NCS TIB 99-3

NATIONAL COMMUNICATIONS SYSTEM

TECHNICAL INFORMATION BULLETIN 99-3

GIGABIT NETWORKING

JANUARY 1999

OFFICE OF THE MANAGER
NATIONAL COMMUNICATIONS SYSTEM
701 SOUTH COURT HOUSE ROAD
ARLINGTON, VA 22204-2198

GIGABIT NETWORKING

Office of the Manager
National Communications System

February 1999

Communication Technologies, Inc.
503 Carlisle Drive, Suite 200
Herndon, Virginia 20170
703-318-7212 (Voice)
703-318-7214 (Fax)
Web: www.comtechnologies.com

NCS TECHNICAL INFORMATION BULLETIN 99-3

GIGABIT NETWORKING

JANUARY 1999

PROJECT OFFICER

DALE BARR, JR.
Electronics Engineer
Technology and Standards Division

APPROVED FOR PUBLICATION:

PETER M. FONASH
Chief, Technology
and Standards Division

FOREWORD

Among the responsibilities assigned to the Office of the Manager, National Communications System, is the management of the Federal Telecommunications Standards Program. Under this program, the NCS, with the assistance of the Federal Telecommunications Standards Committee identifies, develops, and coordinates proposed Federal Standards which either contribute to the interoperability of functionally similar Federal telecommunications systems or to the achievement of a compatible and efficient interface between computer and telecommunications systems. In developing and coordinating these standards, a considerable amount of effort is expended in initiating and pursuing joint standards development efforts with appropriate technical committees of the International Organization for Standardization, the International Telecommunication Union-Telecommunications Standardization Sector, and the American National Standards Institute. This Technical Information Bulletin presents an overview of an effort which is contributing to the development of compatible Federal and national standards in the area of Asynchronous Transfer Mode Standardization. It has been prepared to inform interested Federal and industry activities. Any comments, inputs or statements of requirements which could assist in the advancement of this work are welcome and should be addressed to:

Office of the Manager
National Communications System
Attn: N6
701 S. Court House Road
Arlington, VA 22204-2198

GIGABIT NETWORKING

Abstract

Gigabit networking is expected to grow rapidly in popularity over the next several years. It is anticipated that as the communications infrastructure evolves to support gigabit speeds, new applications such as telemedicine, desktop videoconferencing, scientific modeling, and interactive telecollaboration will evolve to become important commonplace capabilities. These and many other bandwidth intensive applications have great potential to improve National Security/Emergency Preparedness (NS/EP) communications and event management activities. This report examines existing and emerging gigabit networking technologies, addresses issues associated with their applications, and discusses their applicability to NS/EP environments.

Table of Contents

Figures

Tables

Appendices

1 Introduction

In April 1984 the signing of Executive Order (E.O.) 12472, *Assignment of National Security and Emergency Preparedness (NS/EP) Telecommunications Functions*, changed the mission focus of the National Communications System (NCS). Since that time the NCS has been assisting the President and the Executive Office of the President (EOP) in exercising wartime and non-wartime emergency telecommunications and in coordinating the planning for, and provisioning of, NS/EP communications for the Federal Government under all circumstances. In this regard, the Office of the Manager, NCS (OMNCS), particularly its Technology and Standards Division (N6), always seeks to improve the Federal Government's ability to respond to national security and emergency situations. As part of this mission the N6 division identifies new technologies that enhance NS/EP communications capabilities and ensures key NS/EP features such as priority, interoperability, reliability, availability, and security are supported by emerging standards. In concert with this approach, the N6 manages the Federal Telecommunications Standards Program. Additionally, the N6 division directs efforts in both NS/EP management and applications services

Gigabit networking is one of the new technologies identified by the N6 division for enhancing NS/EP communications. Accordingly, this report provides an analysis of the gigabit networking technology. The report covers the basic concepts of gigabit networking, provides an overview of gigabit networking technologies, applications of gigabit networking into NS/EP environment, and issues associated with the use or applications of gigabit networking.

1.1 NS/EP Event Characteristics

In order to more fully appreciate the importance of the NCS mission with respect to influencing the evolution of technologies, such as gigabit networking, it is helpful to focus on some of the characteristics that are associated with an NS/EP event. Essentially, NS/EP events are characterized as extreme events. They cause significant disruption and place lives and property at risk. Some NS/EP events arise from natural disasters such as earthquakes, hurricanes, fires, and floods. Man-created NS/EP events can be accidental, such as oil spills or the release of toxic substances, or they may be intentional, such as bombings by terrorists. An NS/EP event requires an immediate and a coordinated application of resources, facilities, and efforts beyond those customarily available to handle routine problems.

The challenges confronting NS/EP managers will often be extreme in several dimensions simultaneously. They may require extraordinary quantities of resources including: search and rescue teams, medical assistance, food, and shelter. The situational demands are highly diverse creating the need for cooperation among many differently trained professions. Many kinds of NS/EP events will be largely unpredictable in terms of location, time, and the specific resources needed. Furthermore, the urgency associated with NS/EP crises events often involves numerous associated implications such as the need to rapidly identify, collect, and integrate crucial information about the developing situation. It is also important to have instant access to tools and resources that are not cumbersome or difficult to use under stressful conditions. Lastly, it can be important to have the capability to make projections and estimations, as well as to initiate actions in the face of uncertainty and with a known incompleteness of information.

Today, NS/EP requirements research needs to begin to include a simultaneous focus on applications-motivated computer science and engineering topics that combine NS/EP event management with other information technologies. For example, these can include:

a) communications resources such as rapidly deployable, self-configuring wireless networks that are useful for coordinating the actions of multiple response teams;
b) decision support systems and tools to assist NS/EP crisis managers in making decisions in the absence of complete, and/or reliable information;
c) computer simulations of phenomena such as hurricanes and fires that can deliver useful predictions and other information to NS/EP managers rapidly are required.

It is clear from the aforementioned examples that NS/EP experimental testbeds and simulators for critical event management-related research and development must remain high priority items. NS/EP event testbeds that provide a realistic application setting, such as simulation combined with field-based training exercises, can serve as model implementation environments for new technologies and can serve as sources of feedback to identify and refine the application objectives. Application users, including federal, state, and local civilian NS/EP personnel should participate in testbeds and simulation activities to the maximum practical extent possible.

To derive the full benefits of application-specific computing and communications technologies, NCS must capitalize on the increasingly interconnected nature of national-scale applications. In application areas such as NS/EP event management and health care the widespread interconnection of computing and information resources and the people who use them over networks has made it feasible, and increasingly common, for resources to be called on in unforeseen ways. NS/EP event management, in particular, illustrates the value of being able to integrate highly diverse resources whose usefulness in an unusual situation could not have been anticipated in advance. Therefore, interoperability and standards must continue to play a crucial enabling role in order to facilitate dynamic multivendor applications and communications.

In recognition of the importance that evolving applications and high-end computer technologies will have on the future of both gigabit networking and NS/EP event management, this paper provides a brief examination of both the communications infrastructure changes and, with equal importance, the associated NS/EP user application requirements.

2 Background

Over the last few years several factors have placed a severe demand on computer networks for high speed and large bandwidth. Some of these factors are:

a. Growing number of data network users: In recent years there has been a tremendous increase in the number of data network users. This growing number of users is significantly increasing traffic on networks. Many enterprises have experienced 200% to 300% traffic growth per year as a result of this increase. Presently, there are serious examinations under way with respect to future Internet requirements to support over 1 billion users.

b. Changing network traffic pattern: The explosive growth of Internet, enterprise intranets, World Wide Web, and a demand to communicate from anywhere to anywhere has radically changed the traffic pattern of many corporate backbones and Wide Area Networks (WANs). More and more network users are constantly accessing servers from many sub-networks in various geographic locations. As a result, the majority of network traffic runs over corporate backbones and WANs rather than Local Area Networks (LANs). For example, when hurricane Bonnie struck the Carolina coast, an unprecedented number of people tried to access Federal Emergency Management Agency's (FEMA's) web site. The number of files being accessed on the FEMA's web site daily rose from 200,000 to almost two million. The agency's T-1 communications line was not capable of handling the additional volume. To handle the situation, the agency had to expand the site's capacity to 10 Mb/s within a few hours.

c. Faster desktop computers and network servers: The development of faster Central Processing Units (CPUs), faster memory, and high-speed buses have provided desktop computers, workstations, and network servers the capability of processing raw Input/Output (I/O) with a throughput in the order of 100's of Mb/s to Gb/s. When these faster computers or servers are connected to a relatively slower network, a bottleneck occurs between the computers or servers and the network. For example, a 1,000 Mega Instructions Per Second (MIPS) computer needs 1,000 Mb/s (MIPS=Mb/s) I/O. Therefore, a relatively high number of CPU interrupts per transfer happens if a computer is required to adapt itself to a slower network.

In the future NS/EP managers will require the assistance of supercomputer level network servers for a vast array of emergency situation assisting applications. Current supercomputer computation research suggests that the next major advance will be a sustained rate of one petaflops (pflops — 10^{15} floating-point operations per second). This milestone is expected to be achieved by 2007, assuming that certain key technologies continue to progress at current rates. It is further estimated that the required applications and data storage capacity associated with these systems will be in the exabyte (10^{18}) range.

d. **Bandwidth-intensive applications:** There is an enormous increase in the use of bandwidth-intensive applications. These applications range from scientific modeling, medical data transfers, and data warehousing to desktop video conferencing. As shown in Table 1, these applications demand higher-bandwidth for desktops, servers and backbones.

Application	Data Types / Size	Network Traffic Implication	Network Need
Scientific Modeling, Engineering	Data files; 100’s of Megabytes to Gigabytes	Large files increase bandwidth required	Higher bandwidth for desktops, servers, and backbone
Publications, Medical Data Transfer	Data files; 100’s of Megabytes to Gigabytes	Large files increase bandwidth required	Higher bandwidth for desktops, servers, and backbone
Internet/Intranet	Data files now; Audio now; Video will emerge; High Transaction Rate; Large Files, 1 MB to 100 MB	Large files increase bandwidth required; Low transmission latency; Class of service reservation; High volume of data streams	Higher bandwidth for servers and backbone; Low latency
Data Warehouse	Data files; Gigabytes to terabytes	Large files increase bandwidth required; Search and access require low latency	Higher bandwidth for servers and backbone; Low latency
Network Backup	Data files; Gigabytes to terabytes	Large number of large files; Transmitted during fixed time period	Higher bandwidth for servers and backbone; Low latency
Desktop Video Conferencing, Interactive Whiteboard	Constant Data Stream; 1.5 to 3.5 Mb/s at the desktop	Class of service reservation; High volume of data streams	Higher bandwidth for servers and backbone; Low latency; Predictable Latency

Table 1: Applications Driving Network Growth

Consequently, computer networks operating at very high speeds are becoming important for interconnecting faster desktops and servers, clusters of workstations and other high-end equipment. Today's advanced technologies in fiber optics, computing systems and networking have made the development of such networks possible. This is ascertained by existing and emerging gigabit networking technologies. These technologies provide connections between workstations, servers, and peripherals at a signalling rate of at least 1,000 megabits per second (Mb/s). Gigabit Ethernet, Fibre Channel, High Performance Parallel Interface (HIPPI)-800, HIPPI 6400 (also known as Gigabyte System Network (GSN)), Asynchronous Transfer Mode (ATM)-OC12 (622.08 Mb/s), and ATM-OC48 (2.488 Gb/s) are all examples of technologies that operate at, or above, these emerging speeds.

3 Gigabit Networking Technologies

This section provides an overview of existing and emerging gigabit networking technologies such as fibre channel, HIPPI, HIPPI-Serial, HIPP-6400, Gigabit Ethernet, and ATM.

3.1 Fibre Channel

Fibre channel is a highly reliable gigabit networking technology. It is designed to significantly improve the speed at which data are transferred between workstations, mainframes, supercomputers, and storage devices. The fibre channel was originally intended to support fiber optic cable only. A few years ago copper support was added to fibre channel's feature set. To lessen the association with fiber optic, an International Organization for Standardization (ISO) task force changed the spelling of "fiber" to the French spelling "fibre". Since then "fibre" has been used as a generic term that can indicate either an optical or a copper cable. The fibre channel supports transfer rates of 133 Mb/s, 266 Mb/s, 530 Mb/s, 1 Gb/s, 2 Gb/s, and 4 Gb/s. The following table summarizes supported distances and speed over the respective media types.

Media	Speed	Distance
Coax/Twinx – Copper	1.0625 Gigabit	24 Meters
	266 Megabit	47 Meters
9 Micrometer Single mode Fiber	1.0625 Gigabit	10 Kilometers
50 Micrometer Multi-mode Fiber	1.0625 Gigabit	300 Meters
	266 Megabit	2 Kilometer
62.5 Micrometer Multi-mode Fiber	266 Megabit	1 Kilometer
	132 Megabit	500 Meters

Table 2: Supported Distances and Speed

Fibre channel combines the best features of channels and networks. It combines simplicity, repeatable performance and guaranteed delivery features of channels with the connectivity, distance, and protocol multiplexing features of networks. The fibre channel supports different protocols such as Internet Protocol (IP), Fiber Distributed Data Interface (FDDI), Serial HIPPI, Small Computer System Interface (SCSI) and Intelligent Peripheral Interface (IPI). It also supports multiple interoperable topologies such as point-to-point, switched fabric, and arbitrated loop. Fibre channel networks provide enterprises with new levels of performance and reliability. Typical network applications of fibre channel include gigabit LANs, enterprise backbones, digital audio/video networks, quick-response networks for imaging applications, and high performance Computer Aided Design/Computer Aided Engineering (CAD/CAE) networks.

3.1.1 Fibre Channel Architecture

As shown in Figure 1, fibre channel architecture is defined as a five-layered stack covering various functions. These layers are called FC-0, FC-1, FC-2, FC-3, and FC-4. They encompass physical media and transmission rates (FC-0), encoding scheme (FC-1), framing protocol and flow control (FC-2), common services (FC-3), and upper layer protocol interfaces (FC-4). All of these layers are briefly described in the following sections.

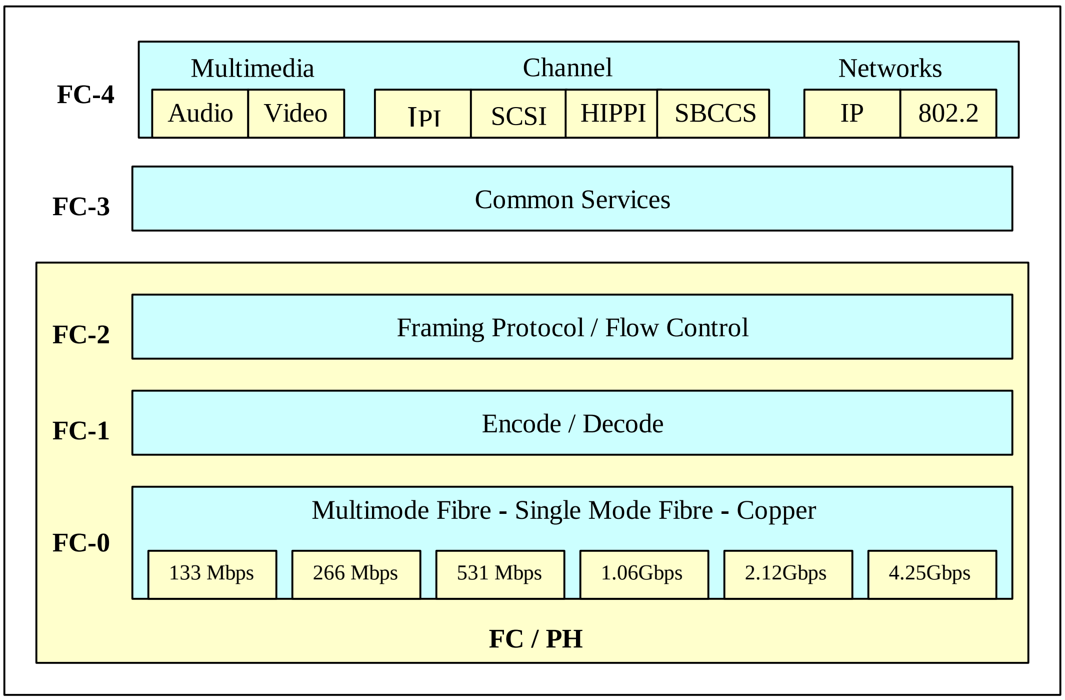

Figure 1: Fibre Channel Architecture

3.1.1.1 FC-0

This layer is also known as the physical interface and media layer. It specifies the physical link of the channel, including cables, connectors, drivers, transmitters, and receivers that can be used with fibre channel.

