SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF COMPUTER SCIENCE AND ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE ENGINEERING

Computer Networks and Design – SCSA1502

Computer Networks and Design – SCSA1502

COURSE OBJECTIVES

- To recognize the principles of the big picture of computer networks.
- To understand the networking environment.
- To know the importance of VPNs.
- To convey the availability of tools and techniques for networking.
- To discuss about evolving technologies in networks.

UNIT 1 FUNDAMENTALS OF NETWORK DESIGN 9 Hrs

Design Principles - Determining Requirements - Analyzing the Existing Network - Preparing the Preliminary Design - Completing the Final Design Development - Deploying the Network - Monitoring and Redesigning – Maintaining - Design Documentation - Modular Network Design - Hierarchical Network Design.

UNIT 2 UNDERLYING LAN CONCEPTS 9 Hrs

LAN connectivity for small businesses – Integration – Token-Ring – Ethernet – ATM LAN emulation – InterLAN Switching – LAN to Mainframe – Building networks.

UNIT 3 VPNS, INTRANETS AND EXTRANETS 9 Hrs

Virtual Network management and planning – VPNs for small businesses – Secure remote access in VPNs – IPSec VPNs – Integrating data centers with Intranets – Implementing and supporting Extranets.

UNIT 4 NETWORKING TOOLS AND TECHNIQUES 9 Hrs

Simulation method for designing multimedia networks – Determining remote bridge and router delays – Network baselining as a planning tool.

UNIT 5 EVOLVING TECHNOLOGIES 9 Hrs

Trends in data communications – Merits of xDSL technology – Preparing for cable modems - Voice and video on the LAN – Internet voice applications – Building IP PBX telephony network – Fax over IP – Videoconferencing over IP networks.

Max.45 Hrs.

COURSE OUTCOMES

On completion of the course, student will be able to
CO1 - Understand the principles of networks.
CO2 - Interpret LAN concepts and design.
CO3 - Gain knowledge in evolving technologies.
CO4 - Clearly outline the logic behind VPNs.
CO5 - Know the importance of tools and techniques in building a network.
CO6 - Understand the underlying working concepts of a real-time network.

TEXT/ REFERENCE BOOKS

1. Gil Held, "Network Design: Principles and Applications (Best Practices)", Auerbach Publications, 1st edition, 2000.
2. Diane Tiare and Catherine Paquet, "Campus Network Design Fundamentals", Pearson Education, 1st edition, 2006.
3. Larry L. Peterson, Bruce S. Davie, "Computer Networks: A Systems Approach", Morgan Kaufmann Publishers Inc., 5th edition, 2012.
4. William Stallings, "Data and Computer Communications", Pearson Education, 8th edition, 2016.
5. James F. Kurose, Keith W. Ross, "Computer Networking - A Top-Down Approach Featuring the Internet", Pearson Education, 6th edition, 2012.

END SEMESTER EXAMINATION QUESTION PAPER PATTERN

Max. Marks : 100 **Exam Duration : 3 Hrs.**

PART A : 10 Questions of 2 marks each-No choic **20 Marks**

PART B : 2 Questions from each unit with internal choice, each carrying 16 marks **80 Marks**

SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF COMPUTER SCIENCE AND ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE ENGINEERING

Computer Networks and Design – SCSA1502

UNIT I - FUNDAMENTALS OF NETWORK-SCSA1502

UNIT I

FUNDAMENTALS OF NETWORK DESIGN

Design Principles - Determining Requirements - Analyzing the Existing Network - Preparing the Preliminary Design - Completing the Final Design Development - Deploying the Network - Monitoring and Redesigning – Maintaining - Design Documentation - Modular Network Design - Hierarchical Network Design.

INTRODUCTION

Network designers ensure that our communications networks can adjust and scale to the demands for new services. To support our network-based economy, designers must work to create networks that are available nearly 100 percent of the time. Information network security must be designed to automatically fend off unexpected security incidents. Using hierarchical network design principles and an organized design methodology, designers create networks that are both manageable and supportable.

NETWORK DESIGN CONCEPTS

What is the design methodology used by network designers?

- Network designers ensure that our communications networks can adjust and scale to the demands for new services.
- To support our network-based economy, designers must work to create networks that are available nearly 100 percent of the time.
- Information network security must be designed to automatically fend off unexpected security incidents.
- Using hierarchical network design principles and an organized design methodology, designers create networks that are both manageable and supportable

Basics of network design

- Network design overview
- The benefits of a hierarchical network design
- Network design methodology

NETWORK DESIGN OVERVIEW

Computers and information networks are critical to the success of businesses, both large and small. They connect people, support applications and services, and provide access to the resources that keep the businesses running. To meet the daily requirements of businesses, networks themselves are becoming quite complex.

Today, the Internet-based economy often demands around-the-clock customer service. This means that business networks must be available nearly 100 percent of the time. They must be smart enough to automatically protect against unexpected security incidents. These business networks must also be able to adjust to changing traffic loads to maintain consistent application

response times. It is no longer practical to construct networks by connecting many standalone components without careful planning and design.

DISCOVERING NETWORK DESIGN BASICS

The sections that follow cover the basics of network design with regard to the following concepts

- Network design overview
- The benefits of a hierarchical network design
- Network design methodology

Network Design Overview

Computers and information networks are critical to the success of businesses, both large and small. They connect people, support applications and services, and provide access to the resources that keep the businesses running. To meet the daily requirements of businesses, networks themselves are becoming quite complex.

NETWORK REQUIREMENTS

Today, the Internet-based economy often demands around-the-clock customer service. This means that business networks must be available nearly 100 percent of the time. They must be smart enough to automatically protect against unexpected security incidents. These business networks must also be able to adjust to changing traffic loads to maintain consistent application response times. It is no longer practical to construct networks by connecting many standalone components without careful planning and design.

Technical Requirements

- Applications that are to run on the network
- Internet connections required
- Addressing restrictions, for example, the use of private Internet Protocol (IP) version 4 (IPv4) addresses
- Support for IP version 6 (IPv6) addresses
- Other protocols that are to run on the network (for example, routing protocols)
- Cabling requirements
- Redundancy requirements
- Use of proprietary equipment and protocols
- Existing equipment that must be supported
- Network services required, including quality of service (QoS) and wireless
- How security is to be integrated into the network
- Network solutions required (for example, voice traffic, content networking, and storage networking)
- Network management
- Support for existing applications while new ones are being phased in
- Bandwidth availability.
- Requirements Related to Business Issues
- Budget- Capital (for new equipment) and operating (for ongoing expenses).

- Schedule - This could include the phasing out of older applications, hiring of new personnel, and so forth.
- People - Considerations include who will install and operate the network, what skills they have, whether they require training, whether any of these tasks will be outsourced, and so forth.
- Legal Issues include any restrictions on the use and storage of data collected
- History Factors include examining the existing network's structure and determining whether any person or group will block changes or additions.
- Policies Consider whether current organizational policies might restrict the network design.
-
- Most businesses actually have only a few requirements for their network
- The network should stay up all the time, even in the event of failed links, equipment failure, and overloaded conditions.
- The network should reliably deliver applications and provide reasonable response times from any host to any host.
- The network should be secure. It should protect the data that is transmitted over it and data stored on the devices that connect to it.
- The network should be easy to modify to adapt to network growth and general business changes.
- Because failures occasionally occur, troubleshooting should be easy. Finding and fixing a problem should not be too time-consuming.

Building a Good Network

Good networks do not happen by accident. They are the result of hard work by network designers and technicians, who identify network requirements and select the best solutions to meet the needs of a business.

The steps required to design a good network are as follows
Step 1. Verify the business goals and technical requirements.
Step 2. Determine the features and functions required to meet the needs identified in Step 1.
Step 3. Perform a network-readiness assessment.
Step 4. Create a solution and site acceptance test plan.
Step 5. Create a project plan.

After the network requirements have been identified, the steps to designing a good network are fol- lowed as the project implementation moves forward. Network users generally do not think in terms of the complexity of the underlying network. They think of the network as a way to access the applications they need, when they need them.

