TECHNOLOGY BRIEF

October 14, 1997

Compaq Computer
Corporation

Strategic Direction for Compaq Fibre Channel-Attached Storage

CONTENTS

EXECUTIVE SUMMARY

As performance of processors and peripherals improves and companies increasingly move to distributed architectures while consolidating servers, high-speed and data intensive network applications (such as transaction processing, decision support, data warehousing, image-based document systems, geophysical mapping, and multimedia) are proliferating. Interconnects between servers and the I/O devices they support have become a management bottleneck. Current interconnects require that I/O devices be located within very close proximity to servers. This limited transmission distance is inadequate for mirrored data sites. A further restriction is the number of I/O devices that can be attached to systems.

New interconnect technology is needed to overcome current I/O and physical limitations and to meet future demands. Nowhere is this need more critical than in the storage subsystem.

This technology brief addresses recent computing trends and customer issues with current storage technology. It provides an overview of Fibre Channel technology and explains why Compaq's strategic direction for high-performance and high-capacity external storage is based on Fibre Channel.

Please direct comments regarding this communication to the ECG Technology Communications Group at this Internet address: TechCom@compaq.com

COMPAQ

NOTICE

Strategic Direction for Compaq Fibre Channel-Attached Storage
Second Edition (October 14, 1997)
Document Number ECG009.1097

STORAGE TRENDS AND CUSTOMER ISSUES

Today storage needs are expanding more rapidly than ever before. Customers are demanding improved storage solutions. Driving their demands are these new trends in enterprise computing:

- the explosion of the Internet
- the need to keep more information on-line
- the need to collect, scan, and track decision support information
- consolidation of servers
- movement of PC Servers into business critical applications
- the growing complexity of applications with more graphics, video, and sound to be stored

Key storage issues for enterprise customers include their current and future needs for distributed storage in conjunction with improved network storage management; increased connectivity and capacity, plus dynamic expansion capabilities; high performance, availability, and reliability; investment protection; and reduced cost of ownership.

Small Computer System Interface (SCSI) technology has carried the storage industry forward for many years. Inherent I/O and physical limitations, however, now prevent SCSI technology from satisfying the expanding needs of enterprise storage.

Although SCSI will remain an important part of data storage solutions for some time, a new interconnect technology must propel future growth of enterprise storage. This technology brief explains why Compaq believes Fibre Channel is the right interconnect technology for building future storage solutions.

TERMINOLOGY AND CONVENTIONS

In this document the term *fibre* (international spelling) refers to a communication medium consisting of either copper or fiber optics. The term *Fibre Channel* is capitalized in accordance with the convention set by the governing standards committee.

COMPAQ'S STRATEGIC DIRECTION FOR EXTERNAL STORAGE

Compaq's strategic direction for high-performance and high-capacity external storage is based on Fibre Channel technology because it provides the means to satisfy all the enterprise storage needs identified above. Fibre Channel is a key technology for the high-speed storage interconnect (that is, processor-to-storage and storage-to-storage communications) and for the serial drive interface (high-performance disk systems). It provides opportunity for the integration of primary and secondary storage as well as for shared storage among multiple servers.

Compaq also has a strong interest in Gigabit Ethernet and Tandem ServerNet. Gigabit Ethernet is a high-speed extension to Ethernet. It leverages the physical level and the encoding used in Fibre Channel. Gigabit Ethernet and Fibre Channel are complementary. Gigabit Ethernet provides the high-speed local area network, while Fibre Channel provides the high-speed storage area network. A Fibre Channel storage area network allows a client attached to a specific processor to access data in any storage device within the storage area network because all storage devices are accessible to all processors.

Compaq is developing ServerNet, on the other hand, as the server node-to-server node interconnect within Compaq clusters because ServerNet features very low latency in server node-to-server node communications. Current implementations of ServerNet use a copper interface. Future implementations will use fiber and leverage the physical level and encoding of Fibre Channel, as does Gigabit Ethernet.

