Enterprise Storage Server Fibre Channel Attachment Version 5.01

Phil Mills

International Business Machines Corporation
Storage Subsystem Strategy and Futures
H65A/050-2
SSG - San Jose
408-256-1486
millsp@us.ibm.com

December 15, 2000

Table of Contents

Preface . . . 3
Introduction to Fibre Channel . . . 4
  Fibre Channel Topologies . . . 4
  Fibre Channel Ports . . . 10
  Fibre Channel Layers . . . 11
  Fibre Channel Classes of Service . . . 14
  Fibre Channel Physical Transmitter and Cable Technologies . . . 15
  Fibre Channel Storage Area Networks (SANs) . . . 15
  Fibre Channel Futures . . . 16
Hosts, Operating Systems, and Fibre Channel Host Bus Adapters Supported . . . 18
Some Differences Between Fibre Channel and SCSI Host Attachment in ESS . . . 19
  LUN Affinity . . . 19
  Targets and LUNs . . . 20
  Assignment of LUNs . . . 20
  LUN Access Modes . . . 23
Considerations for New Fibre Channel Installs . . . 24
  Fibre Channel Host . . . 24
  ESS Fibre Channel Adapters . . . 24
  Fibre Channel Network (a.k.a. Storage Area Network or SAN) . . . 25
  Fiber Type . . . 25
  Configuring for Long Distances . . . 25
  Multipathing for Availability and Performance . . . 27
  Data Sharing . . . 28
Installing Fibre Channel Host Attachment To ESS . . . 29
Migrating From SCSI Hosts To Fibre Channel Attachment . . . 31
Migrating From SAN Data Gateway To Fibre Channel Attachment . . . 32
Performance Considerations . . . 33
  Sustainable Bandwidth . . . 33
  Total Number of Ports . . . 33
  Port Positions in ESS . . . 33
  Loop Versus Fabric . . . 34
  Performance Comparison: Native SCSI vs SDG vs Native Fibre Channel . . . 34
Summary . . . 35
References and Web Sites . . . 36
Notices . . . 38

Preface

The purpose of this paper is to briefly describe what the Fibre Channel interface is all about, and to explain how this new interface is used in the Enterprise Storage Server (ESS). The intent is to give the ESS users who are responsible for planning and configuring customer storage a feel for how they should best plan for using this new attachment in their environment once it is made Generally Available. It will also guide users in understanding (1) the process for migrating from today's parallel SCSI interfaces to the new Fibre Channel interfaces, as well as (2) the process for migrating from the use of the SAN Data Gateway to native Fibre Channel.

Note: The document containing the details of the announcement for Fibre Channel attachment in ESS is the IBM Hardware Announcement 100-089, dated March 28, 2000, and entitled: "IBM Enterprise Storage Server Multiplatform Storage Solution Provides Superior Performance, Function, Scalability, and Availability."

Introduction to Fibre Channel

Fibre Channel is a 100 MB/sec, full-duplex, serial communications technology used to interconnect input/output (I/O) devices and host systems that can be separated by tens of kilometers. It incorporates the best features of traditional I/O interfaces, like throughput and reliability found in SCSI and PCI, with the best features of networking interfaces, like connectivity and scalability found in Ethernet and Token Ring. It provides a new transport mechanism for the delivery of existing commands, and provides an architecture that achieves high performance by allowing a significant amount of processing to be performed in hardware. It can operate with legacy protocols and drivers like SCSI and IP, enabling it to be introduced easily into existing infrastructures.

Fibre Channel transfers information between the sources and the users of the information. This information can include commands, controls, files, graphics, video, and sound. Fibre Channel connections are established between Fibre Channel ports residing in I/O devices, host systems, and the network interconnecting them. The network consists of elements like switches, hubs, bridges, and repeaters, that are used to interconnect the Fibre Channel ports.

Before going into the description of the Fibre Channel topologies, it will be beneficial to discuss the Fibre Channel addressing. Fibre Channel defines a port address that is used to route frames. Each frame that is sent from one Fibre Channel component to another includes the source and destination addresses, each 3 bytes in length. These addresses are divided into three 1-byte fields, as follows:

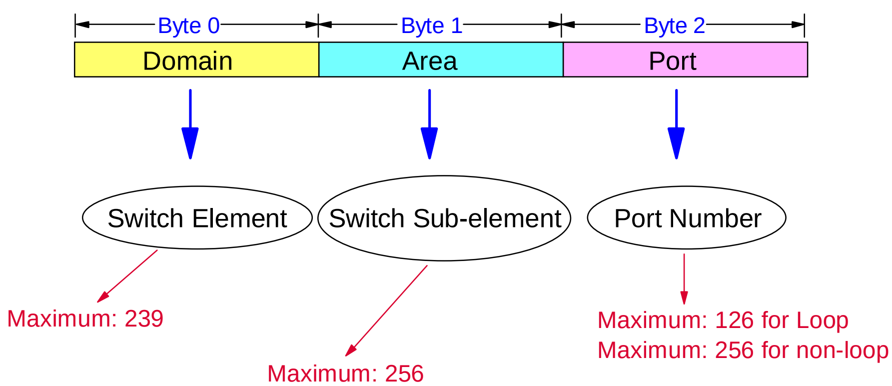

Figure 1. Fibre Channel Addressing Assignment

Therefore, architecturally, you could have 239x256x256=15.6 million nodes in the maximum configuration without loops. However, current implementations use the Area byte for the port number on the switch and then use the Port byte for the address of the node attached to the loop. This results in reducing the maximum nodes to 239x256=61.2 thousand when no loops are configured, and to 239x256x126=7.7 million nodes when loops are configured on every switch port.

Fibre Channel Topologies

There are three Fibre Channel topologies defined in the Fibre Channel architecture. These are Point-to-Point, Switched Fabric, and Arbitrated Loop.

- Point-to-Point

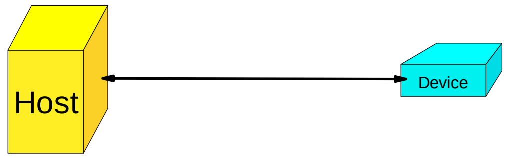

Figure 2. Point-to-Point Topology

This is the simplest Fibre Channel topology. It allows exactly two Fibre Channel end points (i.e., I/O device or host system) to be directly connected via a Fibre Channel cable. This topology can satisfy only the very basic configuration requirements. It does not provide sufficient connectivity to support complex configurations, but it does support the maximum bandwidth capability of Fibre Channel.

ESS fully supports the point-to-point topology. It supports this topology with what has come to be called its point-to-point protocol. This protocol is also used by ESS when directly attached to a fabric in a switched fabric topology (see next topology section).

- Switched Fabric

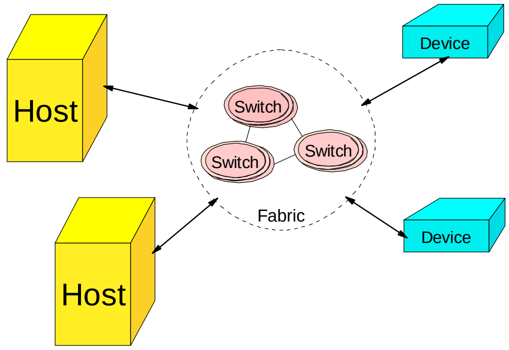

Figure 3. Switched Fabric Topology

In this topology two or more Fibre Channel end points are interconnected through one or more switches. Each switch can support up to 256 ports (per the Fibre Channel architecture), but most switches today support only 8, 16, 32, or 64 ports. When multiple switches are interconnected, they are said to be cascaded. The set of cascaded switches is commonly called a fabric. The term fabric is used to describe a routing structure that receives addressed information and routes it to its appropriate destination. A fabric may consist of one or more switches, functioning as a single routing mechanism. As mentioned earlier, the architected maximum number of end points that may be connected in a fabric is somewhat less than 8 million when maximum use is made of loop configurations, and somewhat less than 16 million when no loops are used. Most fabrics in existence today, though, support only 10s to 100s of ports. The Fibre Channel end points have no awareness of the internal structure of the fabric. Fibre Channel architecture does not support the cascading of independent fabrics where each fabric contains its own separate routing mechanism. The addressing structure within Fibre Channel architecture will only support a single fabric router.

In addition, Fibre Channel architecture currently does not have the appropriate standards in place to guarantee interoperability between switches from different vendors. These standards are being aggressively pursued and should be available soon. Until then, though, most fabrics will be built using switches from a single switch vendor.

Switches (or switched fabrics) also include a function commonly called *Zoning*. This function allows the user to partition the switch ports into port groups. The ports within a port group, or zone, can only communicate with other ports in the same port group (zone). By using zoning, the I/O from one group of hosts and devices can be completely separated from that of any other group, thus preventing the possibility of any interference between the groups. One example where zoning might be used is in an environment where both a production system and a test system coexist. The customer will want to make sure that his test system I/O does not impact his production system I/O. He can accomplish this very simply by setting up two zones within the fabric, one for the hosts and devices in the production system, and the other for the hosts and devices in the test system. He could, of course, accomplish the same result by implementing two separate switch fabrics, but that would be more expensive and also require more complicated configuration management. This zoning just described is known as hard zoning, since it is enforced by the switch. There is another kind of zoning, called soft zoning, that can be enabled as well. It is not enforced by the switch, but rather operates on the honor system. The way this zoning works is that the user assigns nodes to a zone according to the node's World Wide Name - either the World Wide Port Name (WWPN) or the World Wide Node Name (WWNN). This information is captured by the name server, which is a function embedded within the switch. Then, whenever a port communicates with the name server to find out to which nodes it is allowed to connect, the name server will respond only with the nodes that are within that port's zone. Since the standard Fibre Channel device drivers do communicate with the name server in this manner, this type of zoning is adequate for most situations. However, it is possible that a device driver could be designed that would attempt to access nodes not in its list of allowed connections. If this occurred, the switch would neither prevent nor detect the violation.

ESS fully supports the switched fabric topology. As stated previously, it supports this topology with its point-to-point protocol. Through the use of switched fabrics, ESS can attach to multiple Fibre Channel hosts via a single Fibre Channel adapter. ESS supports connections to a maximum of 128 hosts per Fibre Channel port, and a maximum of 512 hosts across all Fibre Channel and SCSI ports configured on one machine.