3.1.1.2 FC-1

FC-1 is also referred to as the transmission protocol layer. It defines a dc-balanced 8B/10B coding scheme for clock recovery, byte synchronization, encoding, and decoding. In this scheme, every 8 bits of data are encoded into a 10-bit transmission character. Character conversion is achieved by forming a transmission character combination in the form Zxx.y. This combination represents a valid transmission character. It is formed by taking an 8-bit information, logically labeled A, B, C, D, E, F, G, H, and using a control character Z. For example, each binary value represented by E, D, C, B, A is converted into the form xx, which is a decimal representation of the binary value, i.e. $2^5 -1 = 31$ that requires two decimal digits. Similarly, H and G are converted into the form y, which is a one digit decimal value. Finally, the control character Z represents data-type (D) or special-type (K). After transmission, D type transmission characters are decoded into one of the 256 eight-bit combinations while K type transmission characters are used for protocol management functions. All other codes besides D or K types are considered as invalid codes.

The physical layer also uses a Runnig Disparity (RD). It is a binary number that is calculated based on the number of 1's and 0's in the two sub-blocks of a transmission character. The first six bits form the first sub-block while the last four bits form the second sub-block of a transmission character. A new RD value is circulated at the transmitter and receiver and if the RD value is not the same, a disparity violation condition is indicated.

3.1.1.3 FC-2

FC-2 is called the framing and signalling layer. It mainly defines the framing structure of the data to be sent from one port to another. Additionally, it provides the following functions.

1. A 32-bit Cyclic Redundancy Check (CRC) to detect transmission errors and ensure data integrity.

2. Various classes of service to provide and manage circuit switching, frame switching, and fractional bandwidth virtual circuits. The supported service classes are Class 1, Class 2, Class 3, Class 4 and Class 6. These service classes are explained in detail in section 3.1.1.5.1.

3. Multiplexing of operations.

4. A credit-based flow control scheme to provide a guaranteed delivery capability. Flow control is either buffer-to-buffer or node-to-node or both depending on which service class is used. Table 3 shows the type of flow control associated with each service class.

Service Class	Flow Control
Class 1	Buffer-to-buffer Node-to-node
Class 2	Buffer-to-buffer Node-to-node
Class 3	Buffer-to-buffer
Class 4	Buffer-to-buffer Node-to-node
Class 6	None

Table 3: Service Classes and Flow Control

5. A set of generic functions common to multiple upper-layer protocols.

6. A built in protocol to aid in managing the operation of the link, control the fibre channel configuration, perform error recovery, and recover link and port status information.

7. Optional headers that may be used for network routing.

8. Control information in the header to assist hardware routing.

9. Segmentation and reassembly of data.

FC-2 defines a set of four building blocks to provide efficient data transport across the links. These building blocks are ordered set, frame, sequence, and exchange.

1. Ordered set

Ordered sets are transmission words. Each word is comprised of four 10-bit transmission characters. These words are used for bit or word synchronization and word boundary alignment. There are three types of ordered sets: Frame delimiters, Primitive signals, and Primitive sequence.

2. Frame

Figure 2 shows a frame structure consisting of a start-of-frame, a frame header, optional headers, payload, a 32-bit CRC, and end-of frame.

4 Bytes	24 Bytes	2112 Bytes Payload		4 Bytes	4 bytes
Start of Frame	Header	64 Bytes Optional Header	2048 Bytes Payload	CRC	End of Frame

Figure 2: Frame Structure

Addressing is done within the frame header and contains source and destination addresses. The payload field is 2112 bytes long, but if the optional header is used then it is 2048 bytes long. A frame can be either a data frame or a link control frame. The data frames are further categorized into link data frames and device data frames. The link control frames are further classified into acknowledge and link response (busy and reject) frames.

3. Sequence

A sequence is a set of one or more related frames for a single operation, flowing in the same direction on the link. Each frame within the sequence is uniquely numbered with a sequence count. While each sequence is uniquely identified by the initiator of the sequence utilizing the sequence identifier field within the frame header. The FC-2 layer is responsible for breaking a sequence into the frame size negotiated between the communicating ports. A sequence also serves as the recovery boundary in fibre channel. When an error is detected, fibre channel identifies the sequence in error and allows that sequence and subsequent sequences to be retransmitted.

4. Exchange

An exchange consists of one or more non-concurrent sequences for a single operation. For example, a single operation may involve a command to read data, transfer data, and then provide a status report of the operation. In this example, all the three phases are separate sequences but they comprise a single exchange. Furthermore, within a single exchange, only a single sequence may be active at any one time. However, sequences for different exchanges may be active at the same time. To manage the exchange, an exchange ID is assigned by both the originator and the responder in an implementation-dependent manner. The exchange IDs are contained within the frame header and are used locally.

3.1.1.4 FC-3

FC-3, also known as the common services layer, defines functions that span multiple ports on a single node or a fabric created by linking multiple switches together. Currently defined or proposed functions include:

a. Multicast: Delivers a single transmission to multiple destination ports.

b. Hunt Groups: A hunt group is a set of associated ports attached to a single node. This set is assigned an alias identifier that allows any frame containing this alias to be routed to any non-busy port within the set. This reduces the chances of reaching a busy port and helps improve efficiency.

c. Striping: Provides higher bandwidth using multiple ports in parallel to transmit a single information unit across multiple levels.

3.1.1.5 FC-4

This layer specifies the application interfaces that can be supported over fibre channel. Both network and channel protocols can be concurrently transported over the same physical interface. Currently specified or proposed network and channel protocols include:

a. Small computer system interface (SCSI);
b. Intelligent peripheral interface (IPI);
c. High performance parallel interface (HIPPI) framing protocol;
d. Internet protocol (IP);
e. Link Encapsulation (FC-LE) using IEEE 802.2;
f. Single byte command code set (SBCCS) mapping;
g. Audio video fast file transfer;
h. Audio video real time stream transfer.

Proprietary protocols can also be supported over the fibre channel, but they are outside the scope of the standard.

3.1.1.5.1 Classes of Service

The FC-2 layer supports five different classes of service. These classes are as follows:

1. Class 1: Acknowledged connection service;
2. Class 2: Acknowledged connectionless service;
3. Class 3: Unacknowledged connectionless service;
4. Class 4: Fractional bandwidth connection oriented service;
5. Class 6: Simplex connection service.

3.1.1.5.1.1 Class 1

This class provides reliable and guaranteed delivery of data. An end-to-end connection between the communicating devices is established through switches. Accordingly, class 1 offers circuit switched dedicated bandwidth connections. It also provides an acknowledgment of receipt for a guaranteed delivery. For this class, the flow control is buffer-to-buffer and node-to-node. The only overhead associated with class 1 is the connection setup and release time. However, in fibre channel it is in microseconds. This class is particularly suitable for large data transfer.

Class 1 also offers special features such as "Camp On" and "Stacked Connect" to make switching service more efficient. The camp on feature enables a switch to monitor a busy port and queue that port for the next connection. As soon as the port is available, the switch establishes the connection. The stacked connect option allows an originating port to queue sequential connection requests with the switch. Additionally, class 1 provides buffered service and dedicated simplex service. The buffered service is used to connect two ports that are operating at different speeds. By default, class 1 connections are bi-directional. However, simplex connections can be used to separate transmit and receive switching. This allows one node to transmit data to another node and to receive data form a third node at the same time.

3.1.1.5.1.2 Class 2

This class offers connectionless service with guaranteed delivery. All frames are switched independently and the path between two interconnected devices is not dedicated. Class 2 uses buffer-to-buffer and node-to-node flow control to eliminate congestion. Because of the connectionless service, this class has more overhead compared to class 1. Moreover, typical class 2 frame latency is less than one microsecond. These factors make this class ideal for small data transfers.

3.1.1.5.1.3 Class 3

This class also offers connectionless service like class 2, but it does not send an acknowledgment to confirm frame delivery. This class is mainly used for multicast and broadcast on networks.

3.1.1.5.1.4 Class 4

This class provides a fractional bandwidth connection oriented service. All connections are bi-directional and support a different set of Quality of Service (QoS) parameters in each direction. The QoS parameters include guaranteed bandwidth and latency. A QoS facilitator (QoSF) function is provided within fibre channel switches to maintain the negotiated QoS on each connection and in each direction. The fibre channel switches pace frames from the source node to the destination to regulate available bandwidth. This class is used for time critical applications such as real time audio and video.

3.1.1.5.1.5 Class 6

This class is similar to class 1, but provides simplex service. Additionally, it supports multicasting and preemption. This class is designed for video broadcast applications and real-time systems that move large quantities of data.

Table 4 summarizes the above referenced service classes and their typical usage.

Service Class	Delivery	Usage
Class 1	Reliable and guaranteed	Large data exchanges
Class 2	Guaranteed	Small data transfer
Class 3	Reliable	Broadcast Multicast
Class 4	Reliable and guaranteed	Real time audio and video
Class 6	Reliable	Video broadcast

Table 4: Fibre Channel Classes of Service

3.1.2 Topologies

Fibre channel specifications define three physical topologies to connect communicating devices. These topologies are point-to-point, arbitrated loop, and switched fabric.

3.1.2.1 Point-to-Point

This topology is very simple and gives the best possible bandwidth of all topologies. It can be configured in two ways. The first is to set up a fixed point-to-point connection between two communicating devices. The second is to set up point-to-point connections using an externally controlled link switch. The ports in this topology are called N_ports and the communication is full duplex.

3.1.2.2 Arbitrated Loop

This topology is similar to token ring where each node on the loop arbitrates for loop access. Once the access is granted, the requesting node gets all the bandwidth to make logical point-to-point connections. An arbitrated loop may connect up to 127 nodes and only one of them may be a switch port. Nodes request control of the loop by sending a signal called a "primitive". If the signal is returned with the requesting node's address, then that node gets the control. If two or more nodes contend for the loop control at the same time, then the node with the lowest address gets the control. The ports in this topology are called NL-ports. The loop is self-configuring and may operate with or without a switch present. A loop node or switch port self-discovers its environment and works with other nodes without manual intervention. If a node fails, the loop is out of operation. This is avoided by using a hub. The hub uses a port bypass circuit (PBC) to ensure that the loop is operational. The PBC either opens the loop to insert an active node or closes the loop to remove failed nodes.

3.1.2.3 Switched Fabric

In this topology, each communicating device is connected to a switch and receives a non-blocking path to any other connection on the switch. This is equivalent to a dedicated connection to every device. If multiple switches are used to accommodate more devices, then the switches are connected together. Fibre channel switches may offer connection-oriented switching, connectionless switching, or both.

3.1.3 Fibre Channel Services

Fibre channel services, typically found in switched topologies, specify a set of functions required by fibre channel protocols. Various servers within the switch handle these functions. These servers are login server, fabric/switch controller, management server, time-server, aliases server, QoS server, name server, and multicast server. The management server uses the most widely implemented Simple Network Management Protocol to mange fibre channel networks. Management Information Base (MIB) data are associated with each fibre channel node, hub, and fabric and the SNMP protocol is used to monitor or modify the MIB data. The SNMP may use IP over FC-PH as its transport mechanism. The QoS server is used to set the desired QoS established for class 4 service on all VCs. The QoS parameters are guaranteed minimum bandwidth and guaranteed maximum end-to-end delay on each VC. These parameters are negotiable for each VC.

3.1.4 Fibre Channel Standardization

Various standards organizations such as the American National Standards Institute (ANSI), the Internet Engineering Task force (IETF), the Fibre Channel Association (FCA), and the Fibre Channel Systems Initiative (FCSI) are working on the standardization of fibre channel. The ANSI, particularly ASC (Accredited Standards Committee) X3T11, is primarily responsible for the fiber channel standardization. The list of approved or proposed standards for fibre channel can be found in Appendix C of this report.

3.2 HIPPI

HIPPI, now known as HIPPI-800, was originally developed at the Los Alamos National Laboratory, Los Alamos, New Mexico, to transfer data between supercomputers and data storage machines at extremely fast rates. In 1987, it was proposed to the ANSI as a standard for transferring data at gigabit speeds. Eventually in 1991, the HIPPI became a standard. It is a connection-oriented protocol and operates in a simplex mode. Therefore, most installations use a pair of connections to achieve a full-duplex mode. The HIPPI uses a parallel data path with copper cables and allows speeds of 800 Mb/s and 1.6 Gb/s for distances of up to 25 meters. The 800 Mb/s HIPPI uses a 32-bit data bus while the 1600 Mb/s HIPPI uses a 64-bit data bus. The HIPPI specifications also define a few reverse-direction control signals for connection setup and flow control. Although the HIPPI was first initiated for supercomputer applications, it could be used to connect different types of devices, such as workstations, monitors, frame buffers, etc. It supports point-to-point connections and uses switches to construct networks.

3.2.1 HIPPI Architecture

A layered concept has been used to define the HIPPI protocol architecture. In relation to the Open Systems Interconnection (OSI) Basic Reference Model, the HIPPI covers the physical layer and a small portion of the data-link layer. Table 5 shows the HIPPI layered architecture covering different protocols. A brief description of these protocols is presented in the sections that follow.

<table>
<tr><td>HIPPI-LE
(Link Encapsulation (LE), Mapping to IEEE 802.2)</td><td>HIPP-FC
(Mapping to Fibre Channel (FC))</td><td>HIPPI-IPI
(Mapping to Intelligent peripheral interface (IPI)</td><td>HIPPI-ATM
(Mapping to ATM)</td></tr>
<tr><td colspan="4">HIPPI-FP
(HIPPI-Framing Protocol (FP))</td></tr>
<tr><td colspan="2">HIPP-PH
(HIPPI - Physical Layer (PH))</td><td colspan="2" rowspan="2">HIPPI-SC
(HIPPI-Switch Control (SC))</td></tr>
<tr><td colspan="2">HIPPI-Serial
(HIPPI-Serial - Physical Layer)</td></tr>
</table>

Table 5: HIPPI Layers

3.2.1.1 HIPPI-PH

HIPPI-PH defines the mechanical, electrical, and signalling specifications of the HIPPI physical layer. It is designed for transmitting digital data at peak data rates of 800 or 1600 Mb/s between data-processing equipment using multiple twisted-pair copper cabling at distances up to 25 meters. The HIPPI-PH signalling protocol is designed to be distance independent, allowing the average data rate to approach the peak data rate, even over distances longer than those specified for the HIPPI-PH.

Figure 3 shows the basic organization of the information, called logical framing hierarchy, on the HIPPI-PH.