Commonality of architecture at the physical level of these three technologies promotes the use of common parts and allows the use of the same infrastructure for Fibre Channel, Gigabit Ethernet, and ServerNet.

WHAT IS FIBRE CHANNEL?

Fibre Channel is the general name of an integrated set of standards being developed by committees accredited by the American National Standards Institute (ANSI). This set of standards defines new protocols for flexible information transfer. Fibre Channel is an industry standard interconnect and high-performance serial I/O protocol that is media independent and supports simultaneous transfer of many different protocols.

Development of the Fibre Channel standards began in 1988. These standards are being developed to meet several objectives:

- to keep pace with increasing host processor performance
- to keep pace with growing data-intensive applications
- to provide a practical and inexpensive means for high-speed transfer of large amounts of data
- to ensure the integrity of data
- to support multiple physical interface alternatives
- to provide a common interface for all data traffic
- to provide a means of transmitting data with very low error rates
- to separate logical protocol from physical interface, allowing transport of multiple protocols over the same interface
- to allow simultaneous transfer of many different protocols over the same interface

Channel and Network Functions

In business computing there are two basic protocols for device communication: channels and networks. A channel is an interface between a host computer and I/O peripherals such as tape drives, disks, and printers. The host system has knowledge of all the peripheral devices attached to it, so this is a structured, predictable, hardware-intensive environment with relatively low software overhead.

A network, on the other hand, comprises distributed devices that may include mainframes, workstations, and file servers. A network has its own protocol and is an unstructured environment because almost any device in the network can communicate with any other device at any time (peer-to-peer communication). This environment has more overhead than a channel because more software support is required to verify access permission, set up sessions, and route transactions correctly.

Fibre Channel is superior as a traditional channel for attaching storage devices in a robust fashion. The use of fiber optics for the transmission media provides for extremely low error rates. Fibre Channel incorporates both a powerful encoding scheme and a strong cyclic redundancy check (CRC) on each message frame, ensuring data integrity. Fibre Channel uses a topology (either a loop or a switch) to provide connectivity. The capability to provide scalable connectivity and the peer-to-peer basis of the Fibre Channel architecture are the key enablers for networking. Fibre

Channel is now the only standard that can perform both the traditional channel and network functions simultaneously on the same port.

Interconnect Topologies

Fibre Channel nodes each have one or more ports that enable external communication. Each port uses two fibres, one for outgoing information and the other for incoming information. The pair of fibres is called a *link*. All the components that connect ports comprise an interconnect topology.

Various topologies are used to provide connectivity between Fibre Channel ports. The two basic topologies used today are Fibre Channel Arbitrated Loop (FC-AL) and the fabric switch. Both are illustrated in Figure 1.

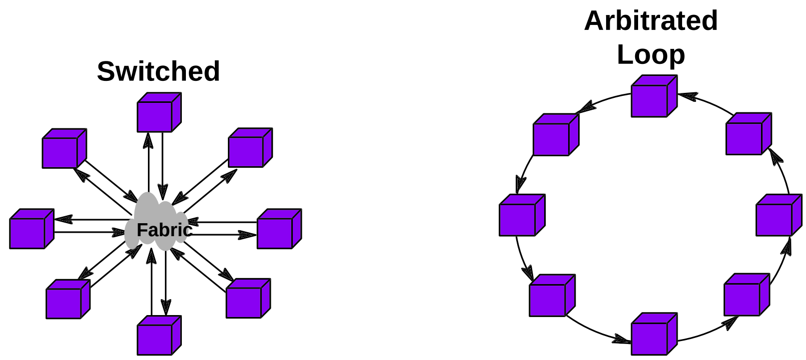

Figure 1. Simplified depiction of non-blocking cross-point switch topology and of FC-AL topology

A fabric switch allows multiple pairs of nodes to communicate with each other simultaneously. Therefore, as more nodes are added, the aggregate data throughput capability can increase incrementally. A fabric switch requires a cross-point switching function and the intelligence to make the connections. A pair of transceivers is required to form the link between the attaching port and the port on the switch. These transceivers add to the cost of the switch.