In switched configurations, because of possible interactions among the Fibre Channel adapters from one or more hosts, it is recommended that zones be established which contain the desired number of ESS ports and only a single host port. These zones can be created as either soft zones or hard zones. The ESS ports can be members of multiple zones, so that zones can be established with each host port to allow the hosts to have shared access to the ESS ports. Creating the zones in this manner will prevent any unwanted communication between the host's ports.

In general, fabrics allow all of its ports to use the full bandwidth of the Fibre Channel architecture simultaneously. This provides for a very high performance interconnection topology.

- Arbitrated Loop

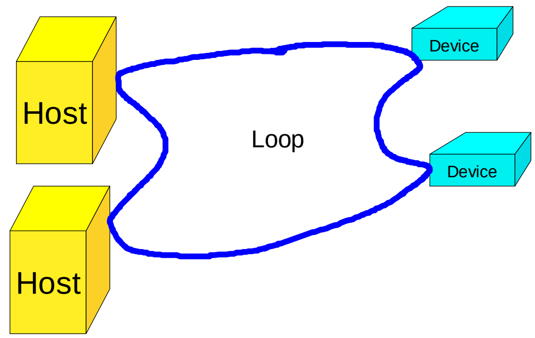

Figure 4. Loop Topology

In this topology two or more (up to a maximum of 126) Fibre Channel end points are interconnected via a looped interface. Information is routed around the loop and repeated by intermediate ports until it arrives at its destination. The Fibre Channel ports that support this topology must contain these routing and repeating functions in addition to all the functions required by the Point-to-Point ports. This topology is called Fibre Channel - Arbitrated Loop, and is often referred to by its acronym FC-AL. All end points share the FC-AL interface and therefore also share the bandwidth of the interface. Only one active connection may be ongoing at a time.

A very commonly used variation on this topology is the Fibre Channel hub, shown in the following figure.

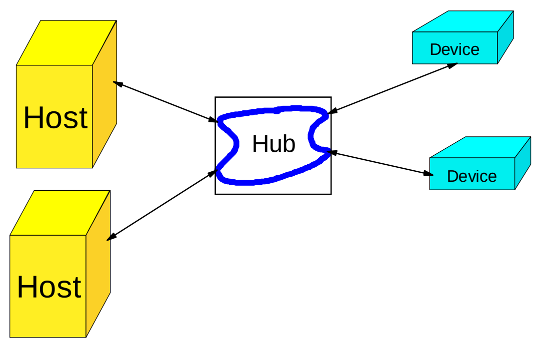

Figure 5. Loop Topology With Hub

A hub incorporates the structure of FC-AL in a package that provides ports physically similar to those of a switch. This allows more centralized cabling and, due to added functionality provided by hubs, increases the reliability and availability of the Arbitrated Loop topology. It does not, however, improve the performance of the loop, since it is still just a single loop that has been repackaged. Today's hubs generally have 8, 16, or 32 ports. Additional ports can be configured by cascading hubs together, but the resulting configuration is still a single loop.

One important point to remember concerning loops, whether or not a hub is utilized, is that the loop goes through an architected Loop Initialization Process (LIP) each time the loop is broken and reconnected or whenever a host or device is added to or removed from the loop. This includes whenever hosts or devices that are attached to the loop are powered on or off. This LIP disrupts any I/O operations currently in progress. Thus, if you have multiple hosts on a loop, then whenever any host is booted up it will cause a LIP and therefore disrupt any ongoing I/O. For this reason it is normally recommended to only have a single host on any loop. Devices have the same effect, but they are generally not booted very often. It is therefore quite common to have multiple devices on a loop.

ESS currently does not support the arbitrated loop topology (when there are more than two Fibre Channel components in the loop). It supports attachment to a hub only as a distance extender where it is attached to a single host. No other devices or hosts may be attached to the hub.

A common variation on the loop topology is the point-to-point loop, as shown in the following figure:

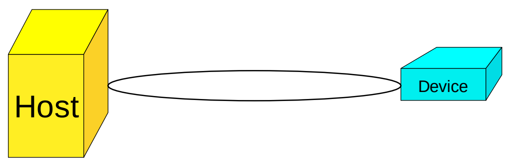

Figure 6. Point-to-Point Loop Topology

This topology is very similar to the point-to-point topology. There is only one I/O initiator (host) and one I/O target (device) on the loop, just like in the point-to-point case. The difference, though, is in the protocol used to communicate between the initiator and the target. In the point-to-point topology described previously, the point-to-point protocol was used. In this topology, the loop protocol is used. These protocols are significantly different from each other. The point-to-point protocol is architected for the point-to-point and fabric topologies. The loop protocol is architected for a point-to-point topology with exactly two Fibre Channel end points connected together, and also for a loop topology with anywhere from 3 to 126 Fibre Channel end points connected together.

Therefore, when just two Fibre Channel end points are connected together, either the point-to-point protocol or the loop protocol can be utilized, but both end points must use the same protocol. Most host Fibre Channel adapters will default to using the loop protocol whenever they are not directly connected to a fabric. In these cases the ESS adapter must also be configured for the loop protocol. And, of course, when directly connected to a fabric, the adapters will use the point-to-point protocol.

Combining Fabric and Loop Topologies

In addition to the three basic Fibre Channel topologies Point-to-Point, Fabric, and Loop, it is also possible to have combinations of these where fabrics and loops exist together, as shown in the following figure. In these configurations, the Fibre Channel hosts and devices *attached to the loop* can operate in either one of two ways. If they support what is called the "Public Loop Profile", then they can communicate with any of the other hosts and/or devices attached to either the loop or the fabric. However, if they support what is called the "Private Loop Profile", then they can only communicate with other hosts and/or devices attached to the loop and can't communicate to any of the those attached to the fabric. These two different operational modes apply *only* to the hosts and devices attached to the loop.

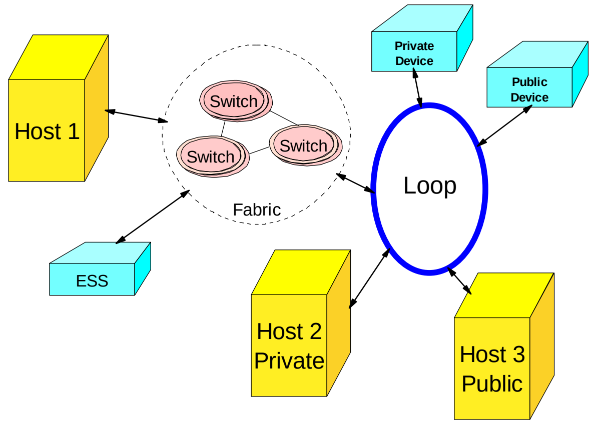

Figure 7. Fabric-Loop Combination

In the configuration shown above, the ESS is attached to the fabric and therefore utilizes its point-to-point protocol. It can successfully communicate with Host 1, which is also attached to the fabric, and with Host 3, which is a public-loop host attached to the loop. It cannot, however, communicate with Host 2, since it is a private-loop host attached to the loop. For the two devices on the loop, the public-loop device can communicate with all three hosts, but the private-loop device can only communicate with Hosts 2 and 3.

There are a few exceptions to what was just described. First, some switches have implemented a function called Address Translation, where the switch port that is attached to the loop will translate the addresses for private-loop devices so that they can communicate to other hosts attached to the switch. This address translation works in both directions with the switch translating the private-loop device's 1-byte Arbitrated Loop Physical Address (AL_PA) to a 3-byte fabric address for the hosts, and then translating the host's 3-byte fabric address to a 1-byte AL_PA for the private-loop device. This address translation only works for private-loop devices, and specifically does not work for private-loop hosts. The IBM 2109 Fibre Channel switch supports this address translation function. In the preceeding figure, if the fabric is composed of IBM 2109 switches, then the private-loop device would be able to communicate with Host 1 as well as Hosts 2 and 3 due to the address translation performed by the switch.

Another function implemented by some switches is called Quickloop. This function does for the private-loop hosts essentially what the Address Translation function does for the private-loop devices. That is, it allows the private-loop hosts to attach to devices through a switch. The way this function works is that the switch "joins" the looplets (loops attached to switch ports) from all ports that have been Quickloop-enabled into a single arbitrated loop. All hosts and devices on any of these "looplets" can communicate with each other just as though they were connected into a Fibre Channel hub. The only real difference is that the Quickloop will allow multiple connections to take place at the same time, as long as they are on different looplets. The IBM 2109 Fibre Channel switch supports Quickloop as a priced feature. In addition, the IBM 3534-1RU Managed Hub, supported by Netfinity servers for attachment to low and mid-range storage, is basically a low cost switch with all ports Quickloop enabled. This product is sometimes referred to as a loop switch.

Another point to remember is that a loop may only be connected to one switch at a time. When a switch is connected to a loop, the switch port is automatically assigned the task of "loop master". Each loop can have only one "loop master", so each loop can have only one active switch attachment.

If two switches are connected simultaneously to a single loop, the switches will automatically detect this, and one of the switch ports will remain inactive.

Fibre Channel Ports

There are five basic kinds of ports defined in the Fibre Channel architecture, as well as some vendor-specific variations. The five basic ports are as follows:

- **Node Ports, N_ports**
  These ports are found in Fibre Channel Nodes, which are defined to be the source or destination of Information Units (IUs). I/O devices and host systems interconnected in point-to-point or switched topologies use N_ports for their connections. N_ports can only attach to other N_ports or to F_ports. The ESS Fibre Channel adapters support the N_port functionality when connected directly to a host or to a fabric.

- **Node-Loop Ports, NL_ports**
  These ports are just like the N_ports described above, except that they connect to a Fibre Channel Arbitrated Loop (FC-AL) topology. NL_ports can only attach to other NL_ports or to FL_ports. The ESS Fibre Channel adapters support the NL_port functionality when connected directly to a loop.

- **Fabric Ports, F_ports**
  These ports are found in Fibre Channel Switched Fabrics. They are not the source or destination of IUs, but instead function only as a “middleman” to relay the IUs from the sender to the receiver. F_ports can only attach to N_ports. The ESS Fibre Channel adapters do not support the F_port functionality, which is found only in fabrics.

- **Fabric-Loop Ports, FL_ports**
  These ports are just like the F_ports described above, except that they connect to an FC-AL topology. FL_ports can only attach to NL_ports. The ESS Fibre Channel adapters do not support the FL_port functionality, which is found only in fabrics.