Network Requirements

Most businesses actually have only a few requirements for their network

The network should stay up all the time, even in the event of failed links, equipment failure, and overloaded conditions. The network should reliably deliver applications and provide reasonable response times from any host to any host. The network should be secure. It should

protect the data that is transmitted over it and data stored on the devices that connect to it. The network should be easy to modify to adapt to network growth and general business changes. Because failures occasionally occur, troubleshooting should be easy. Finding and fixing a problem should not be too time-consuming.

FUNDAMENTAL DESIGN GOALS

When examined carefully, these requirements translate into four fundamental network design goals

Scalability Scalable network designs can grow to include new user groups and remote sites and can support new applications without impacting the level of service delivered to existing users.

Availability A network designed for availability is one that delivers consistent, reliable performance, 24 hours a day, 7 days a week. In addition, the failure of a single link or piece of equipment should not significantly impact network performance.

Security Security is a feature that must be designed into the network, not added on after the network is complete. Planning the location of security devices, filters, and firewall features is critical to safeguarding network resources.

Manageability No matter how good the initial network design is, the available network staff must be able to manage and support the network. A network that is too complex or difficult to maintain cannot function effectively and efficiently.

NETWORK DESIGN

Implementation Components

Implementation of a network design consists of several phases (install hardware, configure systems, launch into production, and so on).

Each phase consists of several steps, and each step should contain, but be not limited to, the following documentation

- Description of the step
- Reference to design documents
- Detailed implementation guidelines
- Detailed roll-back guidelines in case of failure
- Estimated time needed for implementation

Analysing the Existing Network

The second step of the design methodology is characterizing the existing network and sites

The following sections present insights into the process of examining an existing network and sites and describe the tools used to gather the data, assess the network, and analyze the network.

- Customer input Review existing documentation about the network, and use verbal input from the customer to obtain a first impression about the network
- Network audit Perform a network audit, also called an assessment, which reveals details of the network and augments the customer's description
- Traffic analysis If possible, use traffic analysis to provide information about the applications and protocols used and to reveal any shortcomings in the network.
- Customer input includes all pertinent network and site documentation.

- Some items the designer could request, depending on the scope of the project, include the following
- Site contact information (especially needed if remote deployments are planned)
- Existing network infrastructure (from physical diagrams and documents, and site surveys as needed), including the following
- Locations and types of servers, Locations and types of network devices Cabling that is currently in place Environmental controls, including heating, ventilation, and air conditioning requirements, and filtration Locations of power receptacles etc
- Existing network infrastructure from logical topology diagrams, routing protocols in use, and the infrastructure services supported, such as voice, storage, and wireless services
- Network topology Includes devices, physical and logical links, external connections, bandwidth of connections, frame types (data link encapsulations), IP addressing, routing protocols, and so forth.
- Network services Includes security, QoS, high availability, voice, storage, wireless, and so forth.
- Network applications Examples include unified messaging and video delivery
- The Second important step is network audit or assessment
- It is used to collect information about an existing network
- An audit provides details such as
- A list of network devices
- Hardware specifications and versions, and software versions of network devices
- Configurations of network devices
- Output of various auditing tools to verify and augment the existing documentation
- Link, CPU, and memory utilization of network devices
- A list of unused ports, modules, and slots in network devices, (to be used to understand whether the network is expandable)

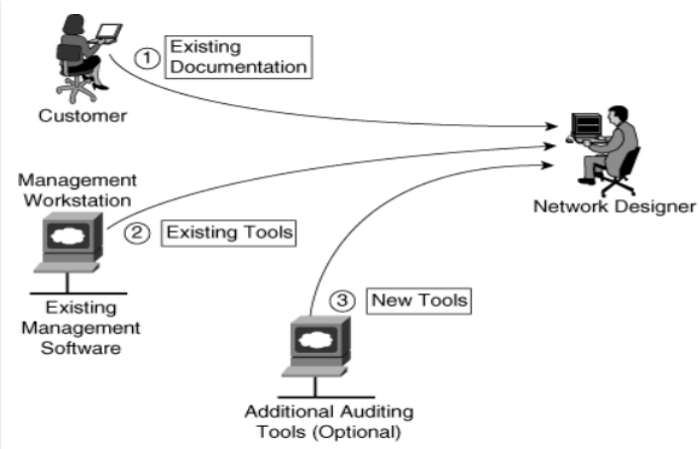

Figure 1 Existing Network system auditing sources

PREPARING THE PRELIMINARY DESIGN

- Preliminary design involves considering all the network requirements and constraints (including the budget), and determining viable alternative solutions.
- The network owner consulted, and together an optimal solution is chosen
- This solution is later developed into the final design

- (Both the preliminary design and final design are done Using PDIOO)

FOLLOWING A DESIGN METHODOLOGY CAN HAVE MANY ADVANTAGES

- It ensures that no step is missed when the process is followed
- It provides a framework for the design process deliverables
- It encourages consistency in the creative process, enabling network designers to set appropriate deadlines and maintain customer and manager satisfaction.
- It allows customers and managers to validate that the designers have thought about how to meet their requirements

Preliminary Design Step involves in the preparation of detailed Documentation of the network

It is achieved by collecting information from customer, inspecting the site, and accessing the network using Automated Tools

Different Approaches

1. Top Down Approach

A top-down approach to network design means that requirements are considered first, with the applications and network solutions that will run on the network driving the design

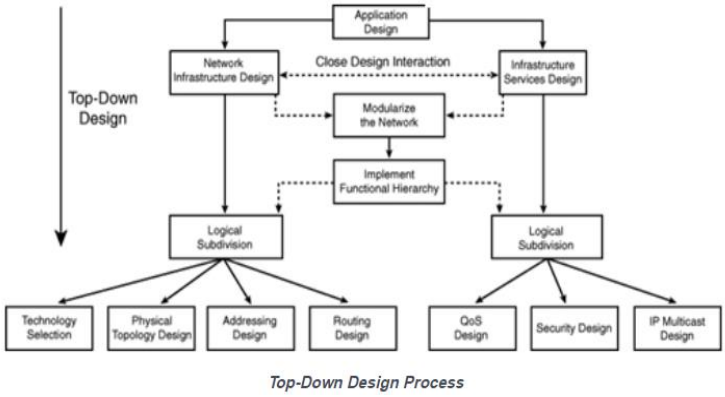

Figure 2 Top-Down Approach

Bottom- Up Approach

A bottom-up approach would first select devices, features, cabling, and so on, and then try to fit the applications onto this network

Issues in Bottom-Up Approach

A bottom-up approach can lead to redesign if the applications are not accommodated properly.

This approach can also result in increased costs by including features or devices that are not required

A bottom-up approach would first select devices, features, cabling, and so on, and then try to fit the applications onto this network

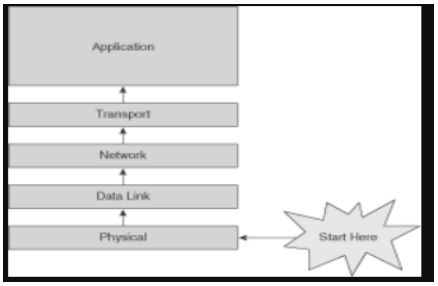

Figure 3 Bottom-Up Approach

Issues in Bottom-Up Approach

A bottom-up approach can lead to redesign if the applications are not accommodated properly. This approach can also result in increased costs by including features or devices that are not required

THE BENEFITS OF A HIERARCHICAL NETWORK DESIGN

To meet the four fundamental design goals, a network must be built on an architecture that allows for both flexibility and growth.

Hierarchical Network Design

In networking, a hierarchical design is used to group devices into multiple networks. The networks are organized in a layered approach. The hierarchical design model has three basic layers

Core layer Connects distribution layer devices

Distribution layer Interconnects the smaller local networks

Access layer Provides connectivity for network hosts and end devices

Hierarchical networks have advantages over flat network designs. The benefit of dividing a flat network into smaller, more manageable hierarchical blocks is that local traffic remains local. Only traffic destined for other networks is moved to a higher layer. Layer 2 devices in a flat network provide little opportunity to control broadcasts or to filter undesirable traffic. As more devices and applications are added to a flat network, response times degrade until the network becomes unusable. Figures 1-1 and 1-2 show the advantages of a hierarchical network design versus a flat network design.