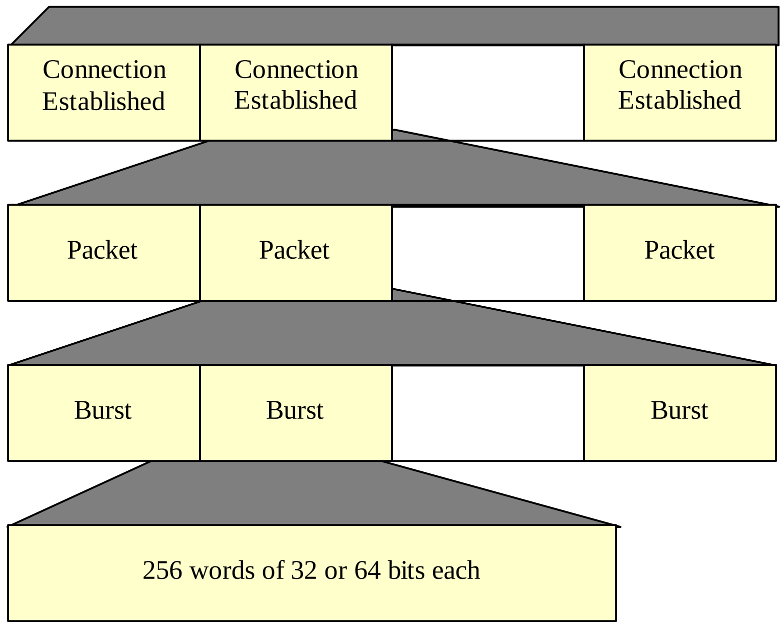

Figure 3: Logical Framing Hierarchy

As specified in the HIPPI-PH, once a connection is established packets can be sent from the source to the destination. Each packet contains one or more bursts. Each burst contains 1 to 256 words and a maximum of 1 Kbyte of data. Therefore, a full burst is either 1 Kbytes with the 32-bit HIPPI-PH or 2 Kbytes with the 64-bit HIPPI-PH. Bursts that contain less than 256 words are called short bursts. A short burst can be either the first or last burst of a multi-burst packet. A packet shall contain no more than one short burst. The packets are organized into classes based on the order of short bursts and full bursts. These classes are:

a. Class 1: Includes all packets that are composed of a single short burst;

b. Class 2: Includes all packets that are composed of a short burst followed by one or more full bursts;

c. Class 3: Includes all packets that are composed of one or more full bursts followed by a short burst;

d. Class 4: Includes all packets that are composed of one or more full bursts and no short bursts.

3.2.1.1.1 Error Detection and Recovery

For error detection, HIPPI-PH uses byte parity and a parity-based Length-Longitudinal Redundancy Check (LLRC) on each burst. The HIPPI-PH does not define any error correction capabilities. That is left for upper layer protocols.

a. Byte Parity: In this method, an odd-parity bit is transmitted with each byte of a word. For example, four parity bits are transmitted with each 32-bit word. Hence an undetected error in a word would require a 2-bit error, with both bits being in the same byte.

b. LLRC: This method implements even parity across the individual bits of multiple words in a burst. For example, bit 23 of the LLRC is the even parity of bit 23 of each word in the burst. As mentioned earlier, a burst is usually 256 words in length, but short bursts may contain fewer words. Hence, the LLRC would not detect errors where the same bit in an even number of words was incorrect.

The LLRC calculation also includes the length of bursts. Therefore, the LLRC would detect cases where a word was dropped or added. In other words, the length received was not the same as what was transmitted.

3.2.1.2 HIPPI-SC

A major limitation of the HIPPI-PH is that it supports only a single point-to-point connection. HIPPI-SC was designed to overcome this limitation. It defines the control for HIPPI physical layer switches. HIPPI physical layer switches may be used to give the equivalent of multi-drop capability, connecting together multiple data processing devices. It supports both source routing and destination addressing. Each point-to-point connection is established by the address associated in the I-field. An I-field is a 32-bit field that is sent as part of the sequence of the physical layer operations establishing a connection from a source to a destination. HIPPI-SC is closely coupled with HIPPI-PH.

3.2.1.3 HIPPI-FP

HIPPI-FP defines the format and content of HIPPI packets. As shown in figure 4, HIPPI packets are composed of three areas: Header Area, D1 Area and D2 Area. The header area consists of two words or 64 bits and must be contained completely in the first burst. There are seven variables contained in this area. They are ULP-id, P, B, Reserved, D1 Area Size, D2 offset, and D2 size. The D2 Size is contained in the second word.

1. ULP-id (8 bits): Designates the destination ULP to which the packet is to be delivered.

2. P (1 bit): Designates that a D1_Data_Set is present (P = 1) in this packet

3. B (1 bit): B = 0 designates that the D2_Area starts at or before the beginning of the second burst of the packet. B = 1 designates that the D2_Area starts at the beginning of the second HIPPI-PH burst of the packet.

4. D1 area size (8-bits): Designates the size of the D1_Area.

5. D2_Offset (3 bits): Designates the number of offset bytes from the start of the D2_Area to the first bytes of the D2_Data_Set.

6. Reserved (11 bits): All the reserved bits are transmitted as zeros.

7. D2_Size (32 bits): Designates the length, in bytes, of the D2_Data_Set portion of the packet.

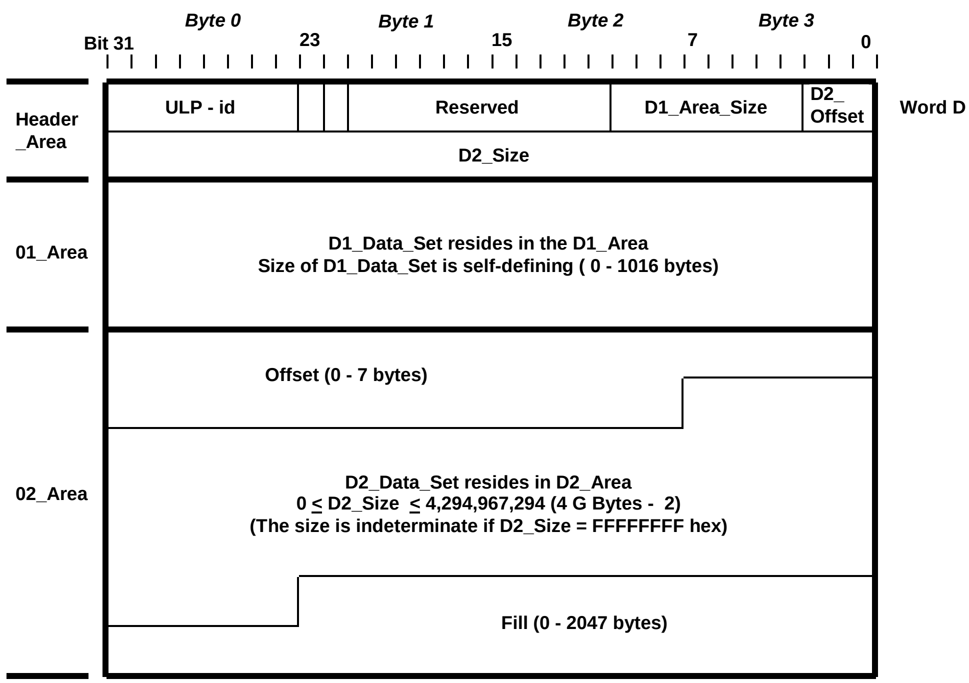

Figure 4: HIPPI-FP Packet Format

HIPPI-FP also splits higher layer packets into the 1 Kbytes or 2 Kbytes required by the HIPPI-PH.

3.2.1.4 Mapping to Other Protocols

Various specifications have been developed or are under development to map HIPPI to other protocols. These specifications are given below.

- **HIPPI-LE:** It provides an encapsulation of ISO/IEC 8802-2 (IEEE Std 802.2) Logical Link Control Protocol Data Units over the HIPPI.

- **HIPPI-FC:** It defines the frame format and protocol definitions required to transfer information for upper-layer protocols that use the HIPPI-FP when using a lower-layer serial link interface operating according to the Fibre Channel – Physical and Signalling Interface (FC-PH) and Fibre Channel – Enhanced Physical (FC-EP) requirements.

- **HIPPI-ATM:** It specifies the frame formats and protocol definitions for encapsulation of HIPPI-PH packets for transfer over Asynchronous Transfer Mode (ATM).

3.3 HIPPI-Serial

HIPPI supports 800 Mb/s or 1600 Mb/s for distances of up to only 25 meters. HIPPI-Serial is defined to overcome this limitation. It specifies a physical-level interface for transmitting digital data at 800 Mb/s or 1600 Mb/s serially over fiber-optic cables across distances of up to 10 km. The signalling sequences and protocol specified are compatible with HIPPI-PH. HIPPI-Serial may be integrated as a host's native interface, or used as an external extender for HIPPI-PH ports.

3.4 HIPPI-6400

HIPPI-6400 (now GSN) is the next generation HIPPI standard. It provides much more than an upgrade path for HIPPI networks. The HIPPI-6400 also embeds many of the high-latency protocol functions such as error detection, retransmission, flow control, and logical level message-size based multiplexing into the physical specification to improve throughput.

3.4.1 HIPPI-6400 Architecture

HIPPI-6400 architecture is similar to that of HIPPI and covers various protocols. These protocols are specified by the HIPPI-6400 working group and are described in three documents: HIPPI-6400-PH (physical layer), HIPPI-6400-SC (switch control), and HIPPI-ST (scheduled transfers). A brief description of these protocols is presented in the following sections.

3.4.1.1 HIPPI-6400-PH

HIPPI-6400-PH specifies a physical-level, point-to-point, full duplex, link interface for reliable, flow-controlled, transmission of user data at 6400 Mb/s, per direction, across distances of up to 1 km. It also specifies a parallel copper cable interface for distances of up to 40 meter.

3.4.1.1.1 HIPPI-6400 Messages

HIPPI-6400 messages consist of one or more micropackets. The micropackets are the basic units of data transfer from source to destination on HIPPI-6400 links. Each micropacket contains 32 bytes of data and 64 bits of control information. The control information in each micropacket includes parameters to select a Virtual Channel (VC), detect missing micropackets, indicate the last micropacket of a message, pass credit information from the source to the destination, link-level and end-to-end checksums, etc. Each micropacket is also given a sequence number for acknowledgment and retransmission purposes.

Mpacket Transfer Order ⇒			
	1	Bytes 0-7, Header Information	c63-c00
	2	Bytes 8-39, Message Data	c63-c00
	3	Bytes 40-71, Message Data	c63-c00
	-	- - - - - - - - - - - - - - - - - -	- - - - - -
	-	- - - - - - - - - - - - - - - - - -	- - - - - -
	n	Last Bytes of Message Data	c63-c00

Figure 5: HIPPI-6400 Message Format

As shown in figure 5, each HIPPI-6400 message starts with a special micropacket called a header micropacket. The header micropacket contains information to route data through HIPPI-6400 switches. The last micropacket of the message is marked with the TAIL bit.

3.4.1.1.2 Virtual Channels

HIPPI-6400-PH specifies four virtual channels, VC0, VC1, VC2, and VC3 in each direction on each link. These VCs are assigned to specific message sizes and transfer methods. According to the current assignment, VC0 is used for control messages, VC1 and VC2 are used for IP sized transmission (<128 Kbytes) and VC3 is used for scheduled transfer. All of the micropackets of a message are transmitted on a single VC. Messages to a final destination are delivered in order on a single VC. Multiple messages may be out of order if sent over different VCs even if the VCs are in the same physical link

3.4.1.1.3 Error Detection and Recovery

To ensure reliable delivery between two communicating devices, HIPPI-6400 employs two 16-bit CRC's. They are Link CRC (LCRC) and End-to-end CRC (ECRC). The LCRC uses the CRC-16 polynomial, $x^{16} + x^{12} + x^5 + 1$, and covers all of the data and control bits of micropackets in a message. The ECRC uses the polynomial $x^{16} + x^{12} + x^3 + x + 1$ and covers the data bits of all of the micropackets in a message. The CRCs in each micropacket are checked at the destination and correct micropackets are acknowledged while erroneous micropackets are discarded. Retransmission is used to send discarded micropackets back to the destination. It is on a hop-by-hop basis and the Go-back-N method is used for retransmission.

3.4.1.1.4 Network Management

HIPPI-6400 specifies "admin micropackets" for managing HIPPI-6400 networks. All admin micropackets consist of a single micropacket and use VC1 for requests and VC2 for responses. These micropackets perform various functions including switch and end-host auto-configuration, reconfiguring switch routing tables, Address Resolution Protocol (ARP), Reverse Address Resolution Protocol (RARP), providing logical addresses to end-points, and an optimized subset of the Simple Network Management Protocol (SNMP). Switches and host controllers take admin micropackets out of the data stream and process the requested command.

3.4.1.2 HIPPI-6400-SC

HIPPI-6400-SC defines the control for HIPPI-6400 physical layer switches. The HIPPI-6400 physical layer switches may be used to obtain the equivalent of multi-drop capability. The following are the major characteristics of HIPPI-6400-SC:

- Defines procedures for admin micropackets to automate Universal LAN Address (ULA) assignment;
- Enables spanning of multiple physical layer switches within a fabric;
- Includes support for physical layer switches with differing numbers of ports, all within the same fabric;
- Specifies reserved ULAs to aid address self-discovery, switch management, and switch control;
- Provides support for 4 virtual channels.

3.4.1.3 HIPPI-ST

HIPPI-ST specifies a connection-oriented data transfer protocol called Scheduled Transfer (ST) protocol. The protocol uses small control messages in order to pre-allocate buffers at the receiver before the actual data transfer begins. This allows the data to be moved immediately from the physical network into the end device’s memory. It uses Scheduled Transfer Units (STUs) as basic units for data transfer. The ST protocol also provides for retransmitting partial transfers for error recovery and mappings onto HIPPI-6400-PH, HIPPI-FP (for HIPPI-800 traffic), Ethernet, ATM, and Fibre Channel lower-layer protocols.

3.4.2 HIPPI, HIPPI-Serial, and HIPPI-6400 Standardization

Various standards organizations such as ANSI, ISO, and the National Committee for Information Technology Standardization (NCITS) are working on the standardization of HIPPI, HIPPI-Serial, and HIPPI-6400 to ensure interoperability and reliability. A number of standards have been developed or proposed for HIPPI, HIPPI-Serial and HIPPI-6400 and can be found in Appendix D of this report.

3.5 Gigabit Ethernet

For local area networking Ethernet technology has become the universally accepted standard. More than 85 percent of all installed network connections were Ethernet by the end of 1997. This represents over 118 million interconnected PCs, workstations and servers. The remaining network connections are a combination of Token Ring, Fiber Distributed Data Interface (FDDI), Asynchronous Transfer Mode (ATM) and other protocols. All popular operating systems and applications are Ethernet-compatible, as are upper-layer protocol stacks such as Transmission Control Protocol/Internet Protocol (TCP/IP), IPX, NetBEUI and DECnet. By the end of 1998 Ethernet technology accounted for 86 percent of all shipments. In contrast, ATM, FDDI/CDDI and Token Ring network interface card shipments combined were expected to reach just 5 million in 1998; only 10 percent of the total. ATM, FDDI/CDDI and Token Ring hub ports were expected to reach 4 million, approximately 7 percent of the total. Ethernet dominance is expected to continue well into the next century. However, as mentioned before, the demand for high-speed network connections is proliferating at a very rapid pace. Today, Fast Ethernet or 100-Mb/s Ethernet has become a commonplace solution to meet this demand. The growing use of Fast Ethernet connections has created a clear need for an even higher-speed network technology at the backbone and server level.