The Fibre Channel Arbitrated Loop is a serial interface that creates logical point-to-point connections between ports with the minimum number of transceivers and without a centralized switching function. FC-AL therefore provides a lower cost solution. The bandwidth of a Fibre Channel loop is shared by all ports on the loop. A single pair of ports on the loop communicates at one time, while the other ports on the loop act as repeaters.

Hubs are useful in configuring the Arbitrated Loop. A hub contains several ports that are internally connected in a loop. Each port is fitted with a port bypass switch to maintain the continuity of the loop should a controller or device attached to the port be powered off or malfunction. A hub port can be given the ability to accept either electrical or optical input. This capability is useful in configuration. For instance, if it were desirable to locate the hub and controllers some distance from the server, an optical connection (long wave or short wave) could be used between the server and hub while copper connections could be used between the hub and controllers. Hubs can be cascaded to provide additional ports for more connectivity.

Structure of Transmitted Information

In network terminology, a block of information to be transmitted or processed is called a *packet*. A packet is sent as a series of one or more frames. A *frame* is a block of bytes plus control information.

With Fibre Channel, the structure of transmitted information is a bit more complex. Figure 2 illustrates the structure of information transmitted via Fibre Channel interconnects. The lowest level of the structure is the frame, which is encapsulated with special ordered sets that define the start of frame (SOF) and end of frame (EOF). Contents of the frame include the frame header (which contains additional controls, the source address, and the destination address); up to 2112 bytes of payload; and a cyclic redundancy check field (CRC). The payload number, 2112, stems from the requirement to handle 2048 bytes of actual payload plus an optional 64-byte header. Payload can consist of data or control information. The source and destination address fields are each 24 bits long. These long fields provide a very high degree of future connectivity, more than 16 million addresses.

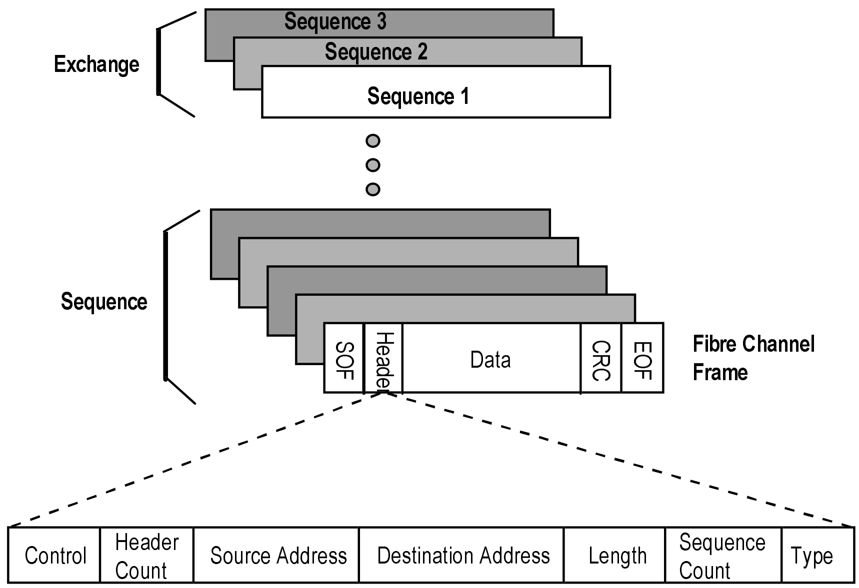

Figure 2. Structure of information transmitted via Fibre Channel interconnects

The second level in the hierarchy is the *sequence*. Each sequence is composed of one or more frames that contain payload and are specific to a particular protocol. The Small Computer System Interface (SCSI) protocol is an example.

The third level in the hierarchy is the *exchange*. An exchange can be equated to a complete I/O transfer, such as a SCSI read operation. An exchange is composed of one or more sequences. For example, the SCSI protocol uses individual sequences for the command phase, data phase, and status phase.

Functional Levels

Five functional levels are included in the Fibre Channel standard: FC-0 through FC-4. These levels are descriptive only. To allow efficiency and tradeoffs in implementation and high performance, a physical interface between these levels is intentionally <u>not</u> part of this architecture.