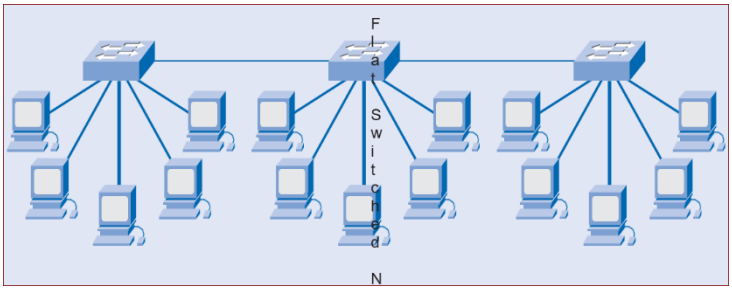

Figure 4 Flat Network

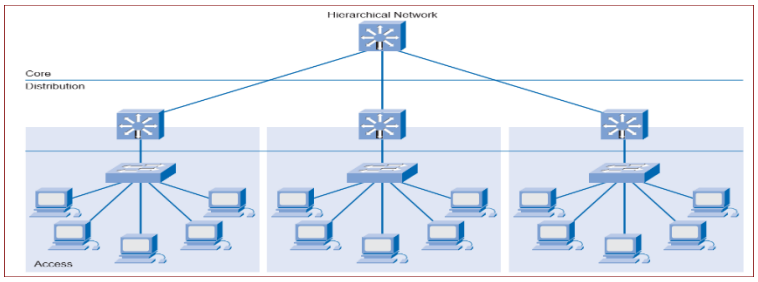

Figure 5 Hierarchical Network (Three Separate Broadcast Domains)

Modular Design of Cisco Enterprise Architectures

The Cisco Enterprise Architectures (see Figure 1-3) can be used to further divide the three-layer hierarchical design into modular areas. The modules represent areas that have different physical or logical connectivity. They designate where different functions occur in the network. This modularity enables flexibility in network design. It facilitates implementation and troubleshooting. Three areas of focus in modular network design are as follows

Enterprise campus This area contains the network elements required for independent operation within a single campus or branch location. This is where the building access, building distribution, and campus core are located.

Server farm A component of the enterprise campus, the data centre server farm protects the server resources and provides redundant, reliable high-speed connectivity.

Enterprise edge As traffic comes into the campus network, this area filters traffic from the external resources and routes it into the enterprise network. It contains all the elements required for efficient and secure communication between the enterprise campus and remote locations, remote users, and the Internet.

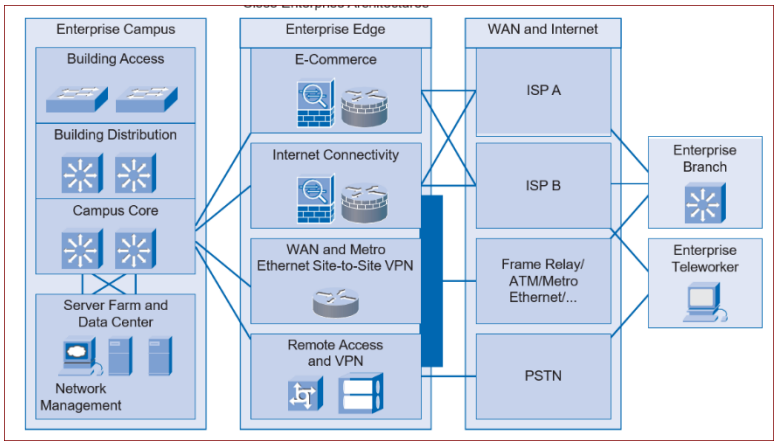

Figure 6 Cisco Enterprise Architectures

The modular framework of the Cisco Enterprise Architectures as depicted in Figure 1-4 has the following design advantages It creates a deterministic network with clearly defined boundaries between modules. This provides clear demarcation points so that the network designer knows exactly where the traffic originates and where it flows. It eases the design task by making each module independent. The designer can focus on the needs of each area

separately. It provides scalability by allowing enterprises to add modules easily. As network complexity grows, the designer can add new functional modules. It enables the designer to add services and solutions without changing the underlying network design.

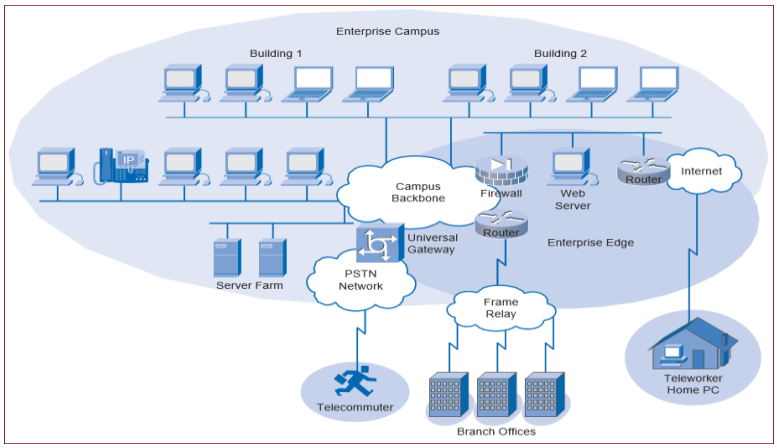

Figure7 Enterprise Campus

Interactive Activity 1-1 Match the Characteristics of the Hierarchal Model and the Cisco Enterprise Architecture (1.1.2) In this interactive activity, you match the characteristics of the hierarchal model and the Cisco Enterprise Architecture to their correct location. Use file ia-112 on the CD-ROM that accompanies this book to perform this interactive activity

NETWORK DESIGN METHODOLOGIES

Large network design projects are normally divided into three distinct steps
Step 1. Identify the network requirements.
Step 2. Characterize the existing network.
Step 3. Design the network topology and solutions.

Step 1 Identifying Network Requirements

The network designer works closely with the customer to document the goals of the project. Figure 1-5 depicts a meeting between the designer and the business owner. Goals are usually separated into two categories
Business goals Focus on how the network can make the business more successful
Technical requirements Focus on how the technology is implemented within the network

Step 2 Characterizing the Existing Network

Information about the current network and services is gathered and analysed. It is necessary to compare the functionality of the existing network with the defined goals of the new project. The designer determines whether any existing equipment, infrastructure, and protocols can be reused, and what new equipment and protocols are needed to complete the design.

Step 3 Designing the Network Topology

A common strategy for network design is to take a top-down approach. In this approach, the network applications and service requirements are identified, and then the network is

designed to support them. When the design is complete, a prototype or proof-of-concept test is performed. This approach ensures that the new design functions as expected before it is implemented.

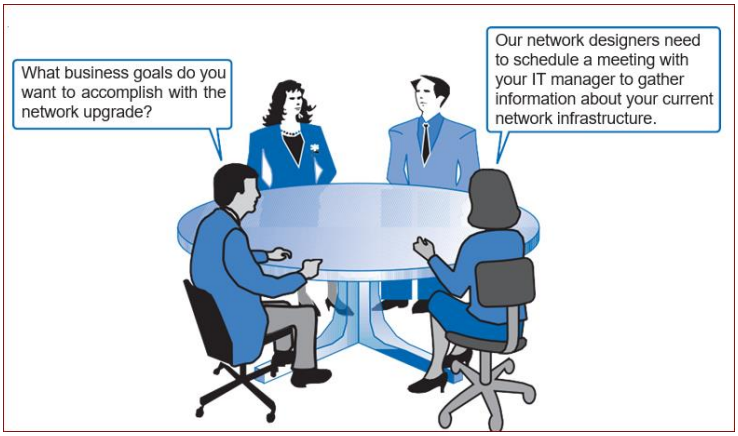

Figure 8 Enterprise Campus

A common mistake made by network designers is the failure to correctly determine the scope of the network design project.