Gigabit Ethernet is one such technology providing seamless integration with Ethernet and Fast Ethernet. It transfers data at one gigabit per second (1,000 Mb/s) and is designed to deliver the same benefits as Fast Ethernet: seamless integration with existing Ethernet networks, dramatically higher performance and a familiar management environment. Gigabit Ethernet uses a mixture of proven protocol technologies adopted from both the original IEEE 802.3 Ethernet specification and the ANSI X3T11 Fibre Channel specification.

3.5.1 Gigabit Ethernet Architecture

Gigabit Ethernet architecture is based on the layer concept defined in the OSI Basic Reference Model. It covers the data link layer and physical layer of the reference model. Figure 6 shows an architectural model of Gigabit Ethernet specified in the IEEE 802.3z.

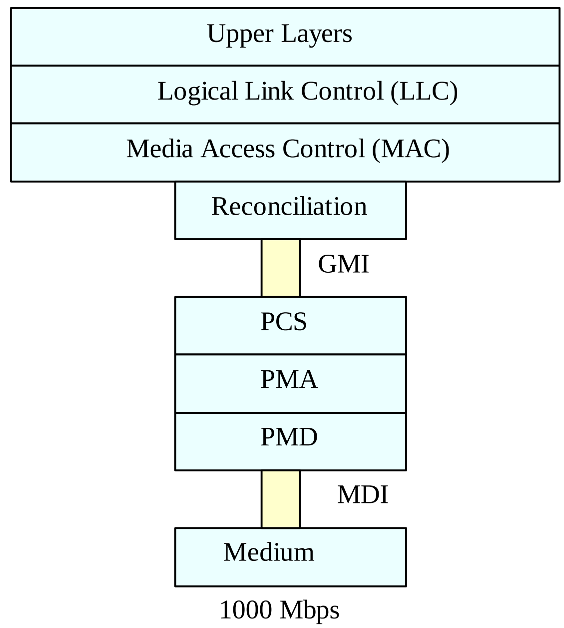

Figure 6: Gigabit Ethernet Architecture

3.5.1.1 Physical Layer

The Gigabit Ethernet physical layer is intended to support the following four media types.

a. Long-wave (LW) laser over single-mode and multimode fiber (1,000Base-LX);
b. Short-wave (SW) laser over multimode fibre (1,000Base-SX);
c. 150-Ohm balanced shielded copper cable (1,000Base-CX);
d. Unshielded Twisted Pair (UTP) category 5 cable (1,000Base-T).

Currently, three transmission media, 1,000Base-LX, 1,000Base-SX, and 1,000base-CX, have been specified in the IEEE 802.3z for Gigabit Ethernet implementations. The IEEE 802.3ab committee is working on the specifications for the 1,000Base-T transmission medium. That specification is expected to be available in 1999. The following table provides the supported distances over these media types.

<table>
<tr><th colspan="2">Media Types</th><th>Distance</th></tr>
<tr><td rowspan="3">1,000Base-LX</td><td>Single Mode Fiber</td><td>3000 m</td></tr>
<tr><td>50u Multimode Fiber</td><td>550 m</td></tr>
<tr><td>62.5u Multimode Fiber</td><td>440 m</td></tr>
<tr><td rowspan="2">1,000Base-SX</td><td>50u Multimode Fiber</td><td>550 m</td></tr>
<tr><td>62.5u Multimode Fiber</td><td>260 m</td></tr>
<tr><td colspan="2">1,000Base-CX</td><td>25 m</td></tr>
<tr><td colspan="2">1,000Base-T</td><td>100 m</td></tr>
</table>

Table 6: Media Types and Distances

The physical layer is divided into three sub-layers: Physical Coding Sub-layer (PCS), Physical Medium Attachment (PMA), and Physical Medium Dependent (PMD)

3.5.1.1.1 Physical Coding Sub-layer

This sub-layer provides a uniform interface to the reconciliation sub-layer for all physical media. It employs the 8B/10B encoding scheme used in fibre channel. The PCS also manages the auto-negotiation process by which the network interface communicates with the network to determine the network speed (10, 100, or 1,000 Mb/s) and the mode of operation (half-duplex or full-duplex).

3.5.1.1.2 Physical Medium Attachment sub-layer

The physical medium attachment sub-layer is identical to the PMA for fibre channel. It provides a medium-independent means for the PCS to support various physical media. The PMA sub-layer uses a serializer/deserializer to support multiple encoding schemes and to allow the presentation of these schemes to the upper layers.

3.5.1.1.3 Physical Medium Dependent sub-layer

The PMD sub-layer maps the physical medium to the PCS. It defines the physical layer signalling and actual attachments, such as connectors, for various media.

3.5.1.1.4 Interfaces

As shown in figure 6, Gigabit Ethernet architecture specifies two types of interfaces: Gigabit Media Independent Interface (GMII) and Medium Dependent Interface (MDI)

- **GMII:** The GMII is the interface between the MAC layer and the physical layer. It is an extension of the Media Independent Interface (MII) used in Fast Ethernet. Various media types such as shielded and unshielded twisted pair cables, single mode and

Multimode fibre cables, can be connected using GMII. Additionally, it is capable of supporting both full-duplex and half-duplex modes at 10, 100, and 1,000 Mb/s,

- **MDI:** The MDI is the actual physical interface.

3.5.1.2 MAC Layer

The MAC layer of Gigabit Ethernet is similar to that of standard Ethernet and Fast Ethernet. It supports both full-duplex (today) and half-duplex (future) transmissions. In full-duplex mode, the Gigabit Ethernet MAC uses the IEEE 802.3x Full-Duplex specification. The use of full-duplex mode eliminates collisions on the medium; therefore, flow control is not required. However, a frame-based flow control method has been proposed in IEEE 802.3x as an option. The first Gigabit Ethernet products support only full-duplex transmission.

In half-duplex mode, the Gigabit Ethernet MAC supports the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) access method, standardized in the original IEEE 802.3 Ethernet specification. Implementations of CSMA/CD for the Gigabit Ethernet are the same as that of Ethernet and Fast Ethernet. They allow the creation of shared gigabit Ethernet via hubs or half-duplex point-to-point connections.

3.5.1.2.1 Carrier Extension

Gigabit Ethernet maintains the minimum and maximum frame sizes (64 to 1514 bytes) of Ethernet. This is achieved by using a function called "Carrier Extension" which is not available in Ethernet and Fast Ethernet. Figure 7 shows Ethernet frame format with carrier extension.

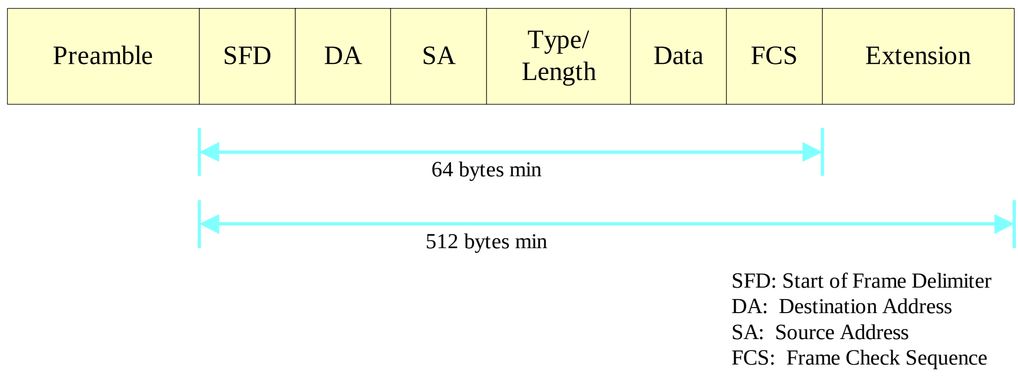

Figure 7: **Ethernet Frame Format with Carrier Extension**

If the frame is shorter than 512 bytes, then it is padded with extension bytes. Up to 448 padding bytes may be used for small frames. These bytes are special symbols, which cannot occur in the payload. This process is called carrier extension. The extension symbols are removed from the frame before the Frame Check Sequence (FCS) is checked by the receiver. Hence, the FCS is calculated only on the original frame.

Carrier extension is a simple function, but the use of padding bytes wastes bandwidth. This results in low throughput. To improve the throughput, a packet bursting function is used.

3.5.1.2.2 Packet Bursting

Packet bursting is a carrier extension plus a burst of packets. When a host has a number of packets to transmit, only the first packet is padded to the slot time if necessary using the carrier extension function. All subsequent packets are sent back-to-back with a minimum inter-packet gap until a burst timer (of 1500 bytes) expires. This function substantially improves the throughput.

3.5.2 Topologies

With Gigabit Ethernet, switched and shared topology can be used to connect various communicating devices. Switched topology supports longer distances and provides high throughput. Shared topology supports short distances, but provides low cost connections. Thus, the choice of topology depends on the network connection objectives. For example, the switched topology would be suitable for campus or building backbones, while the shared topology would be appropriate for server and desktop connections.

3.5.3 QoS on Gigabit Ethernet

Applications such as voice, video, and multimedia over LANs and WANs request specific QoS from network connections. Current Gigabit Ethernet specifications do not address QoS support for such applications, although, IEEE 802.1p, IEEE 802.1Q, and IETF's Resource Reservation Protocol (RSVP) could be used to support QoS over Gigabit Ethernet. The IEEE 802.1p and IEEE 802.1Q provides a means for tagging packets with an indication of the priority or class of service desired for the packet. These tags allow applications to communicate the priority of packets to internetworking devices. The RSVP is used to reserve resources and its support can be achieved by mapping RSVP sessions into 802.1p service classes. However, a major drawback of using RSVP is that each network component in the chain between a client and a server must support RSVP for delivering the requested QoS.

3.5.4 Gigabit Ethernet Standardization

IEEE 802.3 working group and Gigabit Ethernet Alliance, an open forum whose purpose is to promote industry cooperation in the development of Gigabit Ethernet, are working on the Gigabit Ethernet standardization. The IEEE approved the Gigabit Ethernet Standard 802.3z in June 1998. The IEEE 802.3z standard is not a stand-alone specification; it specifies changes and additions to the 802.3 standard for 1,000 Mb/s operation. Currently, the IEEE P802.3ab Task Force is working on the 802.3ab standard to define a Gigabit Ethernet physical layer for operation of over Category 5 UTP cable. Additional information regarding IEEE 802 standards can be found in Appendix F of this report.

3.6 ATM

ATM is emerging as a next-generation technology for multi-media communications. It enables simultaneous transmission of voice, video, and data over the same transmission medium. To facilitate the integration of voice, video, and data traffic, it uses the following five service categories.

1. **Constant Bit Rate (CBR):** This service category is used for connections that request a static amount of bandwidth for the connection lifetime. It is designed to support real time applications such as voice, video, and circuit emulation.

2. **Real-time Variable bit Rate (rt-VBR):** The rt-VBR category is intended for real time applications such as voice and video.

3. **Non-real time Variable Bit rate (nrt-VBR):** This category is used for non-real time applications having bursty traffic characteristics. These applications are less sensitive to delay.

4. **Unspecified Bit Rate (UBR):** The UBR service category is specified for non-real time applications not requiring tightly constrained delay and delay variations. This is a best effort service and does not provide any traffic related service guarantees. Example applications are file transfer and E-mail.

5. **Available Bit Rate (ABR):** This category is intended for sources having the ability to increase or decrease their information rate if the network requires them to do so. The bandwidth allocated by the network may vary during the lifetime of a connection.

ATM is a connection-oriented technology and provides Permanent Virtual Connections (PVCs) or Switched Virtual Connections (SVCs) to transfer the information between a source and a destination. The information transfer is accomplished using fixed-size (53-byte) packets called cells. Each cell consists of a 5-byte header field and a 48-byte payload field. The header field carries control information while the payload field carries user information.

3.6.1 ATM Protocol Architecture

As shown in figure 8, the ATM protocol is specified using various layers. These layers are the Physical layer, the ATM Layer, and the ATM Adaptation Layer (AAL).

ATM Adaptation Layer
ATM Layer
Physical Layer

Figure 8: ATM Protocol Architecture

3.6.1.1 Physical Layer

The physical layer provides for transmission of ATM cells over a physical medium that connects two ATM devices. It is divided into two sub-layers: the Physical Media Dependent (PMD) sub-layer and the Transmission Convergence (TC) sub-layer. The TC sub-layer transforms the flow of cells into a steady flow of bits for transmission over the physical medium. The PMD sub-layer provides for the actual transmission of the bits in the ATM cells. Table 7 shows the physical interfaces specified by the ATM Forum, the International Telecommunications Union – Telecommunications Standardization Sector (ITU-T) and the ANSI.

1. DS1 at 1.544 Mb/s	6. UTP at 25.6 Mb/s, 155.52 Mb/s	11. STS-48c/OC48 at 2.488 Gb/s
2. E1 at 2.048 Mb/s	7. STP at 155 Mb/s	12. STS-192cm/OC48n at 2.488n Gb/s
3. DS2 at 6.312 Mb/s	8. STS-1/OC1 at 51.84 Mb/s	13. POF at 155 Mb/s
4. DS3 at 44.736 Mb/s	9. STS-3/OC3 at 155.52 Mb/s	14. MMF Lasers at 155 Mb/s
5. E3 at 34.368 Mb/s	10. STS-12c/OC12 at 622.08 Mb/s	

Table 7: Physical Interfaces

3.6.1.2 ATM Layer

The ATM layer performs all the functions relating to the routing and multiplexing of ATM cells over virtual connections. These connections are either pre-established (PVCs) or are set up on demand (SVCs) using dynamic signalling. Its primary function is to assign a header to each 48-byte segment received from the AAL and remove that header before sending the ATM cells to the appropriate AAL. It also provides encoding, physical layer services, ATM layer user services, ATM layer management, and traffic and congestion control. To perform these functions, the ATM layer uses Virtual Path Connection identifiers (VPIs) and Virtual Channel Connection Identifiers (VCIs). The ATM Forum has specified valid VPIs and VCIs.

3.6.1.3 AAL

The AAL provides various service classes for the transport of information generated by higher layers. The AAL converts the submitted information into streams of 48-byte segments and transports them in the payload field of multiple ATM cells. Similarly, upon receiving the stream of ATM cells, it converts the 48-byte information field contained within each cell into the required form for delivery to the particular higher layer. Currently specified service classes are AAL 1, AAL 2, AAL 3, AAL 4, and AAL 5.

1. **AAL 1:** This class provides a connection-oriented service and supports the CBR service category.

2. **AAL 2:** This class provides a connection-oriented service. It is mainly used to transfer the rt-VBR service category.

3. **AAL 3:** This class provides a connection-oriented service and supports the nrt-VBR service category.

4. **AAL 4**: This service class offers a connectionless service. It is used for nrt-VBR service category.

5. **AAL 5:** This class provides a connection-oriented service and supports CBR and rt-VBR service categories.