FC-0 defines the physical characteristics of the interface and media. To allow for maximum flexibility, use of existing media and different technologies, and meeting a wide variety of system requirements, the standard includes many variants. For example, it includes copper and fiber optic media with speeds of 12.5 megabytes per second (MB/s) and continuously doubling to 106.25 MB/s, or 1.0625 gigabits per second (Gb/s), commonly referred to as full speed. Tables 1 and 2 list some of the Fibre Channel media options available today. Speeds of 2 and 4 Gb/s have also been addressed by the standards working group responsible for Fibre Channel.

TABLE 1: OPTICAL MEDIA INCLUDED IN THE FIBRE CHANNEL STANDARD

Fibre Type	Laser Type	Effective Rate MB/s	Distance km	Baud Rate Mbaud
9-micron single mode	Long wave	100	up to 10	1062.5
	Long wave	25	up to 10	265.6
50-micron multimode	Short wave	100	up to 0.5	1062.5
	Short wave	25	up to 2	265.6
62.5-micron multimode	Short wave	100	up to 0.175	1062.5
	Short wave	25	up to 0.7	265.6
	Long wave LED*	25	up to 1.5	265.6

* Light Emitting Diode

TABLE 2: ELECTRICAL MEDIA INCLUDED IN THE FIBRE CHANNEL STANDARD

Conductor Type	Electrical Type	Effective Rate MB/s	Distance Meters	Baud Rate Mbaud
Video coax	ECL*	100	up to 25	1062.5
	ECL	25	up to 75	265.6
Miniature coax	ECL	100	up 10	1062.5
	ECL	25	up to 25	265.6
Shielded twisted pair	ECL	25	up to 50	265.6

* Emitter Coupled Logic

FC-1 defines the encoding/decoding and transmission protocol. The dc-balanced 8B/10B code has excellent properties for transmission and allows for low-cost component design, clock recovery, and error detection. An 8-bit byte is encoded into 10 bits for transmission and is then decoded at the receiving end. Some of the unused code points, which have special properties, are used to form special characters. These characters are used to form ordered sets for signaling and frame delineation.

FC-2 is the framing and signaling protocol level. It defines how data is transported from one port to the next. The structure of information transmitted via Fibre Channel interconnects was described in the previous section and Figure 2.

FC-3 provides common services such as striping of data or multicast operations. This level is not used in current implementations.

FC-4 specifies the mapping of upper level protocols to the lower levels of Fibre Channel. Examples of current mappings include these protocols:

- Small Computer System Interface (SCSI)
- Intelligent Peripheral Interface 3 (IPI-3)
- High Performance Parallel Interface (HIPPI)
- Internet Protocol (IP)
- IEEE 802.2
- Single Byte Command Code Set Mapping (SBCCS)

Classes of Service

The Fibre Channel standard currently describes three classes of service. Others are being developed to address needs for guaranteed bandwidth.

Class 1 is a switched circuit connection, like that used in telephone systems. Once a connection is established, the connecting parties/ports can use the entire bandwidth in both directions. The connection is established using a special start of frame (SOF), the source address, and the destination address in the frame header. The connection remains open for frame transmission until it is terminated by recognition of the End of Frame Disconnect (EOF) on the last frame of transmission. This type of service accommodates the traditional channel environment, ensuring rapid, uninterrupted communication with an extremely high degree of data integrity checking and positive/negative acknowledgment of frame reception.

Class 2 provides for frame switching. Frames to different recipients can be transmitted—multiplexed in time—to different recipients. Likewise, a recipient may receive frames from many different sources, multiplexed in time. Class 2 provides the same high degree of checking and positive/negative acknowledgment of frame reception that Class 1 provides.

Class 3 is much like Class 2 but with one major difference: there is no acknowledgment of frame reception. Today's networks operate in this mode at the hardware level. Class 3 is also used for the attachment of disk storage on FC-AL. The topology of the loop ensures in-order delivery; and the SCSI protocol that is used implicitly provides positive/negative acknowledgment, ensuring data integrity.

