data Gateway Performance Considerations

The Data Gateway maximizes the use of host server Ultra SCSI 40 MB/s bandwidth potential. It appears as a single SCSI target, which minimizes the SCSI arbitration overhead. In many cases, the host server bus speed may become a bottleneck.

The Data Gateway provides high performance with scalable SSA 80 Serial Loop architecture. Up to 35 MB/sec and 5,100 IOPS with sequential reads and writes has been measured by Vicom Systems with IBM test guidelines.

SSA 80 provides full duplex serial disk loop performance.

With multiple hosts, multi-domain SSA spatial reuse can multiple individual Data Gateway performance up to eight times!

Note that 7133 Advanced Models D40 and T40 operate at SSA 80, not SSA 160 speeds because all of the disk reads and writes are initiated by the Data Gateway.

Instant Copy has some special performance factors to consider.

- The Instant Copy disk supports parallel backup processing as soon as it is split off from its disk mirroring group and accessed by the backup host server.
- Since the disk rebuild process is used to resync the Instant Copy disk, users need to understand that a significant amount of time may be required and the resync process may potentially impact other application processing. This process should be scheduled during off-peak processing periods.

The bottom line is that the SAN Data Gateway disk controller is ideally suited for parallel processing applications.

Foil 56
Separate Disk Controller and Disk Drawers

The purpose of this chart is to address the investment protection question. A major customer concern is "What if my storage requirements change?" and "What if my dedicated disk storage is no longer required and I need to consolidate my disk storage resources in the future?"

The picture at the left shows a disk controller with four Data Gateways in a Rack Mount Enclosure. It shows that the disk controller is separate from the Serial Disk Drawers. This solution provides the ultimate high performance, high availability, high function, and low cost solution for today's application requirements.

The picture on the right shows IBM Versatile Storage Server attached to UNIX and Windows NT servers with Fibre Channel adapters. Notice that the Serial Disk can be consolidated into the Versatile Storage Server if Fibre Channel connectivity is required in the future.

Since Seascape 7133-020 and 7133-D40 Serial Disk are used, they can easily be migrated to a Seascape Versatile Storage Server in the future if storage consolidation become important.

The key point is that your Serial Disk storage investment will be protected.

Speaker Note:
Using May 4, 1999 US announcement list prices for Data Gateway of $6,300, Rack Mount Enclosure of $700, and 7133 Advanced Model D40 36 GB Serial Disk of $6,300 and 16 disks per loop, one can see that the Data Gateway disk controller cost is only about 5 percent of the total Disk Storage System price (less than $0.20/MB for unmirrored system including about $0.01/MB for the Data Gateway). Please check you local country announcement list prices and discount structure.

Additional Information

For specific details about supported platforms, operating system releases, adapters, and configurations, visit the SAN Data Gateway customer accessible Internet site: www.ibm.com/storage/SANGateway. .

Foil 57
VICOM Storage Solutions

This foil shows a subset of the different platforms that VICOM supports. VICOM does tend to support a larger list of platforms than IBM, so if you run into an opportunity for a UNIX or NT platform that IBM does not support, check with VICOM to see if they can be of assistance. For a more complete list, refer to www.ibm.com/storage. The source for this information is VICOM.

Foil 58
IBM Serial Storage Architecture

This is a summary chart which summarizes the points that were made in this presentation.

SSA is not proprietary, thereby providing customers with more options and future investment protection as their applications grow and requirements change. The SSA architecture enables storage solutions for a wide range of applications and end user environments from data mining to transaction processing and even video. IBM SSA solutions also provide a unique "buy and grow" design. Allowing disk systems to move from simple to complex storage environments simply by changing a low cost adapter preserves the customers major investment in storage. IBM SSA is extremely well positioned in the mainstream of the storage trends.

Foil 59
Trademarks

Self explanatory.

IBM Storage Systems Division
Department DKSA
5600 Cottle Road
San Jose, CA 95193
4/99