3.6.2 Traffic and Congestion Control

Once a connection is established, ATM networks guarantee to provide the negotiated QoS. To provide such guarantee, the following controls have been specified.

a. **Connection Admission Control (CAC):** The set of actions taken by the network during the call set-up phase in order to determine whether a connection request can be accepted or should be rejected.

b. **Feedback controls:** The sequence of actions performed by the network and end-systems to regulate the traffic submitted on ATM connections according to the state of the network.

c. **Usage Parameter Control (UPC):** The series of actions taken by the network to monitor and control traffic, in terms of traffic offered and validity of the connection, at the end-system access level. Its main purpose is to protect network resources from malicious as well as unintentional misbehavior, which can affect the QoS of established connections.

d. **Cell Loss Priority (CLP) control:** Enables an end-system to generate traffic flows of cells with cell loss priority marking. As a result, a congested network may selectively discard cells marked with a low priority to protect the QoS objectives of cells with a high priority.

e. **Traffic Shaping:** It is employed to achieve a desired modification of the traffic characteristics.

f. **Network Resource Management (NRM):** Allows for the allocation of network resources in order to separate traffic flows according to service characteristics, to maintain network performance, and to optimize resource utilization. It is mainly concerned with the management of virtual paths (VPs) because they serve as a useful tool for resource management.

g. **Frame discard:** Provides a congested network with a capability to discard cells at the frame (the ATM Adaptation Layer (AAL) protocol data unit) level rather than at the cell level. It may help avoid congestion and can increase throughput.

h. **ABR flow control:** Provides the capability to adaptively share the available bandwidth among participating users.

3.6.3 ATM Network Security

ATM security specifications are still under development. The ATM Forum released the ATM Security Framework, Version 1.0, in February 1998. It basically addresses generic threats, generic security requirements, and functional security requirements for security services to be provided within ATM networks. The framework does not address any specific security aspects of ATM users and applications. Another security specification called ATM Security Specification, Version 1.0 is under final ballot.

3.6.4 ATM Network Management

The ATM network management framework covers three areas: interface management, layer management, and global network management.

a. **Interface management:** Interface management deals with the exchange of information at the interface level. It is used primarily for configuration and obtaining status information of ATM interfaces. Currently defined interfaces by the ATM Forum include: (1) User-to-Network Interface (UNI), (2) Broadband Intercarrier Interface (B-ICI), (3) LAN Emulation Interface (LUNI), and (4) Data Exchange Interface (DXI).

b. **Layer management:** Layer management allows testing for continuity and loopback at a segment or an end-to-end virtual circuit level. It is mainly used for end-to-end circuit management and allows checking a user circuit through the network.

c. **Global network management:** Global network management deals with the configuration of an ATM network consisting of one or more switches and the monitoring and control of ATM devices in the network. The ATM Forum network management group has developed an ATM network management reference model to address global network management. The reference model defines the following five management interfaces, M1 through M5.

M1: supports various functions for the management of ATM devices.
M2: defines functions for the management of private ATM networks.
M3: manages interactions between private and public ATM networks.
M4: allows the management of public ATM networks.
M5: manages communications between two public ATM networks.

3.6.5 ATM Standardization

Several organizations such as ATM Forum, ITU-T and ANSI have been actively working to standardize ATM specifications. Numerous specifications have been developed by these standards organizations to provide interoperability, reliability and security. A partial list of the standards pertaining to ATM can be found in Appendix E of this report.

4 TCP/IP on Gigabit Networks

Transmission Control Protocol (TCP) and Internet Protocol (IP) are the most widely implemented protocols. Over the last few years extensive research has been done to determine the suitability of these protocols over high-speed networks. These efforts have resulted in the development of many Request For Comment (RFC) documents such as RFC 1323, RFC 2018, RFC 2001, and RFC 1981. They describe the issues associated with the use of these protocols over the high-speed networks and their possible solutions, often called the extensions to basic TCP/IP. Experimental results have validated that the TCP and the IP can be used over high-speed networks if hosts support and utilize the extensions to basic TCP and IP.

For example, the Department of Electrical Engineering and Computer Science at the University of Kansas, Lawrence, KS, conducted several experiments and simulations involving TCP/IP as a part of the Advanced Communications Technology Satellite (ACTS) ATM Internetworking (AAI) project. The experiments were carried out on OC-3 and OC-12 ATM satellite links. The results indicated that the extensions to basic TCP could be used on high-speed networks to achieve high throughput. Gigabit implementations of TCP/IP were also used in gigabit testbeds such as CASA. At present, there are some TCP implementations that run at gigabit per second rates proving that the speed is not an issue for TCP on gigabit networks. However, a new issue for TCP is the support of real-time applications.

Presently, TCP is not well suited for transferring data with real-time characteristics such as interactive, audio, and video. There are several reasons for this. First, a reliable transmission is inappropriate for delay-sensitive data such as real-time audio and video. For example, in the case of an error, by the time the sender discovers it and retransmits an erroneous packet, at least one round trip elapses. During this time the receiver has to wait for the retransmission incurring a noticeable gap in the transmission. Second, TCP does not support multicast. Third, the TCP congestion control mechanisms decrease the congestion window when packet losses are detected. Audio and video, on the other hand, have natural rates that cannot be suddenly decreased without starving the receiver.

The IETF has developed a Real-time Transport Protocol (RTP) that provides end-to-end delivery services for data with real-time characteristics. The RTP is primarily designed to run on top of User Datagram Protocol (UDP) to make use of its multiplexing and checksum services. However, it may be used with other suitable underlying network or transport protocols such as IPv6 and ATM. The RTP supports data transfer to multiple destinations using multicast distribution if provided by the underlying network. It does not provide any mechanism to ensure timely delivery or QoS service guarantees. The RTP relies on lower-layer services to do so. Additionally, it does not guarantee delivery or prevent out-of-order delivery, nor does it assume that the underlying network is reliable and delivers packets in sequence. The data transport is augmented by a control protocol called RTP control protocol (RTCP). The RTCP allows monitoring of the data delivery and provides minimal control and identification functionality.

5 NS/EP Gigabit Applications and Enabling Technologies

5.1 Gigabit Applications

Gigabit networks, of the types described in the previous sections of this report, are likely candidates to support NS/EP requirements over the next several years. For years computer networking applications have been built over technologies such as 10Mb/s Ethernet, T1 or T3 dedicated digital circuits, and of course the existing digital telephone network that has been evolving under ISDN specifications. Over the next few years the underlying communications infrastructure is expected to undergo widespread changes. Ethernet technology has already evolved to 100Mb/s speeds and both IEEE standards and commercial products are already available that provide 1,000Mb/s speed Ethernet. In a similar manner ATM systems are providing gigabit bandwidth capabilities. These changes hold tremendous potential for supporting new applications that can significantly improve the capabilities for managing NS/EP events.

5.2 Application Affinity Groups

As originally envisioned by the Federal Internetworking Requirements Panel (FIRP) affinity groups can be established where a number of application interests need to be coordinated. Applications that are important to only one agency such as intelligence gathering for the CIA or aerospace engineering for NASA, can be handled within the respective agency and do not require affinity group coordination.

For the purposes of this report we further subdivide an affinity group to have at least one of two dimensions: disciplinary or technology. First are the disciplinary affinity groups. These are a collection of end-user organizations that share common interests such as health care, education, or environment. The agencies comprising the NCS, for example, formed a distinctive NS/EP affinity group. These groups collaborate because they recognize that their applications have a great deal in common; and that by collaboration each will realize their own goals more efficiently.

Second are the technology affinity groups. These groups have as their mission to coordinate and develop the technologies or tools that must reside either within the network, within the applications, or both. For example, many applications will require the ability to facilitate collaboration over new gigabit network infrastructures. Examples of both dimensions of affinity groups are depicted in Table 8 below.

Disciplinary Affinity Groups	Applications Technology Affinity Groups
Health Care	Collaborative Technologies
Environment	Distributed Computing
Education	Digital Libraries
Manufacturing	Remote Operations
Crisis Management	Security and Privacy

Table 8: Affinity Groups

In order to take full advantage of the evolving gigabit networking infrastructure the NCS needs to also be proactive in influencing the development of a number of application technologies and/or the respective affinity groups associated with the technology. Especially challenging for the NCS is that during any given NS/EP event it may be necessary to dynamically construct the appropriate matrix of affinity and technology groupings. For example, a given NS/EP event may/may not require the use of digital libraries, or distributed computing technologies. The NCS will also be challenged to identify and concentrate scarce resources on application technology areas that offer the greatest promise of providing the highest NS/EP event management dividends. To assist with that endeavor this section of the report highlights a few key areas in which substantial progress can be made toward developing robust NS/EP applications that take advantage of a gigabit network infrastructure. These include:

1) distributed object computing;
2) telecollaboration;
3) visualization for data analysis;
4) data mining;
5) computer simulation;
6) telemedicine.

5.3 Distributed Object Computing

The usefulness of a high speed gigabit infrastructure can only be realized when managers and users of NS/EP applications can quickly and easily locate and access data and computing resources in the network that are mission essential. In the future those resources may involve, for example, public or private electronic libraries, data conversion services, and other application specific services that may not have been designed exclusively with an NS/EP objective. Since thousands of services, useful during NS/EP conditions, built by hundreds of organizations will exist on the network it is important to facilitate ways of building applications by using software components built by different organizations.

NS/EP distributed object computing research should focus on design patterns, implementation, and experimental analysis of object-oriented techniques that facilitate the development of high-performance, real-time distributed object computing frameworks on parallel processing platforms running over high-speed gigabit networks.

To enable users to dynamically assemble the components they require to build a system capable of handling a particular NS/EP event will require further evolution of distributed computing technologies. One emerging solution to these envisioned dynamic requirements can be seen through the growing popularity of the use of standardized distributed objects. Properly designed object based systems have the unique advantage that changes in one area of the system do not affect other areas. A properly designed set of NS/EP relevant object attributes would enable us to link resources in a modular way to manage complexity and yet allow great extensibility. An emerging industry standard from the Object Management Group (OMG) known as the Common Object Request Broker Architecture (CORBA) can today be used as a key component in building gigabit networked distributed object systems.

The OMG has a Domain Special Interest Group (DSIG) called OMG C4I DSIG. C4I stands for Command, Control, Communications, Computers, and Intelligence and is equivalent to the UK term C2I CIS (Command, Control and Intelligence, Communications and Information Systems). Thus far, no specific CORBA support (e.g., object definitions) for national security or emergency preparedness conditions have been addressed by this or any other OMG SIG. This is one area where progress toward supporting NS/EP requirements could yield significant benefits for the NCS.

5.4 Telecollaboration

Telecollaboration as used within this report refers to a variety of emerging technologies that enable collaborative efforts involving graphics, documents, or other multimedia objects to be discussed and manipulated on a collective basis while simultaneously supporting teleconferencing. For the NCS this emerging technology promises to provide a general, media-rich user interface system that will facilitate enhanced information sharing and teamwork activities, which are so critical to achieving successful NS/EP event management results. In its simplest form this capability should permit peer-to-peer collaboration in a common screen space. More advanced capabilities should also include one-to-many as well as many-to-many collaboration with video teleconferencing.

Today, a number of marketplace products operate using the Internet RFC 1459, Internet Relay Chat (IRC) protocol. A few implementations may also include the features described in the recent Internet Draft: "Extensions to the Internet Relay Chat Protocol (XIRC)," as well as proprietary techniques that integrate word processing and other document management and graphics manipulation products into a collaborative Internet "whiteboard" solution. These techniques can most often be combined with Internet and/or non-Internet telephony to create and enhanced level of information exchange.

Within the context of an NS/EP event management applications system a secondary yet important aspect of the technology should be the ability to develop generalized, media-rich user interfaces that have the capabilities to easily serve the requirements of NS/EP application clients. The Defense Information Infrastructure (DII) Common Operating Environment (COE) provides an excellent example of how standardization on a user rich interface can serve to facilitate common training, applications, and other benefits associated with establishing an executive user desktop that has a set of convenience capabilities for all application areas. This type of enhanced and robust applications solution may be well suited to provide a foundation for extension that could more directly incorporate and satisfy NS/EP event management requirements.

5.5 Visualization For Data Analysis

Within the context of a gigabit networking infrastructure, visualization applications should evolve to play a much more important role in assisting with the management of NS/EP events. Visualization is crucial not only for perusal of extremely large scientific data sets but it also impacts distributed computing, multimedia, data management, and networking efforts as well. As a fundamental tool in the analysis of data, visualization should evolve to become a core tool of an NS/EP event management system. In addition to providing the necessary analytical tools, collaborative visualization tools will need to be developed that take advantage of a gigabit infrastructure to leverage it into a more powerful yet usable commodity.

Visualization spans the technology thrust areas and the application areas. Scientific visualization can build upon the efforts in distributed computing by working closely with the underlying foundations provided by the distributed object environment. Data management requires strong visualization both in terms of data perusal as well as in effective utilization of remote data. Visualization must be integrated with multimedia. This facilitates both object animation as well as active geometric objects, which can be employed to effectively communicate underlying trends within the data.

Much of an NS/EP event data can be presented graphically. Since it is difficult to make many assumptions about the computing power available at an NS/EP event site, a more general model should assume that a rendering is done by a server and a local display is updated only as needed.

With the coming gigabit networking infrastructure and high capacity computers visualization techniques can be applied with respect to weather systems (e.g., including interactive ocean modeling), visualization of near real time wildfire simulations, and earthquake or volcanic eruption visualizations including mantle convection simulations. The use of massively parallel visualization techniques for producing fast stereoscopic images with ray-traced volume rendering and other supercomputer services should be part of the vision for future NS/EP event management scenarios.

5.6 Data Mining

Data mining (also known as Knowledge Discovery in Databases - KDD) can be defined as the nontrivial extraction of implicit, previously unknown, and potentially useful information from data. Data mining uses machine learning, statistical, and visualization techniques to discovery and present knowledge in a form that is easily comprehensible to humans.

With the evolution of both a new gigabit networking infrastructure and the technologies capable of recording, analyzing, and displaying high-resolution digital data for a bewildering array of applications, an information explosion is expected to build rapidly. In a manner similar to the Federal interstate highway system, which led to the explosive growth of automobile use, a gigabit networking infrastructure will inevitably lead to explosive growth in the availability of and demand for electronic information. The viability of such a system will hinge on our ability to continue accommodating Internet levels of explosive growth. Managing the coming gigabit infrastructure growth will additionally be complicated by the multidimensionality of the information. In the next decade, users will be accessing traditional text databases, digital sound, images, video, and numerical archives as well as evolving paradigms for structured data. For NS/EP purposes especially many applications will involve heterogeneous information, both structured documents and raw numerical data, and essentially all users will be simultaneously submitting their requests and reports from remote locations.

The goal that will become common to all NS/EP event users will be to maximize the amount of useful information extracted while minimizing the loss of productivity resulting from wasted time and resources, prohibitive access or transmission charges, and the retrieval of meaningless, incorrect, or corrupted data. To meet these challenges NCS must develop the capabilities to represent and store large information data sets efficiently, to extract and transmit material across a distributed multi-user network, and to analyze, interpret and display such information effectively and meaningfully.

Many large image databases already exist from a variety of imaging applications, including image databases associated with earth orbital satellite systems. Other potential image databases include collections of news photographs, ID photos, fingerprints, and medical images. The ability to efficiently browse these enormous databases will greatly enhance their value.

The storage, retrieval, and manipulation of image data today faces many unsolved problems. For example, the issues related to browsing massive databases of digital imagery includes several related efforts such as the automated segmentation of digital imagery, search and retrieval of digital imagery using image content, and the development of data compression strategies that facilitate progressive coding/transmission.

5.6.1 Data Mining Standardization Efforts

The CRoss Industry Standard Process for Data Mining (CRISP-DM) has as it's general objectives to cater for the needs of data warehouse users by defining a cross industry data mining process and providing tool support, allowing for cheaper, faster, and more reliable data mining, and widespread adoption of the CRISP-DM process model. Additional objectives include: ensure the quality of knowledge discovery project results, reduce skills required for knowledge discovery, capture experience for reuse, general purpose (i.e., widely stable across varying applications) robust (i.e., insensitive to changes in the environment), tool and technique independent solutions.