- **Expansion Ports, E_ports**
  These ports are found in Fibre Channel Switched Fabrics and are used to interconnect the individual switch or routing elements. They are not the source or destination of IUs, but instead function like the F_ports and FL_ports to relay the IUs from one switch or routing element to another. E_ports can only attach to other E_ports or to B_Ports. The ESS Fibre Channel adapters do not support the E_port functionality, which is found only in fabrics or hubs.

A couple of other port types are as follows:

- **Combination Ports, G_Ports**
  These ports are sometimes found in Fibre Channel Switched Fabrics and are used either as E_Ports, when the link is connected to another switch or a bridge device (such as Fibre Channel Back Bone devices), or as F_Ports, when the link is connected to an N_Port for a host or device. This port

automatically determines what mode to run in after determining what it is connected to.

- **Universal Ports, U_Ports**
  These ports are sometimes found in Fibre Channel Switched Fabrics and are used either as E_Ports, when the link is connected to another switch or a bridge device (such as Fibre Channel Back Bone devices), as F_Ports, when the link is connected to an N_Port for a host or device, or as FL_Ports, when the link is connected to an Arbitrated Loop. This port automatically determines the mode in which to run after determining what it is connected to.

- **Bridge Ports, B_Ports**
  These ports are found in bridge devices and are used to connect to a E_Ports on a switch. The ESS Fibre Channel adapters do not support the B_Port functionality.

The Fibre Channel architecture specifies the link characteristics and protocols used between Fibre Channel Ports. The ESS Fibre Channel adapters support the N_port functionality when directly connected to a host or a fabric, and the NL_port functionality when directly connected to a loop. They do not support the functionality of any of the other port types.

Fibre Channel Layers

Fibre Channel has five layers, called FC-0, FC-1, FC-2, FC-3, and FC-4. FC-0 through FC-2 define how Fibre Channel ports interact through their physical links to communicate with other ports. These levels are described in detail in FC-PH, the primary Fibre Channel Specification. FC-2 through FC-4 define how Fibre Channel ports interact with applications in host systems. FC-3 and FC-4 implementations are described in separate documentation specific to the protocol implemented.

- FC-0: Media and Interfaces
  FC-0 defines the media and interface characteristics of the serial links between ports. Copper, multi-mode fiber, and single-mode fiber are the supported media types. Both short wave and long wave laser transmitters are used as the fiber light source. Signals are continuously transferred across the link to keep the transmitter and the receiver in constant synchronization. The transmitter clock uses a crystal oscillator, while the receiver recovers this clock using a phase-locked loop.
  The most common signaling rate today is 1.065 Gbits/sec, producing a throughput data rate of 100 MB/sec after accounting for protocol overheads. This signaling rate can scale up, at least by one order of magnitude to 10 Gbits/sec, without affecting upper-level protocols.
  The ESS Fibre Channel adapters, feature code 3022 in ESS 2105, use short wave laser transmitters, attach to multi-mode fiber only, and have a 1.065 Gbit/sec signaling rate.

- FC-1: Transmission Protocol
  FC-1 specifies IBM's 8b/10b coding scheme for converting eight-bit data bytes into ten-bit transmission signals, or packets, as well as specifying the special character sequences called Ordered Sets which are used for a variety of control functions.
  The ten-bit packets defined in FC-1 are balanced, meaning the number of 1s and 0s transmitted is balanced within one bit at word boundaries. The 256 characters of the eight-bit ASCII data space are converted to a 512 character subset of the 1024 possible characters in the ten-bit space. Each character has two different transmission codes assigned to it, with differing numbers of 1s and 0s. The transmitter keeps track of the difference in the number of 1s and 0s and maintains the required

balance. The receiver also tracks the number of 1s and 0s and indicates a disparity error whenever they get out of balance.

The Ordered Sets are composed of characters in the remaining 512 character subset in the ten-bit space. These character sequences are used for frame delimiters, Idle characters, Arbitration controls, Close controls, Open controls, and Flow controls. In addition, they are also used to perform link level protocols such as Not Operational Sequence (NOS), Off Line Sequence (OLS), Link Reset (LR), Link Reset Response (LRR), Loop Initialization Process (LIP), Loop Port Bypass (LPB), and Loop Port Enable (LPE).

The FC-1 layer generates a Bit Error Rate (BER) of less than 1 bit error in 10^{12} bits. This is approximately three orders of magnitude better than SCSI or Ethernet. It accomplishes this low BER with low-cost, simple, reliable transceivers made possible primarily by the 8b/10b coding scheme, which eliminates DC and virtually all transmission frequencies below twenty percent of the clock rate, thereby significantly reducing hysteresis jitter at the receiver and making accurate detection simpler. This also guarantees a signal transmission frequency high enough to allow a receiver to easily recover the clock with a phase-locked loop circuit.

The FC-1 transmission protocol is independent of the media, the distance, or the data rate, enabling systems to easily scale across a wide range of requirements.

There are no variations in FC-1 support. ESS Fibre Channel adapters fully support the FC-1 layer, as this is a requirement of all Fibre Channel adapters.

- FC-2: Signaling Protocol

FC-2 specifies the basic control mechanisms within Fibre Channel, such as the classes of service, flow control, loop arbitration, and the different kinds of logins that are used. It also specifies the message constructs of the information transfers. This consists of *words*, *frames*, *sequences*, and *exchanges*.

Words are the basic message unit used in information transfers between N_ports. When information is not being transferred between N_ports, special "fill-words" are transferred to keep the transmitter and receiver in constant synchronization, as mentioned before under FC-0. Words are four bytes in length.

Frames are the next higher level message unit and are composed of multiple words. Frames are the data containers for Fibre Channel. They are composed of a special Start-Of-Frame ordered set, a six word header, a payload of 0 to 528 words, a CRC word, and a special End-Of-Frame ordered set. The frame header is used to route frames to their appropriate destinations and to enable the correct association of each frame with its sequence and exchange. The frame payload contains the actual information being transferred. The CRC word validates the contents of both the header and payload.

Sequences are the next higher level message unit and are composed of one or more frames carrying related information. A sequence is unidirectional, from source to destination.

Exchanges are the highest level message unit and are composed of one or more sequences carrying related information. An exchange is bi-directional and has an originator (the port that initiates the exchange) and a responder (the port that responds). In the first frame of the first sequence of the exchange, the originator assigns an Originator eXchange IDentifier (OX_ID). Only the originator needs this value. All frames in this exchange transmitted by the responder will contain this OX_ID. The responder assigns a Responder eXchange IDentifier (RX_ID), meaningful only to the responder, and transmits it in every frame of the exchange. After the originator receives its first frame from the responder, it then uses the responder's RX_ID. From that point forward all frames of the exchange contain the assigned OX_ID and RX_ID values. These IDs can then be used to immediately look up all context information about the exchange. This allows frames from different exchanges to be successfully interspersed with each other, allowing very efficient multiplexing to take place.

The following figure shows the relationships just described:

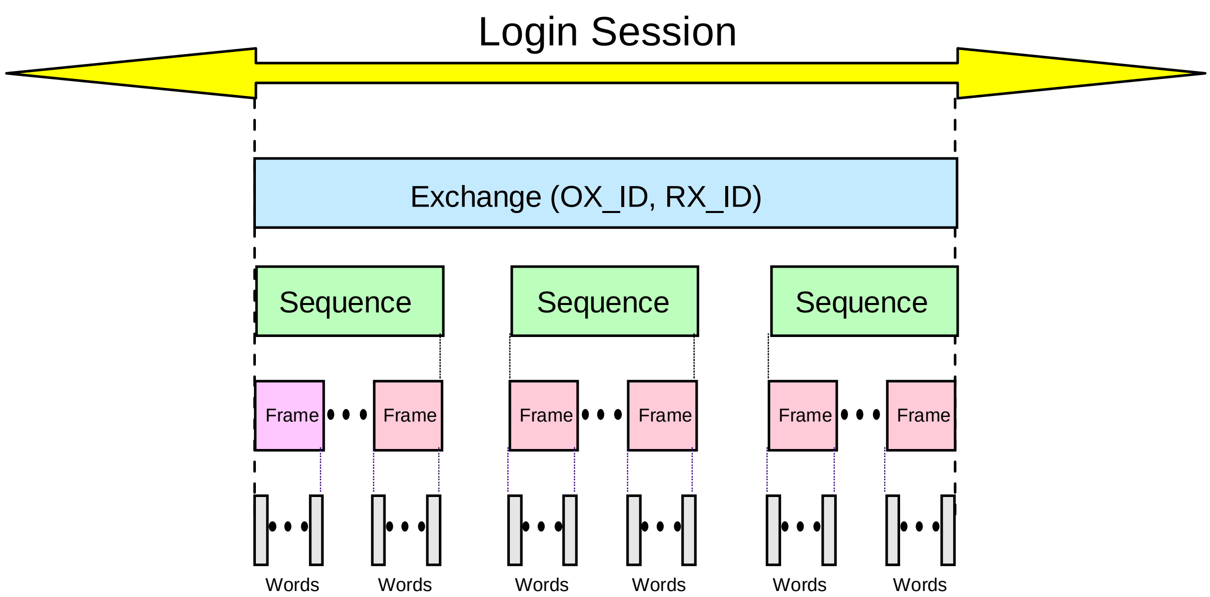

Figure 8. Words, Frames, Sequences, and Exchanges

There are no variations in FC-2 support. ESS Fibre Channel adapters fully support the FC-2 layer, as this is a requirement of all Fibre Channel adapters.

- <u>FC-3: Common Services</u>

  FC-3 is used for defining common services provided by two or more N_ports in a host system. One example is a Hunt Group (special form of multipathing), a group of two or more N_ports in a host system that share a common port address and thereby increase the bandwidth available from port to fabric. There are no FC-3 services currently defined which require any special support by ESS Fibre Channel adapters.

- <u>FC-4: Protocol Mappings</u>

  FC-4 is used to specify the Upper Level Protocol (ULP) that runs on the Fibre Channel interface. There is a unique FC-4 definition for each protocol that gets mapped onto Fibre Channel. The most common protocol in use today is called Fibre Channel Protocol (FCP) and is the SCSI protocol mapped to Fibre Channel. ESS fully supports FCP. Many other protocols are either already documented or are currently in progress. FC-SB-2 is one such protocol. It is being developed by IBM and NCITS (National Committee for Information Technology Standards) and will be supported by ESS in the near future. IBM's implementation of FC-SB-2 is called FICON.