Determining the Scope of the Project

While gathering requirements, the designer identifies the issues that affect the entire network and those that affect only specific portions. By creating a topology similar to Figure 1-6, the designer can isolate areas of concern and identify the scope of the project. Failure to understand the impact of a particular requirement often causes a project scope to expand beyond the original estimate. This oversight can greatly increase the cost and time required to implement the new design.

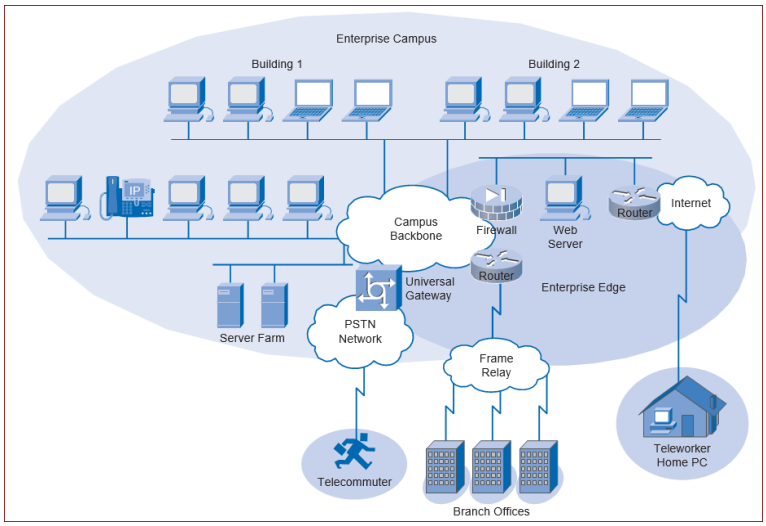

Figure 9 Enterprise Campus

Impacting the Entire Network

Network requirements that impact the entire network include the following

- Adding new network applications and making major changes to existing applications, such as database or Domain Name System (DNS) structure changes
- Improving the efficiency of network addressing or routing protocol changes
- Integrating new security measures
- Adding new network services, such as voice traffic, content networking, and storage networking
- Relocating servers to a data centre server farm

Impacting a Portion of the Network

Requirements that may only affect a portion of the network include the following

- Improving Internet connectivity and adding bandwidth
- Updating access layer LAN cabling
- Providing redundancy for key services
- Supporting wireless access in defined areas
- Upgrading WAN bandwidth

COMPLETING THE FINAL DESIGN DEVELOPMENT

- During the final design stage the detailed architectural and engineering drawings (the blueprints) of all physical components of the Network components are produced.
- In some complex projects, it is necessary to prepare in addition a written final design report.
- Sufficient detail must be provided by the drawings and the report should have reasonably accurate estimates involved in the process.
- All revisions to materials, equipment specifications are made. The updated schedule, cost estimates and specifications should be available in the final design report.
- It is essential to verify at the final design stage that the plan remains economically feasible. If, by some chance it is not, then a decision must be made to revise design solutions or the original concepts, or perhaps terminate the project

Deploying the Network

- Deployment of the network must start with a plan and a schedule.
- Deployment planning starts in the PDIOO Design phase and continues into the Implement phase
- It contains – What to be done? / How to be Done?
- Contingency plans, that is, plans for what happens if a problem occurs during the implementation, should also be included
- Any training required for personnel should be planned during this time
- Any contracts required should be negotiated during this time. Examples include outsourcing, Internet connectivity, maintenance etc.
- If all the above said points are in place the we can proceed with the implementation of the network

Monitoring and Redesigning Phase

- After the network is operating, baseline operational statistics should be gathered so that working status can be identified
- The network should then be monitored for anomalies and problems.
- If problems occurs, or if requirements change or are added, then appropriate design changes must be made and the entire design process should be repeated for that portion of the network.

Note

Monitoring and redesign take place in the PDIOO Operate and Optimize phases, and can lead back into the Plan and Design phases.

Maintaining Design Documentation

The design should be documented throughout the process. Documentation should include the following items

- All the agreed-to requirements and constraints
- The state of the existing network, if any
- Preliminary design options and a brief review of why the final design was chosen
- Final design details
- Results of any pilot or prototype testing
- Deployment plans, schedules, and other implementation details
- Monitoring requirements
- Any other pertinent information

A module is a component of a composite structure Modular network design involves creating modules that can then be put together to meet the requirements of the entire network.

A modular design for a network has many benefits, such as

- It is easier to understand and design smaller, simpler modules rather than an entire network
- It is easier to troubleshoot smaller elements compared to the entire network
- The reuse of blocks saves design time and effort, as well as implementation time and effort
- The reuse of blocks allows the network to grow more easily, providing network scalability
- It is easier to change modules rather than the entire network, providing flexibility of design.

TEXT/ REFERENCE BOOKS

1. Gil Held, "Network Design: Principles and Applications (Best Practices)", Auerbach Publications, 1st edition, 2000.
2. Diane Tiare and Catherine Paquet, "Campus Network Design Fundamentals", Pearson Education, 1st edition, 2006.
3. Larry L. Peterson, Bruce S. Davie, "Computer Networks: A Systems Approach", Morgan Kaufmann Publishers Inc., 5th edition, 2012.
4. William Stallings, "Data and Computer Communications", Pearson Education, 8th edition, 2016.
5. James F. Kurose, Keith W. Ross, "Computer Networking - A Top-Down Approach Featuring the Internet", Pearson Education, 6th edition, 2012.

PART A- 2 MARK QUESTIONS

S.NO	QUESTIONS	CO	LEVEL
1	Identify the design methodology used by network designers?	CO1	3
2	List the technical requirements of computer network.	CO1	4
3	Discuss the steps required to design a good network.	CO1	6
4	Examine various fundamental design goals in the network.	CO1	4
5	Identify various design principles applicable to network.	CO1	3
6	Summarize the benefits of PDIOO Cycle.	CO1	2
7	List the requirements related to business Issues in the computer network.	CO1	4
8	Show the PDIOO life cycle Diagram.	CO1	1
9	How refer design documents explain briefly.	CO1	1
10	List the Network design task and explain briefly.	CO1	1

PART B- 10 MARK QUESTIONS

S.NO	QUESTIONS	CO	LEVEL
1	Analyzing the existing network and explain briefly for each steps.	CO1	4
2	Discuss the preparing for preliminary design with neat diagrams.	CO1	6
3	Discuss the preparing for final design document, Monitoring, Designing and Maintaining Design Documentation.	CO1	6
4	Explain hierarchical network design with neat diagram and write the comparison of flat and hierarchical network.	CO1	5
5	Explain Modular Network Design with examples and also write the benefits of modular network design.	CO1	5

SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF COMPUTER SCIENCE AND ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE ENGINEERING

Computer Networks and Design – SCSA1502

UNIT II -UNDERLYING LAN CONCEPTS-SCSA1502

UNIT II
UNDERLYING LAN CONCEPTS

LAN connectivity for small businesses – Integration – Token-Ring – Ethernet – ATM LAN emulation – InterLAN Switching –LAN to Mainframe – Building networks.

NETWORKS

A network is a set of devices (often referred to as nodes) connected by communication links. A node can be a computer, printer, or any other device capable of sending and/or receiving data generated by other nodes on the network.

Distributed Processing

Most networks use distributed processing, in which a task is divided among multiple computers. Instead of one single large machine being responsible for all aspects of a process, separate computers (usually a personal computer or workstation) handle a subset.

Network Criteria

A network must be able to meet a certain number of criteria. The most important of these areperformance, reliability, and security.

Performance

Performance can be measured in many ways, including transit time and response time.Transit time is the amount of time required for a message to travel from one device to another. Response time is the elapsed time between an inquiry and a response. The performance of a network depends on a number of factors, including the number of users, the type of transmission medium, the capabilities of the connected hardware, and the efficiency of the software. Performance is often evaluated by two networking metrics throughput and delay. We often need more throughput and less delay. However, these two criteria are often contradictory. If we try to send more data to the network, we may increase throughput but we increase the delay because of traffic congestion in the network.

Reliability

In addition to accuracy of delivery, network reliability is measured by the frequency of failure, the time it takes a link to recover from a failure, and the network's robustness in a catastrophe.