5.7 Computer Simulation

The advent of distributed computing environments has provided the scientific community with tools to tackle a cornucopia of illustrious challenge problems addressing national concerns as diverse as atmospheric modeling, energy production, and pollution control. The computer simulations built to address these issues involve teams of researchers working with a mixture of physical models implemented on the latest cutting-edge distributed computing platforms. Although today's grand challenge scientists have access to the world's fastest computers, some of the programming paradigms and languages utilized in these systems literally date back to the days of paper tape and punch cards. The hardware and software must evolve together to capture the complexity of today's emerging NS/EP simulation challenges. Object-Oriented programming languages are promising approaches to solving these application development issues.

If we consider the classical structured programming techniques on a distributed computer with a parallel language extension of C or Fortran, the time required for debugging a moderate sized NS/EP related scientific application is roughly equivalent to the half-life of the latest parallel supercomputer. The issues regarding code portability and reuse grow with the size of the application. However, with high-level objects, such as C++ class libraries, they allow the scientist to express the required abstractions, metrics, and equations directly into the language. High level objects address the issue of code portability between parallel supercomputers by utilizing underlying architecture dependent run-time linked class libraries which encapsulate from the user the node-CPU layout, interprocessor network topology, memory hierarchy, and I/O protocols. These high and low level objects working together allow for code transportation to a variety of parallel architectures with virtually no changes to the original source code. These underlying architecture dependent objects can also mimic parallel computations on serial computers with message-passing modeled as memory read/writes. Thus, a developer who has restricted access to parallel computers is able to perform the necessary development on a workstation and then, with no change to the source code, port the application to a parallel supercomputer.

Today, research teams are focused on developing and combining parallel C++ class libraries to tackle problems in diverse areas, such as numerical research, xerography, and medical imaging. One goal is to build hierarchical classes upon existent classes and, where necessary, create new classes, which together will act as a multi-leveled simulation environment focused on the rapid solution of scientific, engineering, and industrial problems. These solutions can be of direct application and benefit for use during NS/EP events.

Predictive modeling often requires time critical computing where results from a simulation must be produced and interpreted quickly enough to be of use in an ongoing crisis. The time critical nature of these problems requires visualization tools that are tightly integrated with the simulation processing. Visualization tools must be the primary interface to a running simulation for such time critical problems, allowing time critical interpretation of simulation results and allowing NS/EP event management "what if" scenarios to be explored rapidly.

NS/EP scale events will typically be expected to fall beyond the scope and scale utilized by the commercial marketplace. A top down NS/EP focused research and development program for both the telecommunications infrastructure and the applications environment will be required to meet future requirements.

5.8 Telemedicine

As an application layer service telemedicine can be generally defined as the use of telecommunications technologies coupled with medical expertise to facilitate remote health care services. Clearly, few applications hold the potential for providing substantial service delivery improvements to NS/EP environments as telemedicine. Historically, radios and telephones have provided telecommunications capabilities to NS/EP event locations. For the purposes of this paper, telemedicine is referred to as new computer and telecommunication medical technology services that facilitate the communication of medical information from one physical location to another. Today, it is being used in more diverse areas of health care, including pathology, surgery, physical therapy, psychiatry, and continuing medical education. A list of typical applications of telemedicine includes:

1. Electronic management and transport of patient information and records for diagnostic purposes;
2. Image compression for efficient storage and retrieval of image data;
3. Image processing for diagnostic purposes;
4. Electronic processing of health and medical claims;
5. Electronic inventory to support community health care organizations;
6. Teleconferencing for professional training, education, and consultations; and
7. Digital transmission of large 2-D and 3-D medical images.

Telemedicine reflects the convergence of technological advances in a number of fields, including medicine, telecommunications, computer engineering, infomatics, artificial intelligence, robotics, materials science, and perceptual psychology.

Fantastic new breakthroughs in teleradiology and telemedicine are already having a major impact on the way healthcare is being delivered today. For example, during a recent expedition to Mt. Everest the climbers received real-time physical check-ups from doctors in the United States. Today, x-rays of a patient can be examined by a physician hundreds of miles away, and soldiers at a remote location can be simultaneously examined and given treatment options by a team of medical specialists from around the world. In the future soldier's uniforms can be equipped with automated health monitoring data sensors. In November 1998 the Renaissance Cruise line announced a comprehensive wellness at sea program that includes an exclusive relationship with Johns Hopkins Medical Center for a telemedicine link that provides direct access to high-quality, state-of-the-art medical assistance 24 hours/day. The cornerstone of the program is telemedicine technology which provides shipboard physicians access to the Johns Hopkins emergency medicine specialists 24 hours a day, seven days a week via a high-speed satellite video link. The Johns Hopkins physicians can also provide consultations to shipboard doctors about diagnosis, treatment, transfer, and follow-up care.

All of the aforementioned developments strongly suggest that telemedicine should be viewed as a high priority application with direct relevance to the management of NS/EP events.

5.8.1 Telemedicine Standards Activities

Clearly, standardization in telemedicine can have a direct benefit to the efforts of the NCS with respect to NS/EP casualty management.

As part of the ANSI Healthcare Informatics Standards Board (HISB) the Health Level Seven (HL7) Standard, an ANSI approved American National Standard for electronic data exchange in health care, enables disparate computer applications to exchange key sets of clinical and administrative information. Comprised of standard formats, which specify the implementation of interfaces between different computer applications, the HL7 protocols provide the flexibility needed to allow compatibility for specialized data sets that have facility-specific needs.

The origin of "Level Seven" was the Open Systems Interconnection (OSI) Reference Model defined by ISO 7498, which is the application layer. The application layer addresses definitions of the data to be exchanged, the timing of the interchange, and the communication of certain errors to the application. The seventh level supports such functions as security checks, participant identification, availability checks, exchange mechanism negotiations and, most importantly, data exchange structuring.

The scope of the ANSI HISB is very broad and includes:

1. The coordination of standards for healthcare models and electronic healthcare records;
2. The interchange of healthcare data, images, sounds and signals within and between organizations/practices;

3. Healthcare codes and terminology;
4. The communication with diagnostic instruments and healthcare devices;
5. The representation and communication of healthcare protocols, knowledge, and statistical databases, privacy, confidentiality and security of healthcare information;
6. Identifiers for healthcare providers, patients, entities, etc.

All of these aspects have wide applicability during NS/EP events.

5.8.1.1 HL7 Standard Version 2.3

The structure of the Version 2.3 publication is organized into 12 chapters. Each chapter of the standard has its own corresponding technical committee. In addition, technical committees in the areas of syntax, data warehousing, education, implementation, modeling and methodology, and vocabulary are contributing to the development of HL7. Key aspects of this version of the standard include:

1. Control/Query-Message;
2. Definitions and Interchange Protocols;
3. Patient Administration (admit, discharge, transfer, and demographics);
4. Order Entry for Clinical Services and Observations;
5. Pharmacy, Dietary, and Supplies;
6. Financial Management (patient accounting and charges);
7. Observation Report Messages;
8. Health Care Application Master Files;
9. Medical Records/Information Document Management Services and Resources;
10. Scheduling-Appointments, Patient Referrals, Care Referral Messages, and Patient Care-Problem-Oriented Records.

5.8.1.2 HL7 Version 3.0 Specifications

Slated for first ballot in the fall of 1999 with publication of the first chapter expected in January 2000, Version 3.0 of the HL7 standard is being constructed using rigorous analytic and message building techniques and incorporating additional trigger events and message formats with very little optionality, HL7's primary goal for Version 3.0 is to offer a standard that is definite and testable.

Version 3.0 will use an object-oriented development methodology and a Reference Information Model (RIM) to create messages. The RIM is an essential part of the HL7 Version 3.0 development methodology, as it provides an explicit representation of the semantic and lexical connections that exist between the information carried in the fields of HL7 messages.

A draft RIM, incorporating 126 classes and 861 attributes, has been completed and is currently being analyzed by HL7 technical committees. Each technical committee holds the responsibility for the appropriate subsets of the draft RIM based upon the domain of its message development. The HL7 Modeling and Methodology (M&M) Technical Committee oversees this process.

Additionally, "Health Level Seven Messaging over Component Technologies" was recently completed, moving HL7 a significant step ahead in the development of Version 3.0. Developed by the HL7 Object Brokering Technologies Special Interest Group (SIGOBT), the document recommends the use of object brokered (or distributed component) technologies within the scope of the HL7 2.x messaging standard; targeting both CORBA ORB-IIOP and ActiveX DCOM technologies.

As part of the Version 3.0 effort, the Vocabulary Technical Committee is developing methods that will allow HL7 messages to draw upon codes and vocabularies from a variety of sources, and assure that the systems sending and receiving these messages have an unambiguous understanding of the code sources and code value domains they are using.

HL7 was originally founded to develop standards for the electronic interchange of clinical, financial and administrative information among independent health care oriented computer systems; e.g., hospital information systems, clinical laboratory systems, enterprise systems and pharmacy systems. In the last four years, HL7 has tripled to over 1,500 hospital, professional society, health care industry, and individual members including almost all of the major health care systems consultants and vendors. The HL7 standard is supported by most system vendors and used in the majority of large U.S. hospitals today. It is also used in Australia, Austria, Germany, Holland, Israel, Japan, New Zealand and the United Kingdom.

The ANSI has designated HL7 as an accredited standards developer and HL7 has already released its fourth version of the standard. This consensus standard was balloted under ANSI rules

5.8.1.3 Other Telemedicine Related Standards Activities

In addition to U.S. efforts (e.g., HL7, ANSI) a number of organizations around the world are also working toward standardization of telemedicine. There is active work underway by an ISO Member Bodies Task Force on Healthcare Terminology, the European Health Telematics Observatory (EHTO), the International Medical Informatics Association (IMIA), the European Committee for Standardization Technical Committee for Health Informatics (CEN/TC 251), and others. Although the U.S. is in an enviable position with respect to setting telemedicine standards for the world a number of non-US organizations should be monitored for the potential they hold with respect to influencing the direction of technology. For example, the International Medical Informatics Association (IMIA) is a working group formed in Canada in association with other European telemedicine standardization efforts.

The IMIA regionally is divided into four different sections: IMIA-LAC (Latin and Central America), EFMI (Europe) and APAMI (Asian and Pacific Region). The African region (HELINA) is in developing stage. IMIA Working Group (WG) 16 was formed at a general assembly meeting in Vancouver, Canada. The scope of the IMIA standardization in Medical Informatics is as broad as Medical Informatics itself. As such IMIA WG 16, the working group on Standards, has a horizontal range of coverage, comparable to the group on Education and Training in Medical Informatics, that has many links with other IMIA Working Groups.

Being basically both science driven and international, the objectives of IMIA WG 16 are to facilitate the exchange of information between the different standardization bodies of different continents, to ensure broad dissemination, and to create awareness in connection with the standards. One goal of the IMIA WG 16 is to serve as a catalyst for the creation of an ISO technical committee.

5.8.2 Application of Telemedicine to NS/EP Events

Figure 9 is courtesy of the Lawrence Livermore National Labs, Sunrise Project. It is provided to illustrate how an example graphical patient record interface might appear on an NS/EP deployed workstation screen.

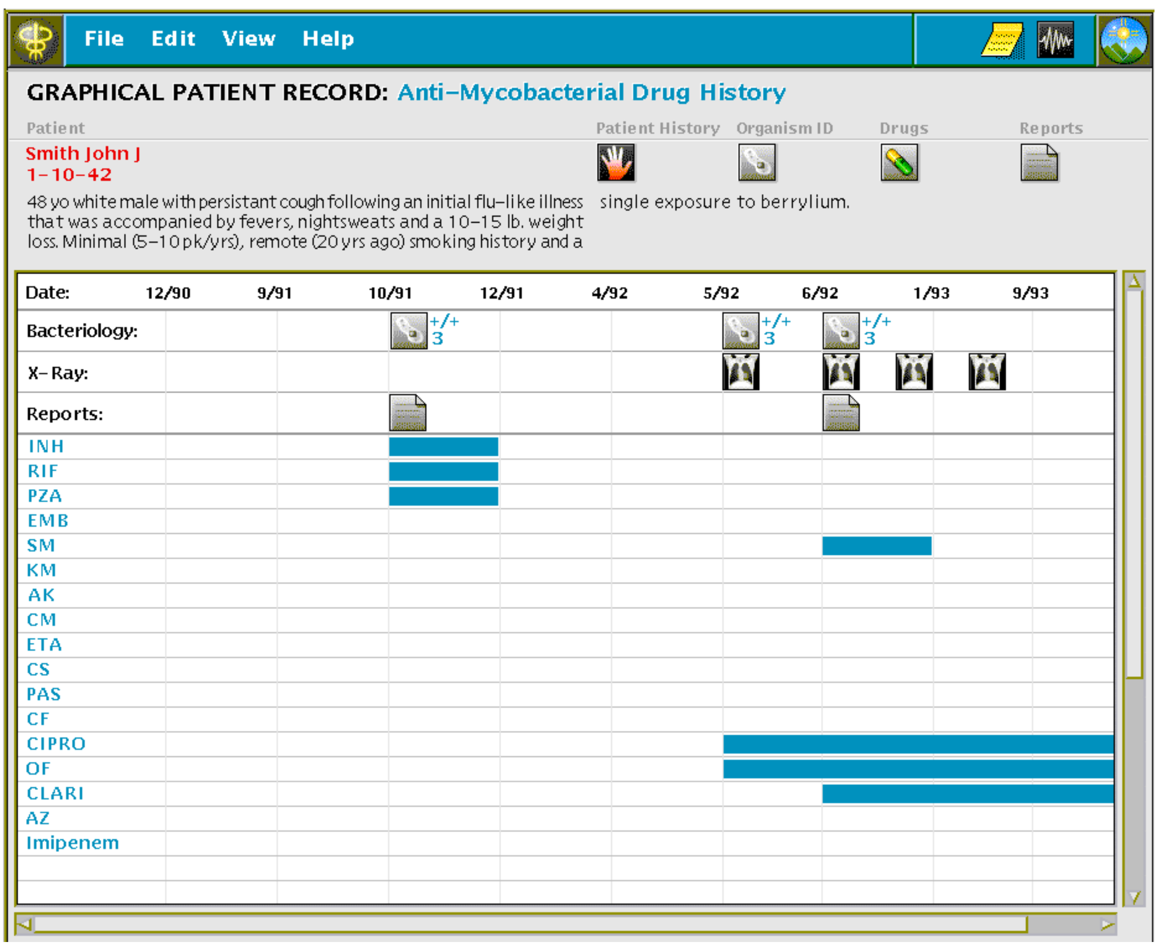

Figure 9: Graphical Patient Record Interface

Subsequent to the application creating the opening screen on an NS/EP deployed workstation screen, databases would be successfully opened and the ability to select patients from a list or via search engine could be presented. Following selection of a patient, and entering information appropriate to the medical emergency condition the NS/EP supporting network connects with appropriate databases from other institutions that provide on-site data regarding the condition and treatment. The data display should support an array of information presented as textual data, image data, audio, and other forms of multimedia presentation. Data regarding prior medication conditions could be transmitted directly to the site. For example, patient history, graphical patient history, graphical patient records, an image history showing Computed Axial Tomography (CAT) scan images as a function of time, image thumbnails showing all CAT history as a thumbnail, the ability to request enlargements of thumbnails, and navigation of visual databases.

5.8.3 Future Application Bandwidth Requirements

With high bandwidth technology applications such as telemedicine it is easy to anticipate that gigabit networking speeds will not be sufficient to satisfy requirements for very long. Already high-speed networking authorities have stated that gigabit rates aren't nearly high enough. Internet Service Providers (ISPs) are reporting growth rates that are anywhere between 15 percent and 30 percent per month. If that growth is extended over a 12 month period it means that annual growth rates are 500 percent to 1,800 percent.