<u>Fibre Channel Classes of Service</u>

There are three basic classes of service, with at least three other variations, defined in the Fibre Channel architecture. The three basic classes of service are:

- Class 1
  Class 1 is a connection-oriented class of service that establishes a dedicated connection with full bandwidth between a pair of N_ports. It provides confirmation of delivery or nondelivery (i.e., notification of nondelivery). This class of service is very similar to that used previously in ESCON. However, because it prevents the efficient use of multiplexing between ports, it is rarely used in Fibre Channel configurations today. ESS Fibre Channel adapters do not support Class 1 service for either the FCP or the FICON protocols.

- Class 2
  Class 2 is a connectionless class of service that provides confirmation of delivery or nondelivery. No bandwidth is allocated or guaranteed. In a fabric topology, it is the responsibility of the fabric itself to provide the necessary control required to prevent bandwidth starvation of ports. Since no connection is established, frames may be sent to different destinations, one after the other. The fabric routes each frame to its appropriate destination. It is also possible for a port that is transmitting a frame to one location to be simultaneously receiving a frame from a different location.
  ESS Fibre Channel adapters will operate in Class 2 mode when required to do so in some initialization sequences, but otherwise operates only in Class 3 mode.

- Class 3
  Class 3 is a connectionless class of service that does not provide confirmation of delivery or nondelivery. If a frame cannot be delivered, it is discarded without notification. No bandwidth is allocated or guaranteed. As in class 2 service, in a fabric topology it is the responsibility of the fabric to prevent bandwidth starvation of ports.
  ESS Fibre Channel adapters fully support Class 3 operations. Both FCP and FICON, the two ULPs being implemented in ESS, use class 3 service a majority of the time.

Some other classes of service that have been defined are as follows:

- Class 4
  Class 4 is a connection-oriented class of service that provides a virtual circuit between a pair of N_Ports with guaranteed bandwidth and latency, and with confirmation of delivery and nondelivery. During a Class 4 operation, all of the resources necessary to provide the guaranteed bandwidth are reserved for that virtual circuit.

- Class 6
  Class 6 is a derivative of Class 1 and is therefore connection-oriented. It provides a one-to-many multicast service with confirmation of delivery and nondelivery.

- Class F
  Class F is a connectionless class of service with notification of nondelivery between E_Ports. It is used for control, coordination, and configuration of the fabric by switches communicating across

Inter-Switch Links (ISLs).

Fibre Channel Physical Transmitter and Cable Technologies

Fibre Channel architecture supports 1 Gbit/sec links (extendible to higher bit rates) in both single-mode and multi-mode optical fiber. The single-mode fiber is only offered in one size - 9 micron. The multi-mode fiber comes in 2 sizes: 50 micron and 62.5 micron.

Fibre Channel architecture supports both short wave laser and long wave laser optical transmitter technologies. As mentioned before these transmitters today operate at 1 Gbit/sec, but technology improvements will allow higher bit rates to be achieved. The short wave laser technology uses a wavelength of 780 nanometers and is only compatible with multi-mode fiber. A maximum distance of 500 meters is supported with 50 micron fiber; 175 meters is supported with 62.5 micron fiber. The transmitter used in ESS's Fibre Channel adapter with feature code 3022 is the short wave laser. IBM will supply a 31 meter 50 micron fiber cable with every ESS short wave laser Fibre Channel adapter. Information relative to ordering fiber cables for ESS can be found at the following web site:

http://www.as.ibm.com/asus/connectivity.html

The long wave laser technology uses a wavelength of 1300 nanometers. It is compatible with both single-mode and multi-mode fiber. With single-mode fiber, a maximum distance of 10 kilometers is supported. With multi-mode fiber approximately 500 meters is supported.

The most important points to remember about the Fibre Channel transmitters are that (1) a fiber link must have the same type of transmitter at each end (i.e., short wave or long wave), and (2) the protocol that flows across the link is completely independent of the transmitter type. In addition, all fiber optic patch panels and cable connectors have an associated light loss which, when incorporated in the customer's Fibre Channel cabling configuration, will decrease the maximum distances capable between the transmitter and receiver.

Fibre Channel Storage Area Networks (SANs)

SANs are networks that connect storage devices to host servers. They are most commonly built upon the Fibre Channel technology as a networking infrastructure. What differentiates SANs from previous interconnection schemes is the basic concept that all (or mostly all) of your storage can be consolidated in one large "storage area" that allows centralized (simplified) management in addition to any-to-any connectivity between host servers and the storage. ESCON can be thought of as the first real SAN. It provided the kind of connectivity commonly found in SANs. However, it was primarily restricted to ESCON hosts and devices. Fibre Channel SANs, though, have the potential to allow the interconnection of open systems and storage (i.e., non-S/390) in the same network as S/390 systems and storage. This is possible because the protocols for both open attachment and S/390 attachment are being mapped to the FC-4 layer of the Fibre Channel architecture. An example of this SAN is shown below:

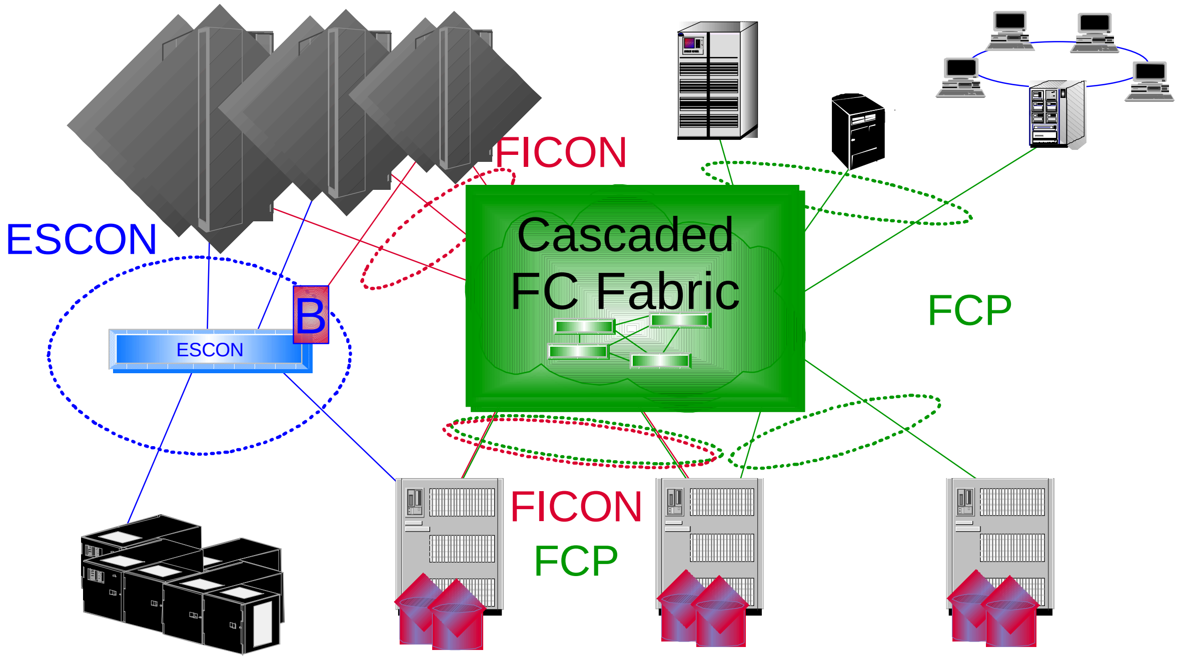

Figure 9. Sample SAN Configuration

For more information on SANs, please visit the following IBM web site:

http://www.storage.ibm.com/ibmsan/basics.htm

Also, there is an excellent IBM Redbook that was published in May, 2000. It is called "Designing an IBM Storage Area Network", document number SG24-5758. It can be found at the following web site:

http://www.redbooks.ibm.com

Fibre Channel Futures

Looking ahead to the next 2-5 years, we can expect that Fibre Channel will incorporate bit rate improvements to 2, 4, and 10 Gbit/sec. *(Note: 4 Gbit/sec is questionable since the technologists currently believe it makes more sense to skip 4 and go directly to 10 Gbit/sec).* The intent is that these higher bit rates will interoperate with today's 1 Gbit/sec technology by introducing a negotiation function at link initialization where the Fibre Channel end points "negotiate" to operate at the fastest bit rate capable by both transmitters.

In addition, the capabilities in fabrics will make some major improvements. Specifically, the number of ports supported by a fabric will increase to meet customer demands, and the software for managing the fabrics will begin to mature.

ESS Fibre Channel Attachment

More near term, native FICON will be delivered by IBM, providing Fibre Channel connectivity to the S/390 platform. FICON is a new high performance I/O interface based on Fibre Channel technology which provides major improvements beyond ESCON in bandwidth, efficiency, distance, addressability, flexibility, scale, and full duplexing capabilities. IBM already offers a Fibre Channel host connection to a FICON-to-ESCON bridge feature for the 9032 Model 5 ESCON Director that allows customers to attach their ESCON devices to a Fibre Channel (FICON) host. IBM plans to deliver native FICON attachment capability on its ESS disk subsystem, as well as on its tape subsystems, in addition to a FICON switch that will have capabilities similar to those found in today's ESCON directors. Also, IBM has completed the full disclosure of the new native FICON (FC-SB-2) FC-4 mapping protocol layer to NCITS for an ANSI (American National Standards Institute) standard. This work has been underway since 1998. Completion and anticipated adoption of this standard in the ANSI committee is targeted for December, 2000.

Hosts, Operating Systems, and Fibre Channel Host Bus Adapters Supported

At ESS's Fibre Channel Attachment General Availability, support will be announced for a specific "set" of hosts, operating systems, Fibre Channel adapters, switches, and hubs. This "set" will be constantly changing as IBM continues testing additional configurations. Please refer to the following IBM Web site for the latest information regarding what is supported:

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

Note: Support for Sun, HP, and Novell Netware was added in ESS code load EC F25863, available 12/15/00.

Some Differences Between Fibre Channel and SCSI Host Attachment in ESS

- **LUN Affinity**

  In SCSI attachment, ESS LUNs have an affinity to ESS SCSI ports, independent of which hosts may be attached to the ports. Therefore, if multiple hosts are attached to a single SCSI port (ESS supports up to four hosts per port), then each host will have exactly the same access to all the LUNs available on that port. When the intent is to configure some LUNs to some hosts and other LUNs to other hosts so that each host is able to access only the LUNs that have been configured to it, then the hosts must be attached to separate SCSI ports. Those LUNs configured to a particular SCSI port are seen by all the hosts attached to that port. The remaining LUNs are "masked" from that port. This is referred to as "LUN Masking".