Security

Network security issues include protecting data from unauthorized access, protecting data from damage and development, and implementing policies and procedures for recovery from breaches and data losses.

PHYSICAL STRUCTURES

Type of Connection

A network is two or more devices connected through links. A link is a communications pathway that transfers data from one device to another. For visualization purposes, it is simplest to imagine any link as a line drawn between two points. For communication to occur, two devices must be connected in some way to the same link at the same time. There are two possible types of connections point-to-point and multipoint.

Point-to-Point

A point-to-point connection provides a dedicated link between two devices. The entire capacity of the link is reserved for transmission between those two devices. Most point-to-point connections use an actual length of wire or cable to connect the two ends, but other options, such as microwave or satellite links, are also possible. When you change television channels by infrared remote control, you are establishing a point-to-point connection between the remote control and the television's control system.

Multipoint

A multipoint (also called multidrop) connection is one in which more than two specific devices share a single link. In a multipoint environment, the capacity of the channel is shared, either spatially or temporally. If several devices can use the link simultaneously, it is a spatially shared connection. If users must take turns, it is a timeshared connection.

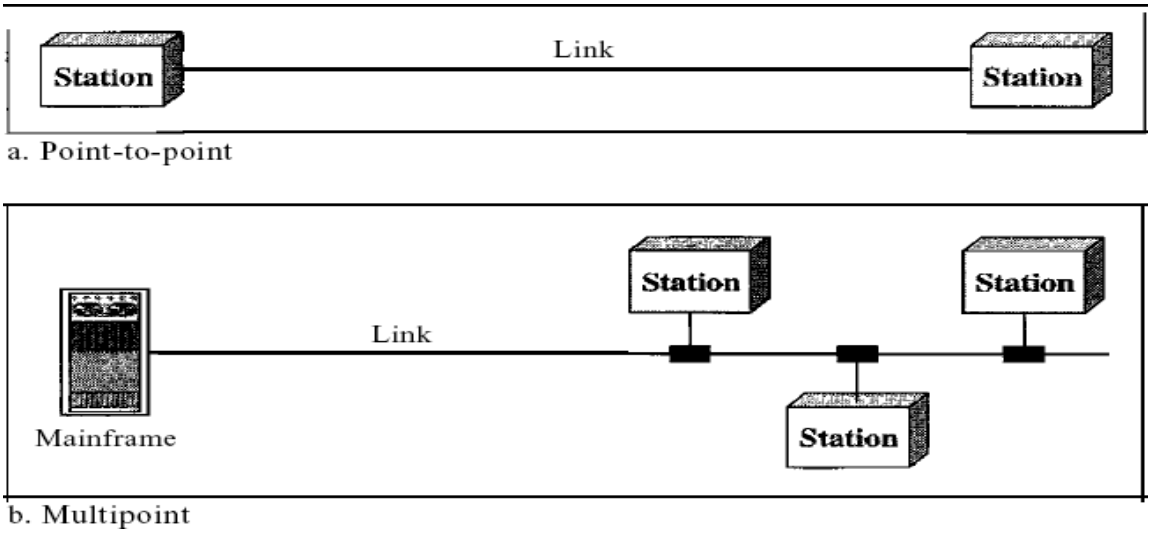

Figure 10 Types of connection

PHYSICAL TOPOLOGY

The term physical topology refers to the way in which a network is laid out physically. One or more devices connect to a link; two or more links form a topology. The topology of a network is the geometric representation of the relationship of all the links and linking devices (usually called nodes) to one another. There are four basic topologies possible mesh, star, bus, and ring

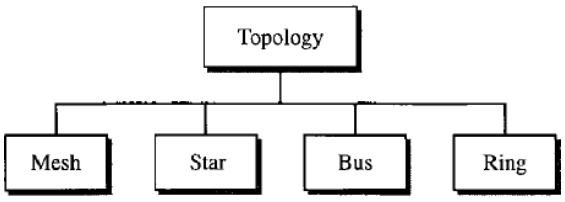

Figure 11 Topologies

Mesh

In a mesh topology, every device has a dedicated point-to-point link to every other device. The term dedicated means that the link carries traffic only between the two devices it connects. To find the number of physical links in a fully connected mesh network with n nodes, we first consider that each node must be connected to every other node. Node 1 must be connected to n - I nodes, node 2 must be connected to n – 1 nodes, and finally node n must be connected to n - 1 nodes. We need n(n - 1) physical links. However, if each physical link allows communication in both directions (duplex mode), we can divide the number of links by 2. In other words, we can say that in a mesh topology, we need n(n -1) /2 duplex-mode links.

To accommodate that many links, every device on the network must have n – 1 input/output (VO) ports to be connected to the other n - 1 stations.

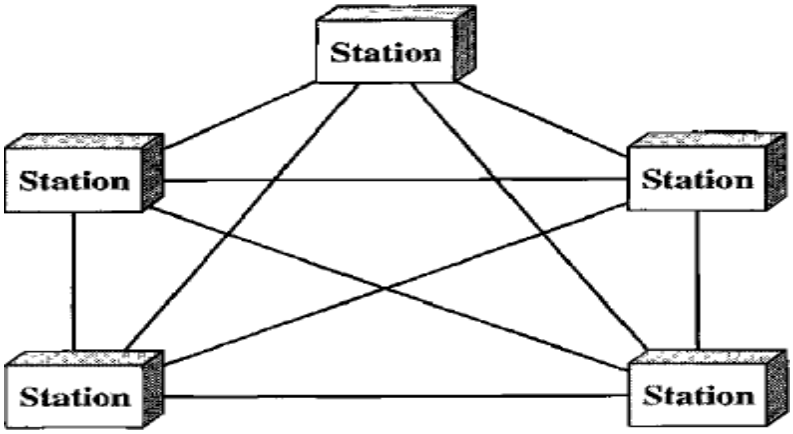

Figure 11 Mesh topology

Advantages

The use of dedicated links guarantees that each connection can carry its own data load, thus eliminating the traffic problems that can occur when links must be shared by multiple devices.

A mesh topology is robust. If one link becomes unusable, it does not incapacitate the entire system. Third, there is the advantage of privacy or security. When every message travels along a dedicated line, only the intended recipient sees it. Physical boundaries prevent other users from gaining access to messages. Finally, point-to-point links make fault identification and

fault isolation easy. Traffic can be routed to avoid links with suspected problems. This facility enables the network manager to discover the precise location of the fault and aids in finding its cause and solution.

Disadvantages

Disadvantage of a mesh are related to the amount of cabling because every device must be connected to every other device, installation and reconnection are difficult.
Second, the sheer bulk of the wiring can be greater than the available space (in walls, ceilings, or floors) can accommodate. Finally, the hardware required to connect each link (I/O ports and cable) can be prohibitively expensive.

For these reasons a mesh topology is usually implemented in a limited fashion, for example, as a backbone connecting the main computers of a hybrid network that can include several other topologies.

Star Topology

In a star topology, each device has a dedicated point-to-point link only to a central controller, usually called a hub. The devices are not directly linked to one another. Unlike a mesh topology, a star topology does not allow direct traffic between devices. The controller acts as an exchange If one device wants to send data to another, it sends the data to the controller, which then relays the data to the other connected device .

A star topology is less expensive than a mesh topology. In a star, each device needs only one linkand one I/O port to connect it to any number of others. This factor also makes it easy to install and reconfigure. Far less cabling needs to be housed, and additions, moves, and deletions involveonly one connection between that device and the hub.

Other advantages include robustness. If one link fails, only that link is affected. All other links remain active. This factor also lends itself to easy fault identification and fault isolation. As long as the hub is working, it can be used to monitor link problems and bypass defective links.