WHY CHOOSE FIBRE CHANNEL?

Enterprise computing customers are eager for interconnect technology that can overcome inherent I/O and physical limitations of current parallel SCSI. Serial interfaces have several advantages over their parallel counterparts that are discussed in the following section. Compaq believes that Fibre Channel is overall the best serial interconnect technology for high-performance, high-availability external storage. Because Fibre Channel is an industry standard, not a proprietary initiative, it enables customers to move more easily to the most cost-effective solutions while preserving their current investments.

Connectivity and Capacity

Fibre Channel increases connectivity. As ever-faster processors and multiprocessing systems are being built with smaller footprints, it becomes more and more difficult to attach enough parallel interfaces to a system to provide the necessary I/O bandwidth and connectivity.

Fibre Channel solves this problem. More than 100 devices or nodes can be attached on a single FC-AL loop. Large subsystems of devices can be attached to nodes on expansion modules, and expansion modules can be cascaded to connect far more subsystems. The scale of connectivity provides the basis for centrally managing primary and secondary storage. In addition, integrating Fibre Channel on the processor board could significantly reduce the number of parallel buses and adapters needed in the future.

Fibre Channel provides greater connectivity per PCI slot, which results in greater capacity. By the second half of 1997, the capacity of typical storage systems using 9-gigabyte, 1.6-inch drives will increase to 432 gigabytes per PCI slot. If six PCI slots were used for storage, then the raw capacity would be 2.6 terabytes.

In late 1998, the capacity per PCI slot will increase to a value still to be identified. A conservative estimate is that capacity per PCI slot will double. Through expansion of hub capability and through cascading of hubs, the number of additional spindles attachable per slot will also grow significantly.

Fibre Channel is the only interconnect that can be used both for traditional channel functions, such as attachment of primary and secondary storage subsystems, and for networking. Both functions—in fact, a multiplicity of protocols—can be incorporated in the same host adapter, saving precious slots. Fibre Channel technology also makes it possible to architect third party transfers in which data from a disk subsystem can be transferred directly to the requester, without having to pass through the server. (In this situation, the server would still maintain control of the operation.) Third party transfers can significantly increase data availability and system performance.

Availability and Reliability

In a typical bus environment with traditional SCSI interfaces, demands on hardware drivers vary with the number of devices on the bus and the cable lengths. Fibre Channel, on the other hand, provides logical point-to-point connections between nodes in which only two wires or fibers connect a device. Use of redundancy techniques in the implementation will enable transfer of data over two separate data paths for increased reliability.

As the Fibre Channel Loop Community notes in its publication *Ultra SCSI to Fibre: The Preferred Performance Path*: "FC-AL is the only interface designed to support simultaneous hot plugability. The FC-AL interface's loop redundancy circuit allows the removal or insertion of multiple drives from an active loop without impacting data throughput. This also provides for better data availability."

Transmission Distance

Increased transmission distance is another significant advantage of Fibre Channel. With parallel SCSI, the maximum distance between a host server and an external storage device is 12 feet, or approximately 3.5 meters. This precludes mirroring of data at an off-site location for disaster recovery. As previously indicated in Table 1, with short-wave lasers and 50-micron multimode fibre, a Fibre Channel interconnect can easily achieve transmission distances of approximately 500 meters. With long wave lasers and 9-micron single mode fibre, Fibre Channel can transfer data for distances of up to approximately 10 kilometers, allowing for remote mirroring of data at very high data rates. The longer communication links of Fibre Channel also allow multiple systems to share the same storage.

Performance

Bandwidth is another important advantage of Fibre Channel. Peak bandwidth of Ultra Wide SCSI is 40 MB/s. Ultra 2 Wide SCSI will offer a maximum bandwidth of 80 MB/s. Fibre Channel bandwidth, currently at 100 MB/s per link in both directions, is targeted to grow to 200 MB/s and 400 MB/s per link in the future.