Although it is unlikely that these growth rates will be sustained indefinitely, improvements in such access technologies as xDSL (digital subscriber line) suggest that at least for the near term we are in a continuing trend. To support this type of growth carriers will have to build much larger capacity trunking capabilities. Indeed to support the present level of growth, petabit-per-second capacity will be needed by 2007 if growth continues at just 15 percent and if it continues at 30 percent petabit capacity will be needed by 2002.

To put the capacity numbers into perspective, consider that a gigabit is equivalent to 10^9 bits/s and a terabit to 10^{12} bits/s. A petabit is equal to 10^{15} bits/s of capacity. Already systems that are operating on the leading edges of research are projecting much larger figures, like exabit (10^{18} bits), zettabit (10^{21}), and yottabit (10^{24}).

6 Gigabit networking issues

Clearly, gigabit networking significantly improves the speed at which information is transmitted between workstations, mainframes, supercomputers, storage devices, and other communicating devices. Besides the speed, there are several other features or issues that also play an important role in determining their suitability for a particular application or environment such as NS/EP. This section addresses some of these features or issues.

6.1 Interoperability

Gigabit networks have become a reality due to the existing and emerging gigabit networking technologies. Each of these technologies has advantages and shortcomings. As a result, it is unlikely that a single technology will be the best for all applications. However, the maximum benefit can be achieved from their interoperability. Therefore, these technologies should be able to interoperate with each other. Not only that, they should also be able to interoperate with existing networking technologies.

Various standards organizations are working on the interoperability of gigabit networking technologies. Gigabit Ethernet is designed to be interoperable with existing Ethernet and Fast Ethernet implementations. The ANSI has developed a standard, ANSI X3.299, for encapsulation of HIPPI-PH packets for transfer over ATM. The ISO is currently processing ANSI standard ANSI X3.254 for sending HIPPI-FP packets over Fibre Channel Physical Layer, as a Draft International Standard (DIS). Today, various vendors are offering solutions that provide interoperability among gigabit networking systems, but numerous interoperability problems in multivendor environments have been encountered.

6.2 Priority

Dissemination of NS/EP information on a priority basis is critical to managing NS/EP conditions. Therefore, priority treatment of NS/EP traffic has been identified as one of the core requirements. The use of and existence of priority may be different for different types of traffic. For example, network control traffic should be given the highest priority to ensure that the network operates efficiently and reliably. NS/EP traffic should be given higher priority over non-NS/EP traffic. Moreover, within the NS/EP traffic there is a need to support multiple levels of priority corresponding to the mission criticality or urgency level associated with the particular NS/EP information. This in turn requires the ability to recognize NS/EP data and to distinguish multiple levels of data urgency.

To be useful in an NS/EP environment, the gigabit networking technologies under considerations must provide for standardized handling of prioritized emergency data. Currently, gigabit networking technologies and standards address only traffic priority to support QoS required for various traffic types. They do not provide any mechanisms to support a prioritized handling service for emergency data.

6.3 Security

During NS/EP conditions, users have requirements to transmit a wide range of information types ranging from unclassified information, to sensitive, or in some instances classified information. Consequently, the availability of trusted and secure communications during these conditions is essential. Typically, this is achieved by implementing various security services such as authentication, data integrity, confidentiality, access control, and non-repudiation. For the NS/EP scenarios, the system may also be required to comply with the guidelines set by U. S. Government agencies for processing classified and/or Sensitive But Unclassified (SBU) information.

Prior to adoption for use in NS/EP environments, gigabit networking technologies must provide the means for secure and reliable communications. However, gigabit technologies are still emerging and their security support is not fully defined yet. For example, ATM security specifications are under development. The ATM Forum released the ATM Security Framework, Version 1.0, in February 1998, addressing generic threats, generic security requirements, and functional security requirements for various security services. Currently, there are some proprietary products available for ATM security, but they were not designed for gigabit speeds. Hence, their suitability at gigabit speeds is for further research. Fiber channel, Gigabit Ethernet, and HIPPI, on the other hand, do provide limited security capabilities.

6.4 QoS

Gigabit networking technologies should support QoS so that applications can request specific qualities or levels of service from the network. Among the gigabit networking technologies available today, only ATM provides mature support for QoS. Once a connection is established, ATM networks guarantee to provide the requested QoS. Gigabit Ethernet specifications do not have any built in support for guaranteed QoS. Currently, gigabit Ethernet will use IEEE 802.1p, 802.1Q, and the IETF's RSVP to render a pseudo QoS service. However, the RSVP is a "best effort" protocol; meaning that the network may acknowledge a QoS request but may not be capable of providing the service. Fibre channel, on the other hand, provides limited support for QoS parameters. Only its class 4 service supports QoS parameters such as guaranteed bandwidth and latency.

7 Conclusions

With today's faster desktops and servers, the rapidly increasing number of network users, enterprise Intranets, and LANs, and increasing use of World Wide Web (WWW), and other more bandwidth-intensive applications, significant pressure for greater bandwidth is being recognized at desktops, servers, hubs, and switches. Most existing networks are not capable of supporting this rapidly increasing demand. The concern, therefore, is how to provide enough speed and bandwidth to satisfy these rapidly growing demands. The result has been a significantly higher level of interest in gigabit networking technologies. These technologies provide connections between workstations, servers, and peripherals at a signalling rate of at least 1 Gb/s.

Advancements in fiber optics, computing systems and networking technologies have made the development of such technologies possible. Gigabit Ethernet, Fibre Channel, HIPPI-800, HIPPI 6400 or GSN, ATM-OC12, and ATM-OC48 are all examples of technologies that can operate at or above 1 Gb/s. Various gigabit network testbeds have proved that gigabit networking is feasible. Standards organizations are actively working on the standardization of these technologies to ensure interoperability, reliability, and security. Many standards have been developed and products have begun to appear in the market. Consequently, gigabit networks have become a reality.

Currently, most installed workstations and personal computers do not have the capacity to use these high bandwidth networks. So, the imminent battle is for the backbones, the network connections between switches and servers in large networks. However, the selection of a specific gigabit networking technology depends on the connection objectives or the requirements of a particular environment. Therefore, it is most unlikely that a single technology will satisfy all of an organization's connectivity requirements. For example, Fibre channel, HIPPI, and HIPP-6400 are best suited for interconnecting workstations, servers, and supercomputers. Gigabit Ethernet is the best fit for the expansion and consolidation of Ethernet and Fast Ethernet Networks. ATM, on the other hand, is the best technology to support time-sensitive data, voice, and video over WANs. Hence, to achieve the maximum benefit, interoperability of these technologies is essential. Besides this, other connection objectives such as priority, QoS support, and security must also be supported.

Many standards have been defined to facilitate interoperability among existing and emerging gigabit networking technologies. Today, various vendors are offering solutions that render interoperability of these technologies, but it is nevertheless fully possible for one vendor to create proprietary extensions to standardized products. Regarding QoS and security support, only ATM has well defined QoS and a security framework. Gigabit Ethernet supports QoS through IEEE 802.1p, IEEE 802.1Q, and RSVP, but the QoS is not guaranteed. Fibre channel also provides limited QoS support. To protect QoS, only traffic-level priority has been supported. Thus, user-level priority is still a very important unresolved NCS NS/EP issue.

Gigabit networks are potential candidates to support NS/EP requirements over the years to come. They offer speed and large bandwidth for medical data transfer and other bandwidth-intensive applications that will help to improve the NS/EP event management process. However, for a successful integration of these technologies in to the NS/EP management environment, they must first support the core NS/EP requirements of interoperability, reliability, security, and priority. Presently, these requirements are not fully supported by commercial off-the-shelf technologies.

8 Recommendations

1. Standards organizations such as IEEE and ANSI are working on the standardization of gigabit networking technologies. The OMNCS should participate in the standards development process to promote NS/EP requirements such as interoperability, reliability, priority, and security. It is further recommended that the NCS continue supporting the service mechanisms that industry is voluntarily providing with respect to providing limited support for NS/EP requirements. These include a variety of communications technology (lower layers infrastructure) mechanisms such as QoS, Resource Reservation Protocol, and others.

2. The OMNCS should also develop schemes for providing user-level priorities and advocate adoption of these by the appropriate standards bodies.

3. Many commercial service providers are planning to migrate towards a gigabit networking infrastructure. The OMNCS should coordinate with such service providers and encourage them to provide support for NS/EP features.

4. Today's Internet is also evolving towards a gigabit infrastructure. For example, government research agencies have already been working to develop advanced networking technologies and applications on testbeds that are 100 to 1,000 faster than today's Internet. The efforts of these agencies have been labeled the Next Generation Internet (NGI) initiative. In addition to the NGI effort, more than 120 U.S. universities are working in partnership with industry and government to develop advanced Internet technologies. This effort is known as Internet2 and it is a project of the University Corporation for Advanced Internet Development (UCAID). The OMNCS should recognize these efforts and look into the possibility of using this evolving infrastructure for NS/EP purposes.

5. It is recommended that the OMNCS recognize that even dramatic increases in bandwidth is no substitute for addressing the core of NS/EP requirements; namely interoperability, security, priority, and applications services that directly support NS/EP event management. Increased communications bandwidth, such as gigabit networks, instead are most often engaged in a "catch up" role with respect to computer architecture advances. Consequently, larger more powerful applications and services will continue to "consume" all of the available communications bandwidth leaving the NCS with inadequate resources during the critical times when it is most needed. This is most likely to continue to be true even as the communications infrastructure technology evolves from gigabit through yottabit speeds.

6. From an NS/EP perspective the current industry directions that support QoS across ATM WAN environments but then downgrade the QoS to non-guaranteed service levels over gigabit Ethernet LANs generally is unsatisfactory with respect to NCS priority requirements. Likewise, priority support must be provided

directly to the NS/EP event management team by means of selection (e.g., urgent, priority, routine) at the application level such as within a Telemedicine application interface. Such support for NCS requirements must be reflected in all middleware components, including CORBA ORBs. The support must then be correspondingly carried down the communications architecture (TCP/IP protocol stack) and ultimately acted upon by switches, routers, and other intervening communications subsystems comprising the NS/EP network.

7. The commercial viability of certain NCS requirements (e.g., priority, security) may be open to debate. Industry may continue to assert that certain NCS requirements (e.g., priority, security) are simply not everyday commercial requirements that warrant a voluntary commercial resolution. Nevertheless, the OMNCS should continue to be actively involved in the standardization process, especially with respect to the application and middleware layers of infrastructure, which today essentially provide no direct or meaningful support for important NS/EP requirements.

   What should not be open to debate is the question of whether or not NS/EP events will happen or not; they will occur. Because NS/EP events are largely neither preventable nor predictable it is recommended that the OMNCS be prepared, in conjunction with appropriate political and regulatory authorities, to step forward with mandates where they are necessary to insure that the best interests of society at large are preserved, even by regulation if required.

Appendix A: Acronyms

AAI	ACTS ATM Internetworking
AAL	ATM Adaptation Layer
ABR	Available Bit Rate
ACTS	Advanced Communications Technology Satellite
ANSI	American National Standards Institute
ARP	Address Resolution Protocol
ATM	Asynchronous Transfer Mode
B-ICI	Broadband Intercarrier Interface
C2I CIS	Command, Control and Intelligence, Communications and Information Systems
CAC	Connection Admission Control
CAD/CAE	Computer Aided Design/ Computer Aided Engineering
CAT	Computed Axial Tomography
CBR	Constant Bit Rate
CDDI	Copper Distributed Data Interface
CLP	Cell Loss Priority
COE	Common Operating Environment
CORBA	Common Object Request Broker Architecture
CPUs	Central Processing Units
CRC	Cyclic Redundancy Check
CRISP-DM	CRoss Industry Standard Process for Data Mining
CSMA/CD	Carrier Sense Multiple Access with Collision Detection
DA	Destination Address
DECnet	Digital Equipment Corporation Network
DII	Defense Information Infrastructure
DIS	Draft International Standard
DSIG	Domain Special Interest Group
DXI	Data Exchange Interface
E.O.	Executive Order
ECRC	End-to-end CRC
EHTO	European Health Telematics Observatory
EOP	Executive Office of the President
FC	Fibre Channel
FCA	Fibre Channel Association
FC-LE	Fibre Channel - Link Encapsulation
FCS	Frame Check Sequence

FCSI	Fibre Channel Systems Initiative
FDDI	Fiber Distributed Data Interface
FEMA	Federal Emergency Management Agency
FIRP	Federal Internetworking Requirements Panel
Gb/s	Giga bits per second
GMII	Gigabit Media Independent Interface
GSN	Gigabyte System Network
HIPPI	High Performance Parallel Interface
HIPPI-PH	HIPPI - Physical layer
HIPPI-SC	HIPPI - Switch Control
HIPPI-ST	HIPPI - Scheduled Transfers
HISB	Healthcare Informatics Standards Board
HL7	Health Level Seven
I/O	Input/Output
ID	Identification
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task force
IMIA	International Medical Informatics Association
IP	Internet Protocol
IPI	Intelligent Peripheral Interface
IPX	Internet Packet eXchange
IRC	Internet Relay Chat
ISO	International Organization for Standardization
ISP	Internet Service Provider
ITU-T	International Telecommunications Union - Telecommunication Standardization Sector
LANs	Local Area Networks
LCRC	Link CRC
LLC	Logical Link Control
LLRC	Length-Longitudinal Redundancy Check
LUNI	LAN Emulation Interface
LW	Long-wave
MAC	Media Access Control
Mb/s	Mega bits per second
MDI	Medium Dependent Interface
MIB	Management Information Base
MIPS	Mega Instructions Per Second
M&M	Modeling and Methodology
NCITS	National Committee for Information Technology Standardization
NCS	National Communications System

NetBEUI	NetBIOS Extended User Interface
NGI	Next Generation Internet
NIC	Network Interface Card
NRM	Network Resource Management
nrt-VBR	non-real time Variable Bit Rate
NS/EP	National Security and Emergency Preparedness
OMG	Object Management Group
OMNCS	Office of the Manager, National Communications System
OSI	Open Systems Interconnection
PBC	Port Bypass Circuit
PCS	Physical Coding Sub-layer
PH	Physical
PMA	Physical Medium Attachment
PMD	Physical Media Dependent
PVCs	Permanent Virtual Connections
QoS	Quality of Service
QoSF	QoS Facilitator
RARP	Reverse Address Resolution Protocol
RD	Runnig Disparity
RFC	Request For Comment
RIM	Reference Information Model
RSVP	Resource Reservation Protocol
RTCP	RTP Control Protocol
RTP	Real-time Transport Protocol
rt-VBR	real-time Variable Bit Rate
SA	Source Address
SBCCS	Single Byte Command Code Set
SBU	Sensitive But Unclassified
SCSI	Small Computer System Interface
SFD	Start of Frame Delimiter
SIGOBT	Object Brokering Technologies Special Interest Group
SNMP	Simple Network Management Protocol
STUs	Scheduled Transfer Units
SVCs	Switched Virtual Connections
SW	Short-Wave
TC	Transmission Convergence
TCP	Transmission Control Protocol
UBR	Unspecified Bit Rate
UCAID	University Corporation for Advanced Internet Development

UDP	User Datagram Protocol
ULA	Universal LAN Address
UNI	User-to-Network Interface
UPC	Usage Parameter Control
UTP	Unshielded Twisted Pair
VC	Virtual Channel
VCIs	Virtual Channel Connection Identifiers
VPIs	Virtual Path Connection Identifiers
VPs	Virtual Paths
WANs	Wide Area Networks
WWW	World Wide Web
XIRC	Extensions to the Internet Relay Chat Protocol

Appendix B: References

1. ATM Forum af-tm-0056.000, Traffic Management Specification, Version 4.0, April 1996

2. ATM Forum af-uni-0010.002, ATM User-Network Interface Specification V 3.1, 1994

3. ATM Forum af-sig-0061.000, UNI Signalling Specification, Version 4.0, July 1996.

4. NCITS Working draft, Information Technology -High-Performance Parallel Interface 6400 Mb/s Physical Layer (HIPPI-6400-PH), July 22, 1998

5. NCITS Working draft, Information Technology -High-Performance Parallel Interface - 6400 Mb/s Physical Switch Control (HIPPI-6400-SC), August 11, 1998

6. NCITS Working draft, Information Technology -Scheduled Transfer Protocol (ST), September 23, 1998

7. NCITS Working draft, Information Technology –High-Performance Parallel Interface –Encapsulation of ISO/IEC 8802-2 (IEEE Std 802.2) Logical Link Control Protocol Data Units (HIPPI-LE), September 21, 1998

8. ANSI X3T11/92-REV 8.2 HIGH-PERFORMANCE PARALLEL INTERFACE - Mechanical, Electrical, and Signalling Protocol Specification (HIPPI-PH), March 3, 1993

9. Charalambous P.Chralambos, Georgios Y. Lazarou, Victor S. Frost, Joseph Evans, and Roelof Jonkman, Experimental and simulation Performance Results of TCP/IP over High-Speed ATM over ACTS, Department of Electrical Engineering & Computer Science, The University of Kansas, Lawrence, KS, September 10, 1997.