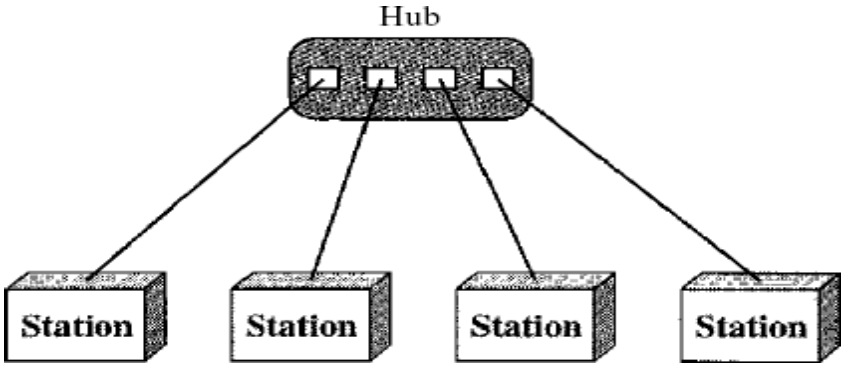

Figure 12 Star topology

One big disadvantage of a star topology is the dependency of the whole topology on one single point, the hub. If the hub goes down, the whole system is dead. Although a star requires far less cable than a mesh, each node must be linked to a central hub. For this reason, often more cabling is required in a star than in some other topologies (such as ring or bus).

Bus Topology

The preceding examples all describe point-to-point connections. A bus topology, on the other hand, is multipoint. One long cable acts as a backbone to link all the devices in a network

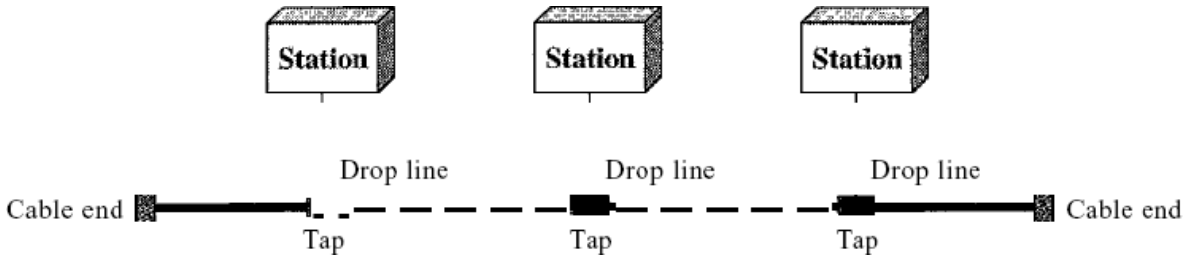

Figure 13 Bus topology

Nodes are connected to the bus cable by drop lines and taps. A drop line is a connection running between the device and the main cable. A tap is a connector that either splices into the main cable or punctures the sheathing of a cable to create a contact with the metallic core. As a signal travels along the backbone, some of its energy is transformed into heat. Therefore, it becomes weaker and weaker as it travels farther and farther. For this reason there is a limit on the number of taps a bus can support and on the distance between those taps.

Advantages of a bus topology include ease of installation. Backbone cable can be laid along the most efficient path, then connected to the nodes by drop lines of various lengths. In this way, a bus uses less cabling than mesh or star topologies. In a star, for example, four network devices inthe same room require four lengths of cable reaching all the way to the hub. In a bus, this redundancy is eliminated. Only the backbone cable stretches through the entire facility. Each drop line has to reach only as far as the nearest point on the backbone.

Disadvantages include difficult reconnection and fault isolation. A bus is usually designed to be optimally efficient at installation. It can therefore be difficult to add new devices. Signal reflection at the taps can cause degradation in quality. This degradation can be controlled by limiting the number and spacing of devices connected to a given length of cable. Adding new devices may therefore require modification or replacement of the backbone.

In addition, a fault or break in the bus cable stops all transmission, even between devices on the same side of the problem. The damaged area reflects signals back in the direction of origin, creating noise in both directions.

Bus topology was the one of the first topologies used in the design of early local area networks. Ethernet LANs can use a bus topology, but they are less popular.

RING TOPOLOGY

In a ring topology, each device has a dedicated point-to-point connection with only the two devices on either side of it. A signal is passed along the ring in one direction, from device to device, until it reaches its destination. Each device in the ring incorporates a repeater. When a device receives a signal intended for another device, its repeater regenerates the bits andpasses them along

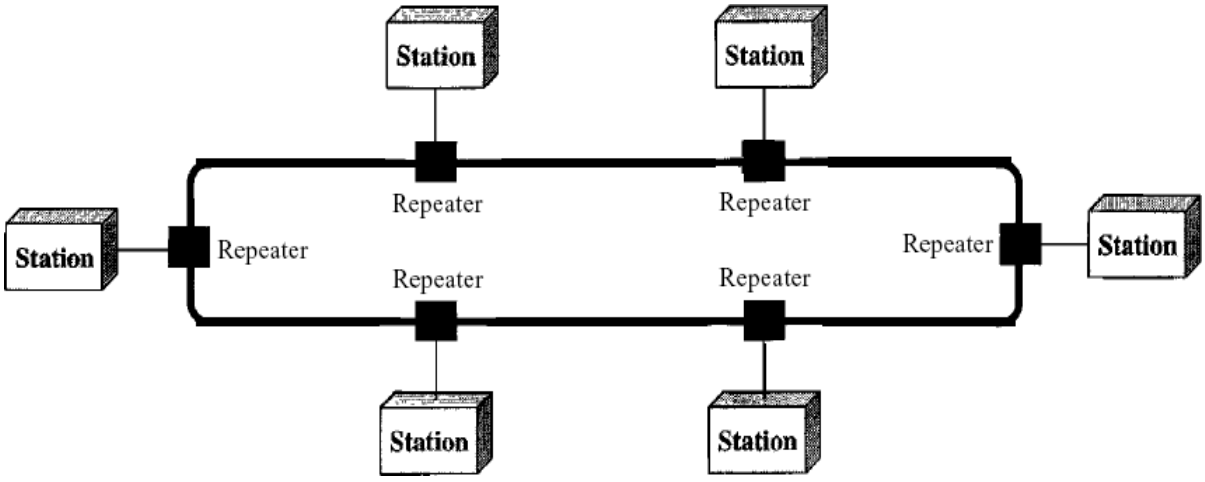

Figure14-1 Ring topology

A ring is relatively easy to install and reconfigure. Each device is linked to only its immediate neighbors (either physically or logically). To add or delete a device requires changing only two connections. The only constraints are media and traffic considerations (maximum ring length and number of devices). In addition, fault isolation is simplified. Generally in a ring, a signal is circulating at all times. If one device does not receive a signal within a specified period, it can issue an alarm. The alarm alerts the network operator to the problem and its location.

However, unidirectional traffic can be a disadvantage. In a simple ring, a break in the ring (such as a disabled station) can disable the entire network. This weakness can be solved by using a dual ring or a switch capable of closing off the break. Ring topology was prevalent when IBM introduced its local-area network Token Ring. Today, the need for higher-speed LANs has made this topology less popular. Hybrid Topology A network can be hybrid. For example, we can have a main star topology with each branch connecting several stations in a bus topology as shown in Figure

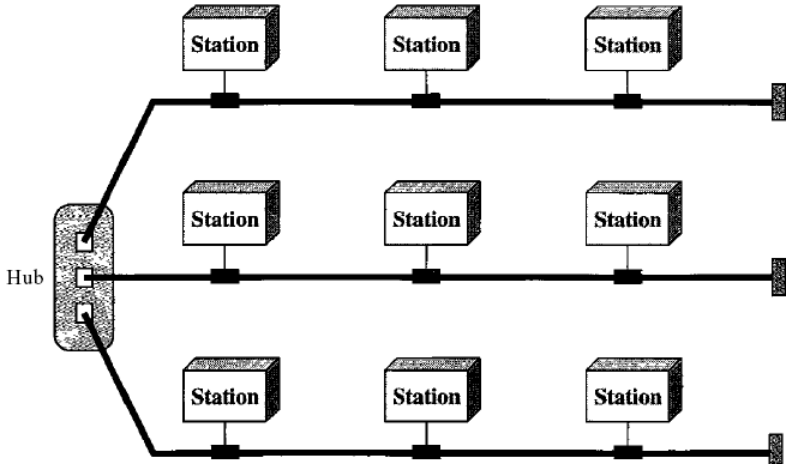

Figure 15 Hybrid Topology

TRANSMISSION MEDIA

For any networking to be effective, raw stream of data is to be transported from one device to other over some medium. Various transmission media can be used for transfer of data. These transmission media may be of two types −

- Guided − In guided media, transmitted data travels through cabling system that has a fixed path. For example, copper wires, fibre optic wires, etc.
- Unguided − In unguided media, transmitted data travels through free space in form of electromagnetic signal. For example, radio waves, lasers, etc.
- Each transmission media has its own advantages and disadvantages in terms of bandwidth, speed, delay, cost per bit, ease of installation and maintenance, etc. Let's discuss some of the most commonly used media in detail.