Cost Effectiveness

Fibre Channel reduces the number of parallel buses and expensive controllers required to support numerous drives and reduces the complexity of cabling for multiple controllers. Moreover, because Fibre Channel cabling is smaller and more flexible than SCSI cable, it also improves cable deployment and management.

As costs come down and as requirements for better scalability increase, there will be a shift from SCSI to FC-AL on the device interface for higher end systems. However, the use of SCSI will persist for years for smaller systems where scaling, distance, and connectivity are not issues. An FC-AL connection to a storage subsystem, whose drives are attached to a SCSI interface, can give distance, scalability, and cost advantages.

Today Gigabit Ethernet uses the optical piece of the Fibre Channel physical level as its base technology. This fact should lead to a common infrastructure that will allow customers to choose the right solution for their applications without having to rewire the premises. The ability to use common parts and infrastructure will lower the cost of the solution.

Flexibility and Scalability

The ability of Fibre Channel to handle different protocols simultaneously on the same physical hardware is a major benefit available today. Equally important, there is lots of room for future enhancements to Fibre Channel. Higher bandwidth, different media, additional protocol mappings, and new topologies are a few of the ways Fibre Channel technology can evolve in the future.

Status of Fibre Channel in the Industry

In its publication *Fibre Channel FAQ* (May 1996), the Fibre Channel Loop Community noted that Fibre Channel is becoming the "defacto connectivity standard for high-speed storage access and server clustering, and is a natural solution for gigabit enterprise backbones, and gigabit LANS for high-speed storage, image, video and mass data transfer applications." Approximately two-thirds of the storage industry now supports the Fibre Channel standard, including all major disk drive vendors. Essential Fibre Channel components are now becoming available: protocol chips, 1-Gb/s transceivers, media modules, cables, connectors, test equipment, and device drivers. Fibre Channel subsystems are now in development: disk drives, drive arrays, host adapters, RAID controllers, fabric switches, and storage subsystems.

There is considerable activity in the Fibre Channel Working Group to enhance and improve the architecture of FC-AL. However, to protect current and future customer investment, maintaining backward compatibility is a firm requirement.

COMPAQ'S INVOLVEMENT WITH FIBRE CHANNEL

Compaq actively participates in the following industry standardization efforts for interconnect technology:

- Technical Committee T10 of the ANSI Accredited Standards Committee NCITS (National Committee for Information Technical Standards), formerly known as X3T10 and responsible for SCSI
- Technical Committee T11 of the ANSI Accredited Standards Committee NCITS, formerly known as X3T11 and responsible for Fibre Channel
- Fibre Channel Association (FCA)
- Fibre Channel Loop Community (FCLC)
- Component Supplier Standardization

To ensure interoperability of its products, Compaq is participating in joint testing with the other members of FCLC and FCA, the University of New Hampshire, and the University of Minnesota.

FIBRE CHANNEL: A NEW PARADIGM

Fibre Channel technology enables a new paradigm. It eliminates the distance barrier in storage configurations and allows storage to be distributed, thereby enabling higher aggregate bandwidth, redundancy, and scalability of capacity, cache, and I/O processing power for the enterprise. At the same time, it maintains capability for enhanced, centralized control and management of both primary and secondary storage. Fibre Channel enables the future coexistence of networking and storage attachment using the same infrastructure and physical components. Fibre Channel also serves as the basis for architecting third party transfers, by which the client can move data directly from or to the source without passing the data through the host bus memory system. Third party transfers free up these valuable resources for use by the host processors.

CONCLUSION

The enterprise computing platform of the future will encompass industry standard hardware, software, and interconnection. Compaq is basing its strategic direction for high-performance, high-availability external storage on the conviction that Fibre Channel is the best interconnect technology to overcome I/O and physical limitations of SCSI technology and to meet the expanding needs for network data storage.

Compaq continues to lead the change to open-system industry standards. As a technology leader, Compaq proactively contributes to development of new industry standards and rapidly incorporates them into innovative products fitting the high-volume business model and offering cost-effective solutions to our customers.