10. Stephen Pink, TCP/IP on Gigabit Networks, Swedish Institute of Computer Science, Kista, Sweden.

11. Oikarien, J. et al., Internet Request for Comments (RFC) 1459. May 1993.

Appendix C: Referenced Fibre Standards

1. ANSI X3.254, Information Technology - Fibre Channel-Mapping - to HIPPI-FC (FC-FP), 1994

2. ANSI X3.303, Fibre Channel Physical and Signalling Interface-3 (FC-PH-3), 1998

3. ANSI X3.272, Information Technology - Fibre Channel - Arbiturated Loop (FC-AL), 1996

4. ANSI X3.287, Information Technology - Fibre Channel - Link Encapsulation (FC-LE), 1996

5. ANSI X3.288, Information Technology - Fibre Channel - Generic Services (FC-GS), 1996

6. ANSI X3.230, Information Technology - Fibre Channel - Physical and Signalling Interface (FC-PH), 1994

7. ANSI X3.289, Information Technology - Fibre Channel - Fabric Generic Requirements (FC-FG), 1996

8. ANSI X3.297, Information Technology - Fibre Channel - Physical and Signalling Interface-2 (FC-PH-2), 1997

9. ANSI NCITS 321-1998 : Information Technology - Fibre Channel - Switch Fabric (FC-SW)

10. ANSI X3.230-1994/AM 1-1996 : Information Technology - Fibre Channel Physical and Signalling Interface (FC-PH) -Amendment 1

11. ANSI X3.271-1996 : Fibre Channel Single-Byte Command Code Sets (SBCCS) Mapping Protocol (FC-SB)

12. ANSI X3.283-1996 : Information Technology - HIGH-PERFORMANCE PARALLEL INTERFACE - Encapsulation of Frames of the Fibre Channel Physical and Signalling Interface (FC-PH Encapsulation) (HIPPI-FC)

13. ANSI X3.269-1996 : Information Technology - Fibre Channel Protocol for SCSI

14. Internet draft, IP and ARP over Fibre Channel", Raj Bhagwat, Murali Rajagopal, Wayne Rickard, 10/01/1998

15. Internet draft, Fibre Channel Interconnect MIB", Kim Banker, 08/06/1998

16. Internet draft, Definitions of Managed Objects for the Fabric Element in Fibre Channel Standard, 11/03/1998

Appendix D: Referenced HPPI Standards

1. ANSI X3.254, Information Technology - Fibre Channel-Mapping to HIPPI-FC (FC-FP), 1994

2. ANSI X3.300, Information Technology - High-Performance Parallel Interface - Serial Specification (HIPPI-Serial), 1997

3. ANSI X3.210, Information Technology - High-Performance Parallel Interface - Framing Protocol (HIPPI-FP), 1992

4. ANSI NCITS 210, Information Systems - High-Performance Parallel Interface - Framing Protocol (HIPPI-FP), 1998

5. ANSI X3.222, Information Technology - High-Performance Parallel Interface - Physical Switch Control (HIPPI-SC), 1997

6. ANSI X3.299, Information Technology - High-Performance Parallel Interface - Mapping to Asynchronous Transfer Mode (HIPPI-ATM), 1997

7. ISO/IEC 11518-2, Information technology - High-Performance Parallel Interface - Part 2: Framing Protocol (HIPPI-FP), 1996

8. ANSI X3.183, Information Systems - High-Performance Parallel Interface - Mechanical, Electrical, and Signalling Protocol Specification (HIPPI-PH), 1991

9. ANSI X3.210, Information Technology - High-Performance Parallel Interface - Framing Protocol (HIPPI-FP) - Amendment 1, 1996

10. ISO/IEC 11518-6, Information technology - High-Performance Parallel Interface - Part 6: Physical Switch Control (HIPPI-SC), 1996

11. ISO/IEC 11518-1, Information technology - High-Performance Parallel Interface - Part 1: Mechanical, electrical and signalling protocol specification (HIPPI-PH), 1995

12. ANSI X3.283, Information Technology - HIGH-PERFORMANCE PARALLEL INTERFACE - Encapsulation of Frames of the Fibre Channel Physical and Signalling Interface (FC-PH Encapsulation) (HIPPI-FC), 1996

13. ANSI X3.218, Information Systems - High-Performance Parallel Interface - Encapsulation of ISO 8802-2 (IEEE Std 802.2) Logical Link Control Protocol Data Units (HIPPI-LE), 1993

14. ISO/IEC 11518-3, Information technology - High-Performance Parallel Interface - Part 3: Encapsulation of ISO/IEC 8802-2 (IEEE Std 802.2) Logical Link Control Protocol Data Units (HIPPI-LE), 1996

15. NCITS Working draft, Information Technology -High-Performance Parallel Interface 6400 Mb/s Physical Layer (HIPPI-6400-PH), July 22, 1998

16. NCITS Working draft, Information Technology -High-Performance Parallel Interface - 6400 Mb/s Physical Switch Control (HIPPI-6400-SC), August 11, 1998

17. NCITS Working draft, Information Technology -Scheduled Transfer Protocol (ST), September 23, 1998

18. IETF Draft RFC 2067, IP over HIPPI, January 1997

19. IETF Proposed RFC 1374, IP and ARP on HIPPI, October 1992

Appendix E: Referenced ISDN Standards

1. ATM Forum af-tm-0056.000, Traffic Management Specification, Version 4.0, April 1996

2. ATM Forum af-uni-0010.002, ATM User-Network Interface Specification V 3.1, 1994

3. ATM Forum af-sig-0061.000, UNI Signalling Specification, Version 4.0, July 1996

4. ATM Forum af-phy-0015.000, ATM Physical Medium Dependent Interface specification for 155 Mb/s over Twisted Pair Cable, September 1994

5. ATM Forum af-phy-0046.000, 622.08 Mb/s Physical Layer, January 1996

6. ATM Forum af-phy-0061.000, DS3 Physical Layer Interface Specification, March 1996

7. ANSI T1.511, Telecommunications - B-ISDN ATM Layer Cell Transfer Performance, 1997

8. ANSI T1.646a, Telecommunications - Broadband ISDN - Physical Layer Specification for User-Network Interfaces Including DS1/ATM, 1997

Appendix F: Ethernet Standards

1. IEEE 802.3z, Media Access Control Parameters, Physical Layers, Repeater and Management Parameters for 1,000 Mb/s Operation, Supplement to Information Technology - Local and Metropolitan Area Networks - Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, 1998

2. IEEE P802.3ab, Standard for Information Technology - Telecommunications and information exchange between systems - Local and Metropolitan Area Networks - Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications - Physical Layer Parameters and Specifications for 1,000 Mb/s Operation over 4 pair of Category 5 Balanced Copper Cabling, Type 1,000BASE-T

3. IEEE 802.1Q, Standard for Virtual Bridged Local Area Networks.

4. IEEE 802.1p, Standard for Local and Metropolitan Area Networks Supplement to Media Access Control (MAC) Bridges: Traffic Class Expediting and Dynamic Multicast Filtering

5. IEEE 802.3x, IEEE Standards for Local and Metropolitan Area Networks: Specification for 802.3 Full Duplex Operation, 1997

Appendix G: Referenced Telemedicine Standards

1. Health Industry Bar Code Supplier Labeling Standard Health Industry Business Communications Council
2. HL7 Version 2.3, Chapter 12 Patient Care, Health Level Seven
3. HL7 Version 2.3, Chapter 11, Patient Referral, Health Level Seven
4. HL7 Version 2.3, Chapter 10 Scheduling, Health Level Seven
5. HL7 Version 2.3, Chapter 9 Medical Records/Information Management, Health Level Seven
6. HL7 Version 2.3, Chapter 8 Master Files, Health Level Seven
7. HL7 Version 2.3, Chapter 7 Observation Reporting, Health Level Seven
8. HL7 Version 2.3, Chapter 6 Financial Management, Health Level Seven
9. HL7 Version 2.3, Chapter 4 Order Entry, Health Level Seven
10. HL7 Version 2.3, Chapter 3 Patient Administration, Health Level Seven
11. HL7 Version 2.3, Chapter 2 Control/Query, Health Level Seven
12. The classification of the Electronic Patient Records Systems by JAHIS, V1.1 JAHIS, Japanese Association of Healthcare Information Systems Industry
13. The standard guidelines of designing security control system for the Electronic Patient Records Systems, V2.0, JAHIS, Japanese Association of Healthcare Information Systems Industry
14. Healthcare Card System Standardization Manual Medical Information System Development Center(MEDIS-DC)
15. Master file of medical treatment for insurance claim MHW, Japan
16. Master file of disease and injury Ministry of healthcare and welfare (MHW), Japan
17. Pharmaceutics products master of social insurance payment fund Ministry of Healthcare and welfare, Japan

18. MEDIS-DC Standards for Electronic Filing of Medical Images MEDIS-DC, Japan

19. The Agreement on Clinical Laboratory Data Communication: JAHIS-DRAFT

20. JAHIS, Japanese Association of Healthcare Information Systems Industry, Japan

21. Standard guidelines for transfer message between automated analyzer and computer Industry, Japan

22. Classification & Coding for Clinical Laboratory Tests, 9th rev. 2nd ed. JSCP, Japan Society of Clinical Pathology, Japan

23. Medical record image text information exchange, Japan

24. Ambulatory Care Information Model NSW Health Department, Australia.

25. Clinical Cancer Management Information Model NSW Health Department, Australia.

26. Personal privacy protection in health care information systems Standards Australia, Australia

27. NSW Health Enterprise Information Model NSW Health Department, Australia

28. NSW Health Standard Technology Products and Acquisition Guidelines NSW Health Department, Australia, Australia

29. NSW Health Standards Framework NSW Health Department, Australia

30. Community Health Information Model NSW Health Department, Australia

31. Implementation of Health Level Seven (HL7) Version 2.2 Part 1: Admission, discharge and transfer Standards, Australia

32. IEEE 1073 Standard for Medical Device Communications, IEEE

33. Australian National Health Information Knowledgebase, National Health Information Management Group - produced by the Australian Institute of Health and Welfare Health, Australia

34. Tentative agreement on clinical laboratory data communication MEDIS-DC, Japan

35. Diagnostic Japanese Terminology in accordance with ICD10 MEDIS-DC (Medical Information System Development Center), Japan

36. NSW Health Information Privacy Code of Practice NSW Health Department, Australia

37. Patient Administration System Information Model, NSW Health Department, Australia

38. CIHI National Continuing Care Data Set Canadian Institute for Health Information, Canada

39. Pharmacy Claim Standard Canadian Pharmaceutical Association, Canada

40. Canadian Classification of Interventions/Classification canadienne des interventions Canadian Institute for Health Information (CIHI), Canada

41. Canadian Classification of Diagnostic, Therapeutic, and Surgical Procedures/Classification canadienne des actes diagnostiques, thérapeutiques et chirurgicaux Statistics, Canada

42. CIHI Discharge Abstract (Dataset for Acute Care), Canadian Institute for Health Information (formerly produced by the Hospital Medical Records Institute), Canada

43. CIHI Ambulatory Care Minimum Data Set Canadian Institute for Health Information, Canada

44. Case Mix Groups Canadian Institute for Health Information, Canada

45. Day Procedure Groups Canadian Institute for Health Information, Canada

46. Teleradiology Standard Ontario Association of Radiologists, Canada

47. CIHI Rehabilitation Minimum Data Set Canadian Institute for Health Information, Canada

48. Guidelines for Management Information Systems in Canadian Health Care Facilities/Guide sur les systèmes d'information de gestion dans les établissements de santé canadiens Canadian Institute for Health Information, Canada

49. Electronic Nomenclature and Classification of Disorders and Encounters in Family Medicine/Codification électronique pour la médecine familiale College of Family Physicians of Canada, in collaboration with the Centre for Information Technology Innovation of Industry Canada and CLINIDATA Inc., Canada

50. Medical Image Management - Part 1: Storage Commitment Service Class CEN/TC 251, Europe

51. Media Interchange for Medical Imaging Communication, CEN/TC 251, Europe

Appendix H: Web Sites

1. Fibre Channel Association (FCA) home page http://www.fibrechannel.com
2. ANSI on line http://web.ansi.org/default_js.htm
3. CERN Fibre Channel home page http://www1.cern.ch/HSI/fcs/
4. Fibre Channel Consortium http://www.iol.unh.edu/consortiums/fc/
5. The Fibre Channel Loop Community http://www.fcloop.org/
6. CERN.High Speed Interconnect project home page http://www1.cern.ch/HSI/
7. CERN HIPPI homepage http://www.cern.ch/HSI/hippi/
8. CERN GSN homepage http://www1.cern.ch/HSI/hippi/gsn/gsnhome.htm
9. CERN ATM page http://www1.cern.ch/HSI/atm/atm.html
10. WideBand Gigabit Networking Alliance (WGNA) http://www.wgna.org/
11. Washington University's Gigabit Network Technology Distribution Program http://boushi.arl.wustl.edu/gigabitkits/kits.html
12. Current Research in High Performance Networking Worldwide http://www.roads.lut.ac.uk/DS-Archive/Gigabit_Review/
13. IEEE Communications Society, Technical Committee on Gigabit Networking (TCGN) http://www.ccrc.wustl.edu/pub/ieee-tcgn/
14. Gigabit Networking Survey http://www.cis.ohio-state.edu/~jain/cis788/gigabit/index.html#Introduction
15. Gigabit Networking Workshop GBN'97 - Online Proceedings http://www.ccrc.wustl.edu/pub/ieee-tcgn/conference/gbn97/
16. Los Alamos National Laboratory http://www.lanl.gov/
17. Gigabit Ethernet Alliance, White Paper : Gigabit Ethernet Overview - May 1997, http://www.gigabit-ethernet.org/technology/whitepapers/gige_97/
18. CASA http://www.nlanr.net/CASA/casa.html
19. Gigabit Ethernet Alliance homepage http://www.gigabit-ethernet.org