Twisted Pair Cable

Copper wires are the most common wires used for transmitting signals because of good performance at low costs. They are most commonly used in telephone lines. However, if two or more wires are lying together, they can interfere with each other's signals. To reduce this electromagnetic interference, pair of copper wires are twisted together in helical shape like a DNA molecule. Such twisted copper wires are called twisted pair. To reduce interference between nearby twisted pairs, the twist rates are different for each pair.

Up to 25 twisted pair are put together in a protective covering to form twisted pair cables that are the backbone of telephone systems and Ethernet networks.

Advantages of twisted pair cable

Twisted pair cable are the oldest and most popular cables all over the world. This is due to the many advantages that they offer −

- Trained personnel easily available due to shallow learning curve
- Can be used for both analog and digital transmissions
- Least expensive for short distances
- Entire network does not go down if a part of network is damaged

Disadvantages of twisted pair cable

With its many advantages, twisted pair cables offer some disadvantages too −

- Signal cannot travel long distances without repeaters
- High error rate for distances greater than 100m
- Very thin and hence breaks easily
- Not suitable for broadband connections

Shielding twisted pair cable

To counter the tendency of twisted pair cables to pick up noise signals, wires are shielded in the following three ways

- Each twisted pair is shielded.
- Set of multiple twisted pairs in the cable is shielded.
- Each twisted pair and then all the pairs are shielded.

Such twisted pairs are called shielded twisted pair (STP) cables. The wires that are not shielded but simply bundled together in a protective sheath are called unshielded twisted pair (UTP) cables. These cables can have maximum length of 100 metres.

Shielding makes the cable bulky, so UTP are more popular than STP. UTP cables are used as the last mile network connection in homes and offices.

Coaxial Cable

Coaxial cables are copper cables with better shielding than twisted pair cables, so that transmitted signals may travel longer distances at higher speeds. A coaxial cable consists of these layers, starting from the innermost −

- Stiff copper wire as core
- Insulating material surrounding the core
- Closely woven braided mesh of conducting material surrounding the insulator
- Protective plastic sheath encasing the wire

Coaxial cables are widely used for cable TV connections and LANs.

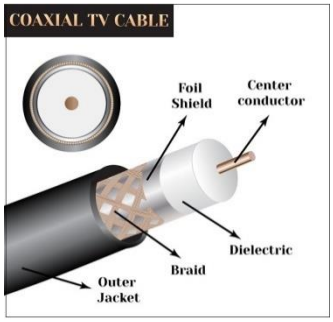

Figure 16 Coaxial cable

Advantages of Coaxial Cables

These are the advantages of coaxial cables −

- Excellent noise immunity
- Signals can travel longer distances at higher speeds, e.g. 1 to 2 Gbps for 1 Km cable
- Can be used for both analog and digital signals
- Inexpensive as compared to fibre optic cables
- Easy to install and maintain

Disadvantages of Coaxial Cables

These are some of the disadvantages of coaxial cables −

- Expensive as compared to twisted pair cables
- Not compatible with twisted pair cables

Optical fibre

Thin glass or plastic threads used to transmit data using light waves are called optical fibre. Light Emitting Diodes (LEDs) or Laser Diodes (LDs) emit light waves at the source, which is read by a detector at the other end. Optical fibre cable has a bundle of such threads or fibres bundled together in a protective covering. Each fibre is made up of these three layers, starting with the innermost layer −

- Core made of high quality silica glass or plastic
- Cladding made of high quality silica glass or plastic, with a lower refractive index than the core
- Protective outer covering called buffer

Note that both core and cladding are made of similar material. However, as refractive index of the cladding is lower, any stray light wave trying to escape the core is reflected back due to total internal reflection.

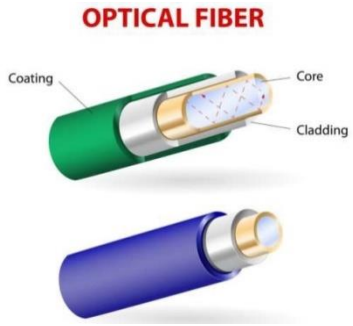

Figure 17 Optical Fibre

Optical fibre is rapidly replacing copper wires in telephone lines, internet communication and even cable TV connections because transmitted data can travel very long distances without weakening. Single node fibre optic cable can have maximum segment length of 2 kms and bandwidth of up to 100 Mbps. Multi-node fibre optic cable can have maximum segment length of 100 kms and bandwidth up to 2 Gbps.

Advantages of Optical Fibre

Optical fibre is fast replacing copper wires because of these advantages that it offers −

- High bandwidth
- Immune to electromagnetic interference
- Suitable for industrial and noisy areas
- Signals carrying data can travel long distances without weakening

Disadvantages of Optical Fibre

- Despite long segment lengths and high bandwidth, using optical fibre may not be a viable option for every one due to these disadvantages −
- Optical fibre cables are expensive
- Sophisticated technology required for manufacturing, installing and maintaining optical fibre cables
- Light waves are unidirectional, so two frequencies are required for full duplex transmission

Infrared

- Low frequency infrared waves are used for very short distance communication like TV remote, wireless speakers, automatic doors, hand held devices etc. Infrared signals can propagate within a room but cannot penetrate walls. However, due to such short range, it is considered to be one of the most secure transmission modes.

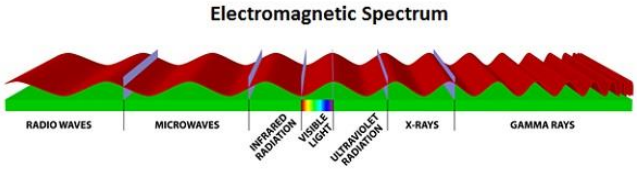

Radio Wave

Transmission of data using radio frequencies is called radio-wave transmission. We all are familiar with radio channels that broadcast entertainment programs. Radio stations transmit radio waves using transmitters, which are received by the receiver installed in our devices.

Both transmitters and receivers use antennas to radiate or capture radio signals. These radio frequencies can also be used for direct voice communication within the allocated range. This range is usually 10 miles.

Figure 18 Radiowave

Advantages of Radio Wave

These are some of the advantages of radio wave transmissions −

- Inexpensive mode of information exchange
- No land needs to be acquired for laying cables

- Installation and maintenance of devices is cheap
- Disadvantages of Radio Wave

These are some of the disadvantages of radio wave transmissions −

- Insecure communication medium
- Prone to weather changes like rain, thunderstorms, etc

CATEGORIES OF NETWORKS

Local Area Networks

Local area networks, generally called LANs, are privately-owned networks within a single building or campus of up to a few kilometres in size. They are widely used to connect personal computers and workstations in company offices and factories to share resources (e.g., printers) and exchange information. LANs are distinguished from other kinds of networks by threecharacteristics

- Their size,
- Their transmission technology, and
- Their topology.

LANs are restricted in size, which means that the worst-case transmission time is bounded and known in advance. Knowing this bound makes it possible to use certain kinds of designs that would not otherwise be possible. It also simplifies network management. LANs may use a transmission technology consisting of a cable to which all the machines are attached, like the telephone company party lines once used in rural areas. Traditional LANs run at speeds of 10 Mbps to 100 Mbps, have low delay (microseconds or nanoseconds), and make very few errors. Newer LANs operate at up to 10 Gbps Various topologies are possible for broadcast LANs. Figure1 shows two of them. In a bus (i.e., a linear cable) network, at any instant at most one machine is the master and is allowed to transmit. All other machines are required to refrain from sending. An arbitration mechanism is needed to resolve conflicts when two or more machines want to transmit simultaneously. The arbitration mechanism may be centralized or distributed. IEEE 802.3, popularly called Ethernet, for example, is a bus-based broadcast network with decentralized control, usually operating at 10 Mbps to 10 Gbps. Computers on an Ethernet can transmit whenever they want to; if two or more packets collide, each computer just waits a random time and tries again later.