  In Fibre Channel attachment, LUNs have an affinity to the host's Fibre Channel adapter (via the adapter's World Wide Unique Identifier, a.k.a. the World Wide Port Name), independent of which ESS Fibre Channel port the host is attached to. Therefore, in a switched fabric configuration where a single Fibre Channel host can have access to multiple Fibre Channel ports on the ESS, the set of LUNs which may be accessed by the Fibre Channel host are the same on each of the ESS ports. The following figure graphically depicts these concepts:

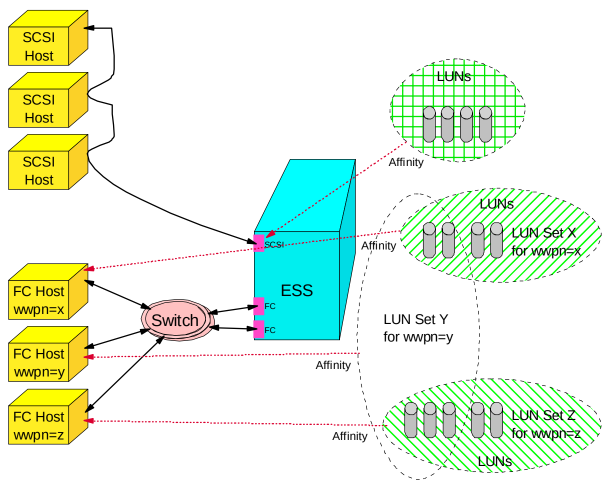

Figure 10. LUN Affinity

  A new function introduced in ESS code load EC F25863, available 12/15/00, called "Access Control by Host by Port", enables the user to restrict a host's access to one or more ports rather than always allowing access to all ports. This new capability significantly increases the configurability of the ESS.

  One result of this implementation is that with Fibre Channel, unlike in SCSI, hosts that are attached to ESS via a fabric to the same Fibre Channel port may not be able to "see" the same LUNs, since the LUN masking can be different for each Fibre Channel host. Another result of this implementation is that when the ESS has multiple Fibre Channel ports attached to a switch or a fabric, as in the figure above, each Fibre Channel host can discover each of these paths. In addition, some hosts are not sophisticated enough to realize that the LUNs accessed over the different paths are the same. For example, in the figure above, the Fibre Channel host with wwpn=x will see the four LUNs in LUN Set X via the first path to the ESS, and will see the same four LUNs via the second path. If the LUNs from each path are not detected as being the same

LUNs, then severe data integrity problems will arise when these LUNs are used by the operating system. Therefore this situation must be avoided.

There are three ways in which to get around this problem. The first is to install the IBM Subsystem Device Driver software (or equivalent software) in the host to handle the multiple paths to the ESS. This is IBM's preferred solution and is more fully described later in this paper in the section entitled "Multipathing for Availability and Performance". The second method is to create Zones in the switch fabric such that each Fibre Channel port from each host is constrained to attach to one and only one Fibre Channel port on the ESS, thereby allowing the host to see the LUNs via one path only. The third method is to utilize the new function "Access Control by Host by Port" to restrict the host's access to a single port. Implementing any of these solutions will avoid this possible configuration problem.

- **Targets and LUNs**

In SCSI attachment, each SCSI bus can attach a combined total of 16 initiators and targets. Since at least one of these attachments must be a host initiator, that leaves a maximum of 15 that can be targets. Shark is capable of exposing all 15 targets on each of its SCSI ports. Also in SCSI attachment, each target can support up to 64 LUNs. The software in many hosts is only capable of supporting 8 or 32 LUNs per target, but the architecture allows for 64. Since ESS supports 64 LUNs per target, it can support 15 x 64 = 960 LUNs per SCSI port.

In Fibre Channel attachment, each Fibre Channel host adapter can architecturally attach up to 2^{56} LUNs. However, the software in some hosts is only capable of supporting 256 LUNs per adapter. Also, in ESS, storage is partitioned in entities called Logical SubSystems, sometimes abbreviated as LSSs. Each LSS supports a maximum of 256 LUNs, and ESS supports up to 16 Open System LSSs. Therefore, the maximum LUNs supported by ESS, across all of its Fibre Channel and SCSI ports, is 16 x 256 = 4096 LUNs. If the software in the Fibre Channel host supports the SCSI command "Report LUNs", then it will support 4096 LUNs per adapter on ESS; otherwise it only supports 256 LUNs per adapter. The hosts that support "Report LUNs" are: RS/6000, AS/400, and NUMA-Q Dynix/ptx; hence they will support up to 4096 LUNs per adapter on ESS. All other host types, including NUMA-Q NT, will only support 256 LUNs per adapter on ESS.

- **Assignment of LUNs**

LUNs, or LUN IDs as they are formally referred to in the FCP documentation, are architecturally eight bytes in length. However, ESS only supports the upper two bytes of this LUN_ID field, which is defined as Level 1 Addressing. For the "Report LUNs" hosts, the ESS LUNs are assigned within the range '5000'x to '5FFF'x. For all other hosts the ESS LUNs are assigned within the range of '0000'x to '00FF'x. For those that are curious, the following explanation provides the reasons for these assignments.

In the two byte LUN_ID that ESS supports (the two high order bytes of the 8-byte LUN_ID field), the high order two bits are used to specify the addressing method - either "Peripheral Device" or "Virtual Device". (The actual meaning of these addressing methods is not important in this discussion.) The remaining 14 bits are used by ESS for the LUN identifier. This is shown in the following figure:

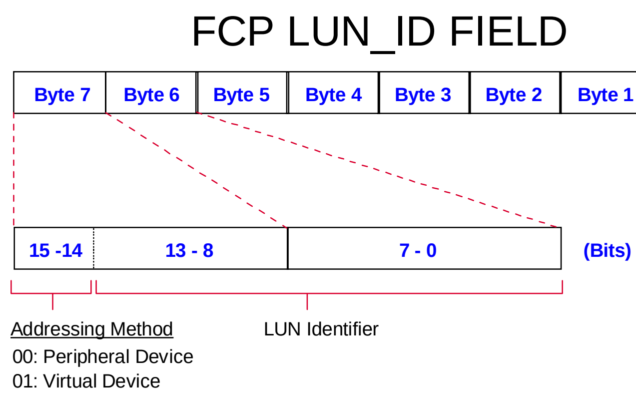

Figure 11. LUN Assignment

Each logical device within ESS has a logical device number and a LUN identifier. ESS automatically assigns both the logical device number and the LUN identifier to the devices. The customer has no control over these assignments. For the logical device number, ESS uses a 14-bit field. The high order bit is reserved for future use; the next five bits are used to identify the logical subsystem (as explained previously), and the remaining eight bits are used to identify the logical device within that logical subsystem. This is shown as follows:

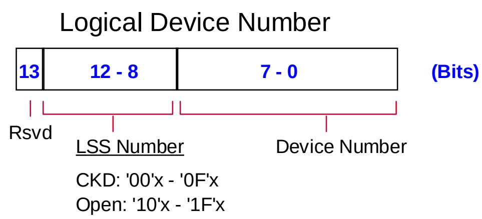

Figure 12. ESS Logical Device Number

ESS supports a maximum of 32 LSSs and 256 logical devices per LSS. 16 of these LSSs can be defined for attachment to FCP (Open) hosts, allowing a total of 16x256= 4096 (4K) FCP devices. The remaining LSSs and the devices within them are reserved for attachment to S/390 (CKD) hosts.

For hosts that support "Report LUNs", the ESS maximum of 4096 LUNs can be assigned. ESS will use its 14-bit logical device number, as described above, as the LUN identifier in the LUN_ID field for all LUNs assigned to these hosts. The LSS number ranges from '10'x to '1F'x (16 - 31 decimal), depending upon which LSS the logical device is contained within. The logical device within the LSS ranges from '00'x to 'FF'x (0 - 255 decimal). Also, ESS uses the "Virtual Device" addressing method for these devices. This addressing method is encoded as '01'b. The result of all of these bit settings is that the 2-byte LUN_ID field used in the FCP frame will range from '5000'x to '5FFF'x. This is shown in the following figure:

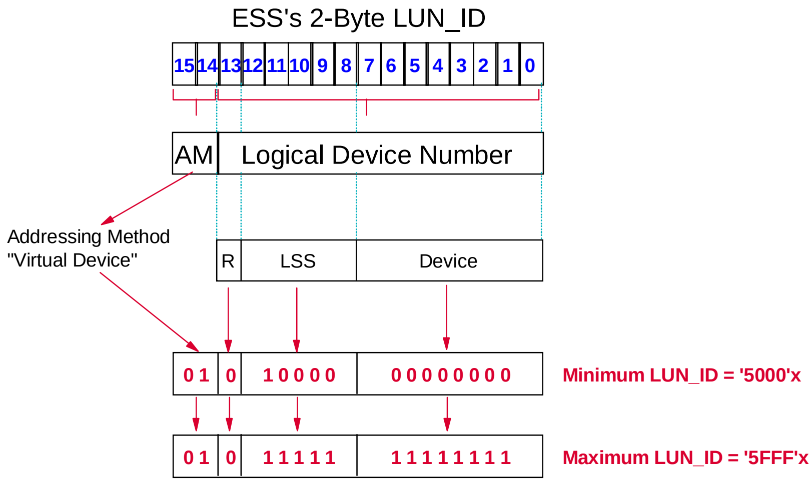

Figure 13. ESS LUN Range for "Report LUNs" Hosts

For hosts that do not support "Report LUNs", ESS must restrict its LUN identifier in the LUN_ID field to the low order byte, supporting LUNs 0 to 255 (decimal) only. Therefore, it cannot use the logical device number as it does for the "Report LUNs" hosts. It will map the logical devices to LUN_IDs in the range '00'x to 'FF'x. The high order byte of the LUN_ID field will be set to '00'x. The addressing method is encoded as '00'b, which indicates "Peripheral Device". This is shown in the following figure:

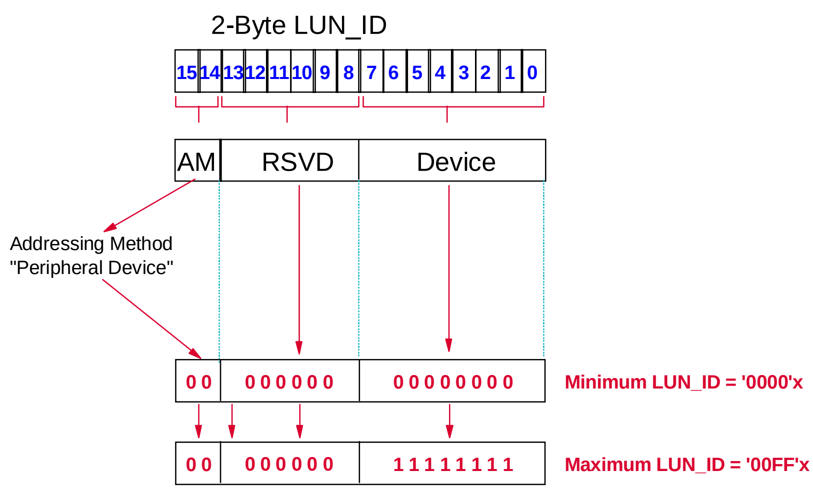

Figure 14. ESS LUN Range for "Non-Report LUNs" Hosts

These LUN assignments for the "Non-Report LUNs" hosts are made independently for each host's Fibre Channel adapter. Thus, a logical device that is configured to two different host Fibre Channel adapters may have one LUN_ID on the first adapter and a different LUN_ID on the second adapter.

- **LUN Access Modes**

In Fibre Channel attachment only, ESS provides an additional level of access security via either the "Access_Any" mode, or the "Access_Restricted" mode. This is set by the Customer Engineer and applies to all Fibre Channel host attachments on the ESS.

There is an Access Profile, which is automatically created when you assign volumes to a host, that is used to modify the effect of these access modes. If no volumes were assigned to a host, then that host will not have an Access Profile created for it.

In "Access_Any" mode, any host's Fibre Channel adapter for which there has been NO Access Profile defined can access all LUNs in the ESS (or the first 256 LUNs in the ESS if the host does not have the "Report LUNs" capability). In "Access_Restricted" mode, any host's Fibre Channel adapter for which there has been NO Access Profile defined, can access NONE of the LUNs in the ESS. In either access mode, a host's Fibre Channel adapter with an Access Profile can see exactly those LUNs defined in the profile, and no others.

The most common setting for most customers is "Access_Restricted". In this setting hosts have access to only those LUNs configured to them. And, with the new "Access Control by Host by Port" function in ESS code load EC F25863, the hosts have access only via the ESS ports that have been configured for them.

Considerations for New Fibre Channel Installs

In ESS the only requirement for Fibre Channel support is the appropriate level of microcode and the Fibre Channel adapters. The microcode upgrade from ESS 1.0 code to ESS 1.1 code will be non-concurrent. The installation of the Fibre Channel adapters is not, by itself, non-concurrent. Hosts sharing the adapter bay where the Fibre Channel adapters are being installed will temporarily lose access to their host attachments while the new adapters are being installed. However, alternate paths (if configured) will remain available. For an in-depth discussion of Fibre Channel installations on ESS, please see the IBM Redpiece "Implementing Fibre Channel Attachment on the ESS", document number SG24-6113. This can be found at the following web site:

http://www.redbooks.ibm.com

The following lists additional considerations that should be taken into account when installing Fibre Channel attachments to ESS.

1. **Fibre Channel Host**
   The first thing that needs to be considered for Fibre Channel attachment is whether or not the Fibre Channel host has the appropriate Fibre Channel Host Bus Adapter(s) (HBA(s)) installed, or at least has slots available for them. Next, the IBM web site at

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

   should be checked to validate that the specific operating system level, SCSI device driver level, and HBA hardware and firmware levels are appropriate for attachment to ESS.

2. **ESS Fibre Channel Adapters**
   You will need to order the new Fibre Channel adapter(s), feature code 3022 on the ESS 2105, and the appropriate microcode upgrade for your ESS. Each adapter contains a single Fibre Channel port. There are a total of 16 adapter slots in the 4 I/O bays of the ESS. You need to consider how many of these slots you want to use for Fibre Channel, how many for SCSI, and how many for ESCON. For maximum availability and serviceability, when installing multiple Fibre Channel adapters for attachment to one or more systems, you should spread these adapters evenly first across the two clusters and then across the two I/O bays in each cluster. In addition, for maximum performance you should install these adapters in the following order:
   - 1st adapter: Bay 1
   - 2nd adapter: Bay 3
   - 3rd adapter: Bay 2
   - 4th adapter: Bay 4

   then repeat this sequence for additional adapters.

   Note: If all 16 adapter slots are filled with SCSI and/or ESCON adapters, then you must first remove some of these to make room for the new Fibre Channel adapters.

   Once installed, these Fibre Channel adapters must be configured. As mentioned earlier, when attaching to a switch the adapter should be configured for the point-to-point protocol. When attaching directly to a host's HBA, the adapter should be configured for the loop protocol, since most HBAs will default to loop mode when no switch is configured.

3. **Fibre Channel Network (a.k.a. Storage Area Network or SAN)**
You must decide exactly how you want to implement your Fibre Channel network. Are you going to connect the ESS to your host(s) via a point-to-point connection or a fabric? As discussed earlier, the fabric topology provides the best connectivity with the best performance. If you implement a fabric, you should consider implementing IBM's Tivoli management software to manage it. Please check the previously mentioned IBM web site for compatible switches and hubs.

4. **Fiber Type**
If you do not already have fiber installed in your establishment, then you need to decide which fiber type to install. The ESS Fibre Channel adapters incorporate short wave laser transmitters. Therefore, you need to install multi-mode fiber for these attachments. The ESS Fibre Channel adapter comes with a 31 meter 50 micron fiber cable. However, this adapter also supports 62.5 micron fiber cable. Also, trunk cables in standard lengths up to 138 meters, and jumper cables in standard lengths up to 61 meters, can be ordered from IBM. Additional information can be found at this web site:
http://www.as.ibm.com/asus/connectivity.html

5. **Configuring for Long Distances**
ESS supports a maximum host attachment distance of 11 kilometers. As explained earlier, in a simple point-to-point topology, using ESS's Fibre Channel adapter to directly connect to a Fibre Channel host, the maximum distance achievable is 500 meters. This assumes that you are using the 50 micron fiber rather than the 62.5 micron fiber (which would only allow you to achieve a distance of 175 meters). If longer distances are required there are several simple solutions:

   - Fabric Attachment
   The maximum distance can be doubled from 500 meters to 1000 meters simply by incorporating a basic (single switch) fabric (IBM 2109 Fibre Channel Switch) in your configuration and locating it half way between the Fibre Channel host and the ESS. This is shown in the following figure:

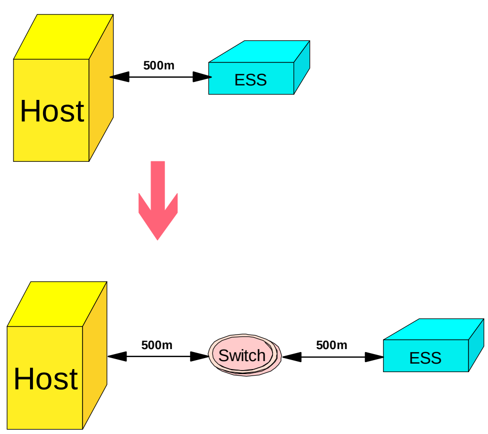

Figure 15. Extending Distance Via Switch

Fabrics with cascaded switches can extend this distance up to the 11 kilometer maximum, since the expansion links interconnecting the individual switches can be up to 10 kilometers long. This is shown in the following figure:

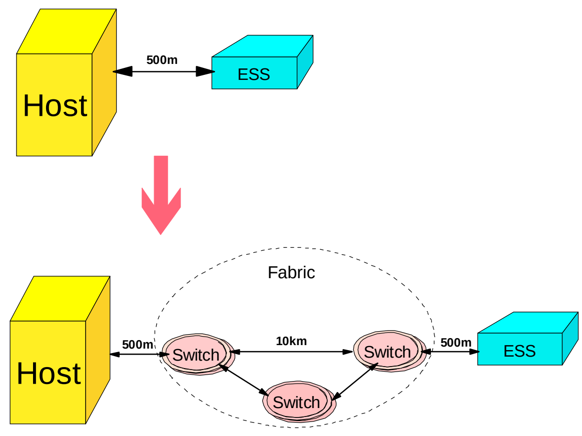

Figure 16. Extending Distance Via Fabric

- <u>Hub Attachment</u>
  Another method that can be used to extend the host to device distance is by using two IBM Fibre Channel Storage Hubs. These hubs can be populated with both short wave laser and long wave laser ports, thereby allowing you to connect the ESS to one of the short wave laser ports on the first hub, connecting the host to one of the short wave laser ports on the second hub, and then connecting the two hubs together using the long wave laser ports. This is shown in the following figure:

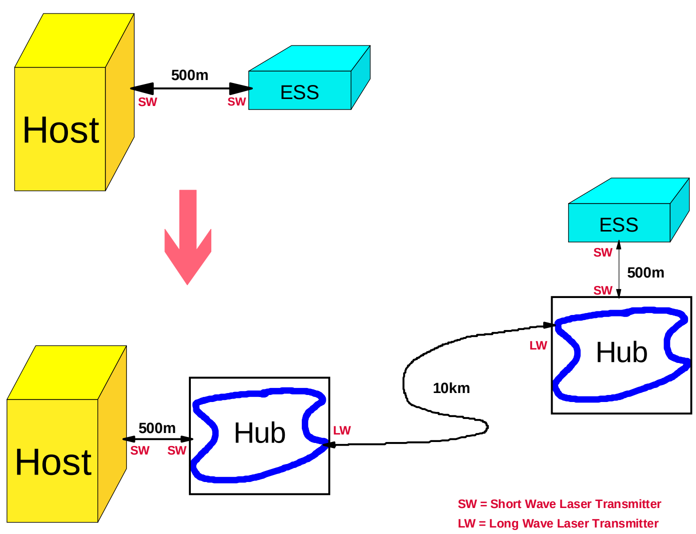

Figure 17. Extending Distance Using Two Hubs

- <u>Fabric or Hub Attachment With Long Wave Laser Transmitters</u>
  In the future, ESS and host Fibre Channel adapters will be offered with long wave laser transmitters for extended distance capability. When these are available it will be possible to get the extended distance capability with direct attachment or by using either a single switch or a single hub as shown in the following figure:

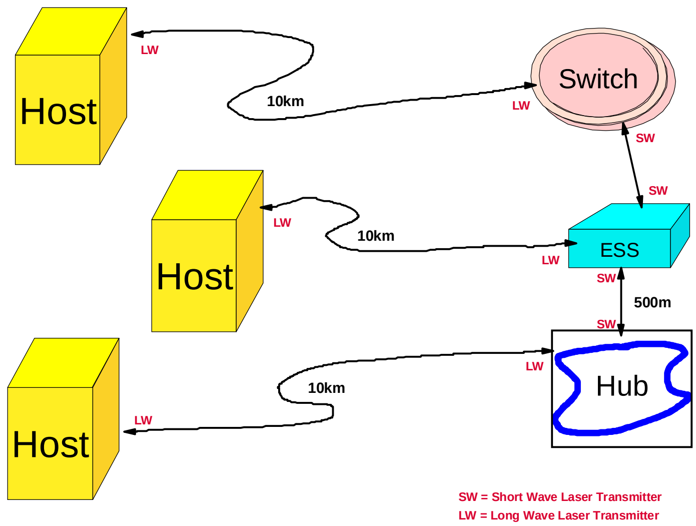

Figure 18. Extending Distance Using Long Wave Transmitters

6. **<u>Multipathing for Availability and Performance</u>**

If you have a requirement for high availability, then you must configure at least two separate Fibre Channel interfaces all the way from the Fibre Channel host to the ESS. At the host these interfaces should be on separate Fibre Channel adapters which are themselves on separate power boundaries, if possible. At the ESS, these interfaces should be on separate Fibre Channel adapters where one is plugged in Cluster 1 and the other is plugged in Cluster 2. Then you must configure each LUN that is to be accessed via multipathing to each path's host Fibre Channel adapter.

In order for the host to take advantage of both paths, you must also install the IBM Subsystem Device Driver (or other vendor software with this same functionality) to manage the paths. This device driver replaces the Dynamic Path Optimizer (DPO) for use with ESS. It runs in the host and will, in the event of failure in one path, automatically switch to the other path. This software is not limited to supporting just two paths - it will support up to 32 separate paths. Therefore you could, for example, configure 8 physical paths from your host to a fabric, then 4 physical paths from the fabric to the ESS. This configuration provides a total of 32 different paths by which the host can communicate with the ESS. The IBM Subsystem Device Driver can use all 32 of these paths.

In addition to providing automatic failover and recovery to an alternate path, the IBM Subsystem Device Driver also performs dynamic load balancing across all configured paths. It will spread the I/O across the paths in such a way that path utilization is maximized. This will obviously improve your total performance since it will either eliminate or drastically reduce the amount of time that any I/O must wait for a path to become available before it can execute.

With SDD 1.2.1, available on 12/15/00, these failover and load balancing functions are now supported for clustered environments in addition to the previously supported single-host environments. The specific clustered environments supported with SDD 1.2.1 are RS/6000 or RS/6000 SP with AIX 4.3.3 and HACMP 4.3.1 or 4.4 and Intel-based NT 4.0 Enterprise Edition

(MSCS). The upgrade from any previous levels of SDD or from DPO requires a reboot of the host system.

For more information on the IBM Subsystem Device Driver, please visit the following web site:

http://www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

7. **Data Sharing**
   Configurations where multiple hosts are actually accessing the same LUNs are called data sharing configurations. These configurations require some software in the hosts to manage concurrent access to the LUNs. This implies (at least with today's implementations) that the sharing hosts use the same operating system and file system, or the same DataBase Management System (DBMS) which supports sharing. In ESS you must configure each LUN that is to be accessed via multiple hosts to each host Fibre Channel adapter.

Installing Fibre Channel Host Attachment To ESS

Note: For the latest information regarding the specific instructions for installing Fibre Channel host attachment on ESS, please refer to the "ESS Host Systems Attachment Guide", IBM publication SC26-7296. For the latest information relating to modifying your host attachment on ESS using the ESS Specialist, please see the "ESS Web Interface User's Guide - Specialist", IBM publication SC26-7346.

In general, the steps required at the ESS to add a new Fibre Channel host attachment to the ESS (i.e., a host that was not previously connected to the ESS**)**, are as follows:

- Install the latest level of ESS microcode supporting Fibre Channel attachment. As mentioned earlier, this microcode upgrade from ESS 1.0 code to ESS 1.1 code is non-concurrent.
- Install the Fibre Channel adapter(s) in the ESS (feature code 3022 in the ESS 2105) and connect the fiber cable(s) from the host, fabric, hub, or fiber patch panel to the adapter(s). If multiple Fibre Channel adapters are to be installed, then in order to perform this step concurrently, only one I/O bay in ESS should be powered down at any one time, allowing the adapters in the other 3 bays to continue providing connectivity to the attached hosts.
- Using the ESS Specialist, navigate to the **Storage Allocation** panel; then to the **Open System Storage** panel; then to the **Configure Host Adapter Ports** panel. Use the **Configure Host Adapter Ports** panel to set up the Fibre Channel adapter.
  - First click on the icon that represents the specific Fibre Channel adapter to be set up
  - The middle of the panel will then display information relative to Fibre Channel hosts
  - The **Fibre Channel Access Mode** must be set appropriately (by the CE) to either "Access_Any" or "Access_Restricted"
  - Use the drop-down menu to select the **Fibre Channel Topology** that will be used
    - Select Point-to-Point if this adapter will be connecting directly to a host's Fibre Channel adapter that is also configured for point-to-point, or if this adapter will be connecting directly to a fabric
    - Select Arbitrated Loop if this adapter will be connecting directly to a host's Fibre Channel adapter that is also configured for arbitrated loop, or to a hub used for distance extension.
- Now, navigate to the **Modify Host Systems** panel. Use the **Modify Host Systems** panel to add the new Fibre Channel host to the list of host systems available for attachment to the storage server. (Note: these steps are required for EACH Host Fibre Channel adapter being attached to ESS.)
  - Enter the host name in the **Nickname** field
  - Enter the host type in the **Host Type** field
  - Enter the host attachment - Fibre Channel Attached - in the **Host Attachment** field
  - Select the World-Wide Port Name from the drop-down list in the **World-Wide Port Name** field
    Note: Refer to the Appendix in "ESS Host Systems Attachment Guide", SC26-7296, for instructions on how to determine the WWPN for a host's Fibre Channel adapter.
  - Enter a description of the host in the **Hostname / IP Address** field
    This is optional but recommended, since when using the ESS Expert software it indicates which host Fibre Channel adapters are contained within the same host machine.
  - Enter the list of ESS Fibre Channel ports by which you want to allow this host to access the ESS in the **Fibre-Channel Ports** field

  - Click **Add** to add the host to the **Host Systems** list on the right side of the panel
  - Click **Perform Configuration Update**
- Continuing to use the ESS Specialist, navigate back to the **Open System Storage** panel and then assign the volumes to this host just as you would for a non-Fibre Channel host.
- If the IBM Subsystem Device Driver is being used, then it should now be installed in the host.
- Reboot the host to complete the install of the IBM Subsystem Device Driver and to invoke its I/O discovery routines, allowing the new devices on the Fibre Channel interfaces to be available for use by the file system.

Migrating From SCSI Hosts To Fibre Channel Attachment

Note: For the latest information regarding modifying your host attachment on ESS using the ESS Specialist, please see the "ESS Web Interface User's Guide - Specialist", IBM publication SC26-7346.

Please refer to the "ESS Fibre Channel Migration Scenarios" document, located at the following web site:

http://www.storage.ibm.com/hardsoft/products/ess/support/essfcmig.pdf

This document describes the steps required for migrating from native SCSI to native Fibre Channel attachment.

Migrating From SAN Data Gateway To Fibre Channel Attachment

Note: For the latest information regarding modifying your host attachment on ESS using the ESS Specialist, please see the "ESS Web Interface User's Guide - Specialist", IBM publication SC26-7346.

The SAN Data Gateway (SDG), IBM product 2108 Model G07, is a protocol converter between SCSI and Fibre Channel interfaces. It is supported by ESS for attachment to Fibre Channel hosts via ESS's SCSI interfaces. SDG was a necessary component for our early ESS customers utilizing Fibre Channel hosts, allowing these hosts to connect to the ESS via SCSI. With ESS's native Fibre Channel attachment, SDG will no longer be required to attach ESS to Fibre Channel hosts. Please note, though, that IBM continues to fully support the SDG in the other environments in which it has been used. For the latest information on the IBM 2108 Model G07, please refer to the SAN Data Gateway web site at:

http://www.ibm.com/storage/SANGateway

The following figure shows a graphical representation of this migration.

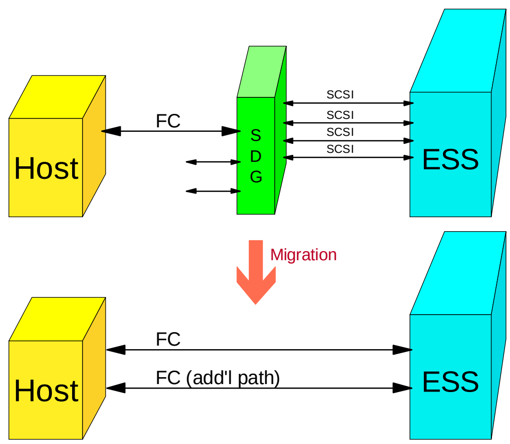

Figure 19. Migration From SDG to Native Fibre Channel

Please refer to the "ESS Fibre Channel Migration Scenarios" document, located at the following web site:

http://www.storage.ibm.com/hardsoft/products/ess/support/essfcmig.pdf

This document describes the steps required for migrating from the SAN Data Gateway to native Fibre Channel attachment.

Performance Considerations

The most obvious and highly touted performance characteristic of Fibre Channel is its bandwidth capability of 100 MB/sec full duplex, for a total of 200 MB/sec. This is far superior to the standard SCSI interfaces being used today, where 20, 40, 80, and 160 MB/sec implementations are available, though only 20 and 40 MB/sec are in common use. However, there are a few other points that need consideration as you think about the performance you can get from your new Fibre Channel attachment.

1. **Sustainable Bandwidth**
   The bandwidth numbers cited above are maximum bandwidth capabilities of the architecture. However, one should never plan workload based on these maximums. In addition, the throughputs will vary based upon type of application, specific adapter, and host attachment. As a general rule-of-thumb, one should base planning decisions assuming a throughput capability of around 60 MB/sec per fiber port configured to an ESS.

2. **Total Number of Ports**
   When the performance requirements of the application(s) is well understood, then it is possible to use the guidelines just presented in the previous section to determine the proper number of host ports. For example, if you know that your application requires 100 MB/sec of aggregate throughput, then you would most likely need two Fibre Channel ports to support this requirement (2 x 60MB/s >or= 100MB/s).
   Sometimes you have little or no information about the exact performance requirements of the application. In these cases you must make educated guesses about the number of ports needed. The following rules-of-thumb can help in these cases:
   - For an ESS subsystem configured entirely for open systems attachment, four fiber ports will usually be adequate to handle most workloads and applications. Most online transaction processing, file serving, and standard applications fall into this category. This assumes that the workload is well balanced across the ports. Naturally, if some of the ports have very little activity while others are very busy, then you need to plan for more ports. (Note: in this case, installing and using the IBM Subsystem Device Driver (or equivalent software) for balancing the workload across the multiple ports can eliminate the need for additional ports.)
   - For some high bandwidth applications in open systems environments, more ports could be necessary. For example, business intelligence applications that rely on massive parallel database scans which are likely to need every last bit of bandwidth might need as many as eight Fibre Channel ports. In this case you should install two adapters in each of the four I/O bays.
   - In some cases, you may not be planning to use the full subsystem capacity for accessibility from Fibre Channel hosts. For example, you might be mixing SCSI and ESCON attachments with your Fibre Channel attachments. In this case it might be useful to plan the number of Fibre Channel ports based on the amount of disk capacity you plan to access through the Fibre Channel ports. A rule-of-thumb for planning is to use one Fibre Channel port for every 400 GB of data, with a two-port minimum to ensure data availability in the event of an interface failure.

3. **Port Positions in ESS**
   As stated above, the Fibre Channel adapters are plugged into the 4 I/O bays in ESS. Each I/O bay is composed of a single bus that interconnects the adapters on that bay. When populating the Fibre

Channel adapters, you want to spread them as evenly as possible across all of the I/O bays in order to get the best performance. These adapters should be installed in the following order:
- 1st adapter: Bay 1
- 2nd adapter: Bay 3
- 3rd adapter: Bay 2
- 4th adapter: Bay 4

then repeat this sequence for additional adapters.

4. **Loop Versus Fabric**
For best system performance you want to interconnect your Fibre Channel components via point-to-point or fabric topologies rather than the loop topology. The reason for this is because each Fibre Channel component connected in a point-to-point or fabric topology can perform at the full Fibre Channel capability, while each Fibre Channel component connected on a loop must share the loop's Fibre Channel bandwidth among all the other components also connected on this loop.

5. **Performance Comparison: Native SCSI vs SDG vs Native Fibre Channel**
The following figure shows how the performance increases as you migrate from native SCSI attachment to the SAN Data Gateway and finally to native Fibre Channel attachment:

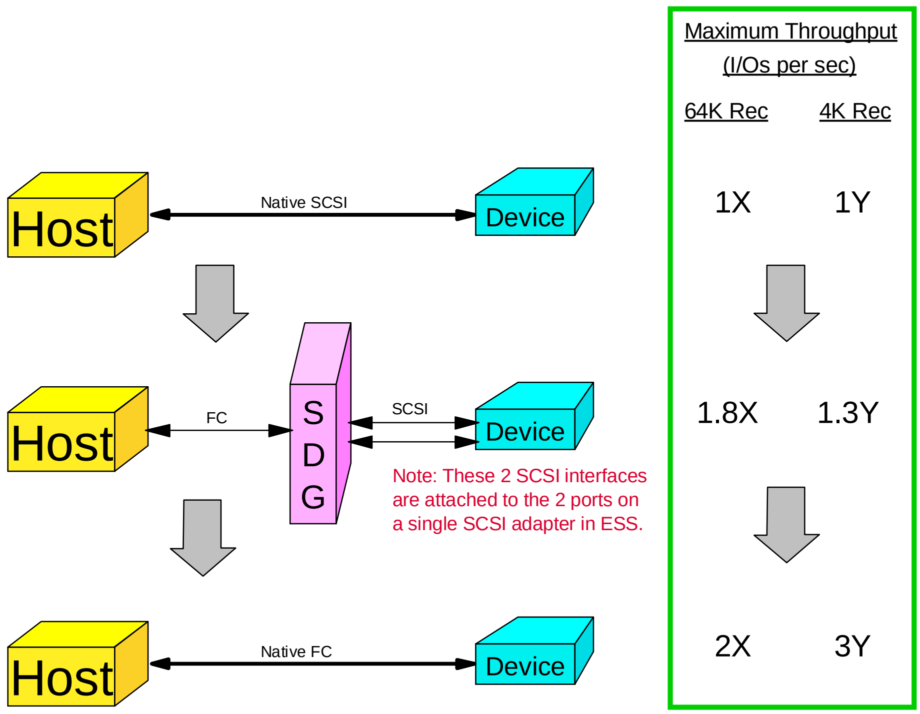

Figure 20. Performance Comparison: SCSI - SDG - FC

Summary

Fibre Channel is the industry's strategic new interface that will allow the customer the most flexibility in achieving his future goals in storage growth. It is so popular, even today, that every major storage and storage peripheral vendor is aggressively designing and shipping products that are Fibre Channel capable. This includes storage subsystems, hosts, switches, hubs, repeaters, and especially the software to manage the varied Fibre Channel configurations.

The Fibre Channel technology promises to avoid any early death due to obsolescence by being built upon a scalable architecture that inherently supports signaling rate improvements of at least one order of magnitude (to 10 Gbits/sec) and probably greater. This scalability was a major objective of the Fibre Channel architecture from its very beginning.

It is, indeed, the very rich connectivity options made possible by Fibre Channel technology that has fueled the concept of Storage Area Networks (SANs) in the last year or two. Fibre Channel technology is the infrastructure of the SAN.

To learn more about the Fibre Channel architecture and its specific usage in ESS, please see the references that are included in the next section. Many of these references can be found on the internet at the following IBM web site:

http://www.ibm.com/storage/hardsoft/products/ess/refinfo.htm

References and Web Sites

1. Fibre Channel Physical and Signaling Interface (FC-PH), ANSI X3.230-1994
2. Fibre Channel Second Generation Physical Interface (FC-PH-2), ANSI X3.297-1997
3. Fibre Channel Third Generation Physical Interface (FC-PH-3), ANSI X3.303-199X, Revision 9.4
4. Fibre Channel Arbitrated Loop (FC-AL), ANSI X3.272-1996
5. Fibre Channel Arbitrated Loop - 2 (FC-AL-2), NCITS Project 1133-D, Revision 7.0
6. Fibre Channel Arbitrated Loop - 3 (FC-AL-3), NCITS Project 1304-D, Revision 1.0
7. Fibre Channel Protocol for SCSI (FCP-2), Second Revision
8. Fibre Channel Private Loop, SCSI Direct Attach (FC-PLDA)
9. The Fibre Channel Consultant - A Comprehensive Introduction (Robert W. Kembel, 1998)
10. The Fibre Channel Consultant - Arbitrated Loop (Robert W. Kembel, 1996)
11. ESS User's Guide, SC26-7295
12. ESS Web Interface User's Guide - Specialist, SC26-7346
13. ESS Host Systems Attachment Guide, SC26-7296
14. ESS Introduction and Planning Guide, SC26-7294
15. ESS SCSI Command Reference, SC26-7297
16. ESS Configuration Planner, SC26-7353
17. ESS Systems Assurance Product Review (SAPR) Guide, SA99-004
18. SAN Data Gateway Installation Guide
19. Sample SAN Configurations for Disk (Mark Bruni)
20. IBM Hardware Announcement 100-089, dated March 28, 2000, and entitled: "IBM Enterprise Storage Server Multiplatform Storage Solution Provides Superior Performance, Function, Scalability, and Availability"
21. ESS Fibre Channel Migration Scenarios (Phil Mills), located at: http://www.storage.ibm.com/hardsoft/products/ess/support/essfcmig.pdf
22. IBM Redbook: "Designing an IBM Storage Area Network", SG24-5758, located at: http://www.redbooks.ibm.com
23. IBM Redpiece: "Implementing Fibre Channel Attachment on the ESS", SG24-6113, located at: http://www.redbooks.ibm.com
24. IBM Web Site for Supported Fibre Channel Server Configurations for the 2105 ESS: http://www.storage.ibm.com/hardsoft/products/ess/supserver
25. IBM Web Site for the 2108-G07 SAN Data Gateway: http://www.ibm.com/storage/SANGateway
26. IBM Web Site for the 2109 Fibre Channel Switch: http://www.storage.ibm.com/hardsoft/products/fcswitch/fcswitch.htm
27. IBM Web Site for the 2103-H07 Fibre Channel Hub: http://www.storage.ibm.com/hardsoft/products/fchub/fchub.htm
28. IBM Web Site for the IBM Subsystem Device Driver: http://www.ibm.com/storage/techsup/swtechsup.nsf/support/sddupdates

29. IBM Web Site on Storage Area Networks:
http://www.storage.ibm.com/ibmsan/basics.htm
30. IBM Web Site for ESS Documentation:
http://www.ibm.com/storage/hardsoft/products/ess/refinfo.htm

Notices

Published: December 15, 2000

IBM does not provide any warranty, guarantee and assurance that the information in this paper is free from errors or omissions. This paper contains discussion of future IBM development plans, which are subject to change or withdrawal at any time without prior notice.

IBM, IBM Enterprise Storage Server, IBM Enterprise Storage Server Specialist, StorWatch, Seascape, RS/6000, AS/400, S/390, FICON, ESCON, FlashCopy, are trademarks of International Business Machines Corporation.

Windows NT is a trademark of Microsoft Corporation in the United States, or other countries, or both. UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

Other company, product, and service names may be trademarks or service marks of others.

Please send comments via e-mail to Phil Mills: millsp@us.ibm.com.