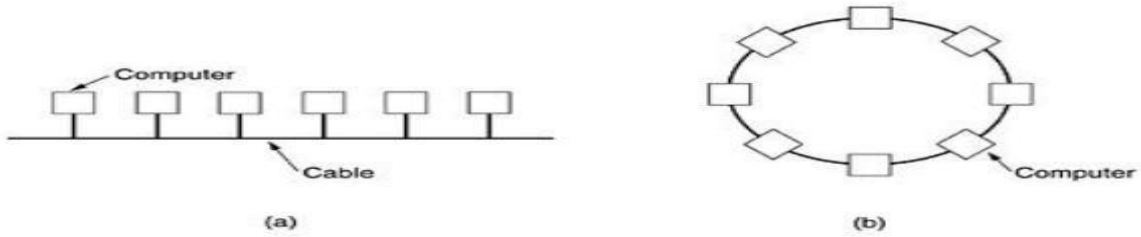

Figure.19 Two broadcast networks . (a) Bus. (b) Ring.

A second type of broadcast system is the ring. In a ring, each bit propagates around on its own, not waiting for the rest of the packet to which it belongs. Typically, each bit

circumnavigates the entire ring in the time it takes to transmit a few bits, often before the complete packet has even been transmitted. As with all other broadcast systems, some rule is needed for arbitrating simultaneous accesses to the ring. Various methods, such as having the machines take turns, are in use. IEEE 802.5 (the IBM token ring), is a ring-based LAN operating at 4 and 16 Mbps. FDDIis another example of a ring network.

Metropolitan Area Network (MAN)

A metropolitan area network, or MAN, covers a city. The best-known example of a MAN is the cable television network available in many cities. This system grew from earlier community antenna systems used in areas with poor over-the-air television reception. In these early systems, a large antenna was placed on top of a nearby hill and signal was then piped to the subscribers' houses. At first, these were locally-designed, ad hoc systems. Then companies began jumping into the business, getting contracts from city governments to wire up an entire city. The next step was television programming and even entire channels designed for cable only. Often these channels were highly specialized, such as all news, all sports, all cooking, all gardening, and so on. But from their inception until the late 1990s, they were intended for television reception only.To a first approximation, a MAN might look something like the system shown in Fig. In this figure both television signals and Internet are fed into the centralized head end for subsequent distribution to people's homes. Cable television is not the only MAN. Recent developments in high-speed wireless Internet access resulted in another MAN, which has been standardized as IEEE 802.16.

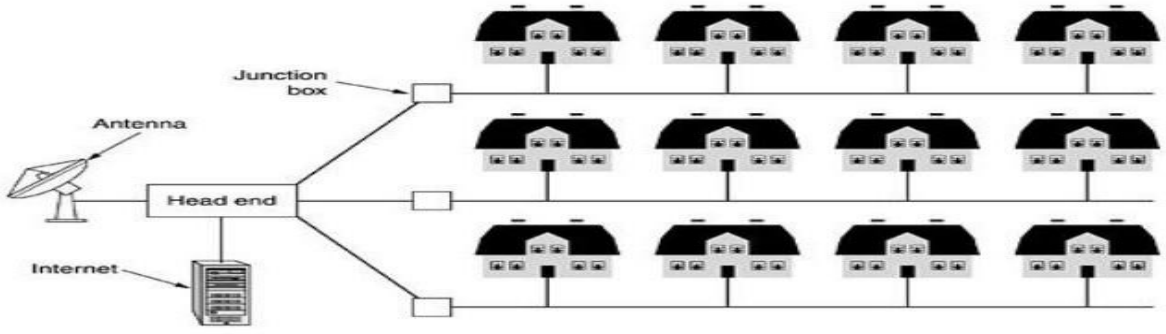

Figure 20 Metropolitan Area Network

A MAN is implemented by a standard called DQDB (Distributed Queue Dual Bus) or IEEE 802.16. DQDB has two unidirectional buses (or cables) to which all the computers are attached.

WIDE AREA NETWORK

A wide area network, or WAN, spans a large geographical area, often a country or continent. It contains a collection of machines intended for running user (i.e., application) programs. These machines are called as hosts. The hosts are connected by a communication subnet, or just subnet for short. The hosts are owned by the customers (e.g., people's personal computers), whereas the communication subnet is typically owned and operated by a telephone company or Internet service provider. The job of the subnet is to carry messages from host to host, just as the telephone system carries words from speaker to listener.

Separation of the pure communication aspects of the network (the subnet) from the application aspects (the hosts), greatly simplifies the complete network design. In most wide area networks, the subnet consists of two distinct components transmission lines and switching elements. Transmission lines move bits between machines. They can be made of copper wire, optical fiber, or even radio links. In most WANs, the network contains numerous transmission lines, each one connecting a pair of routers. If two routers that do not share a transmission line wish to communicate, they must do this indirectly, via other routers. When a packet is sent from one router to another via one or more intermediate routers, the packet is received at each intermediate router in its entirety, stored there until the required output line is free, and then forwarded. A subnet organized according to this principle is called a store-and-forward or packet-switched subnet. Nearly all wide area networks (except those using satellites) have store-and-forward subnets. When the packets are small and all the same size, they are often called cells.

The principle of a packet-switched WAN is so important. Generally, when a process on some host has a message to be sent to a process on some other host, the sending host first cuts the message into packets, each one bearing its number in the sequence. These packets are then injected into the network one at a time in quick succession. The packets are transported individually over the network and deposited at the receiving host, where they are reassembled into the original message and delivered to the receiving process. A stream of packets resulting from some initial message is illustrated in Fig

In this figure, all the packets follow the route ACE, rather than ABDE or ACDE. In some networks all packets from a given message must follow the same route; in others each packed is routed separately. Of course, if ACE is the best route, all packets may be sent along it, even if each packet is individually routed.

Not all WANs are packet switched. A second possibility for a WAN is a satellite system. Each router has an antenna through which it can send and receive. All routers can hear the output from the satellite, and in some cases they can also hear the upward transmissions of their fellow routers to the satellite as well. Sometimes the routers are connected to a substantial point-to-point subnet, with only some of them having a satellite antenna. Satellite networks are inherently broadcast and are most useful when the broadcast property is important.

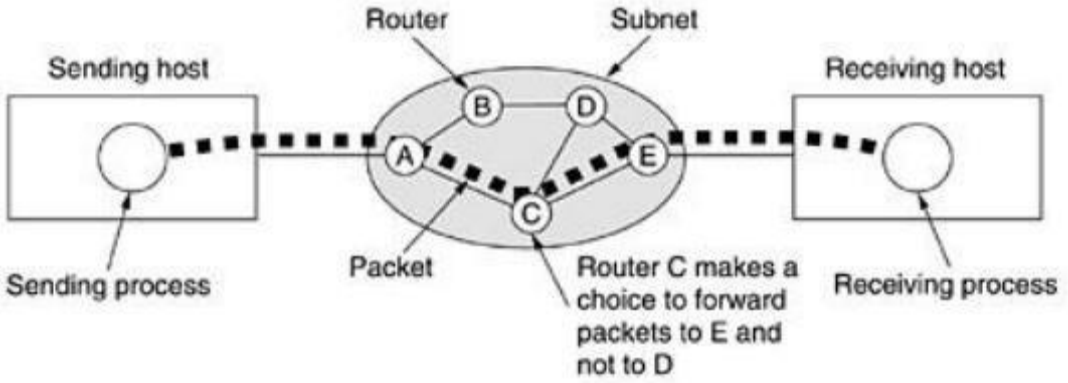

Figure 20 A stream of packets from sender to receiver.

THE INTERNET

The Internet has revolutionized many aspects of our daily lives. It has affected the way we do business as well as the way we spend our leisure time. Count the ways you've used the Internet recently. Perhaps you've sent electronic mail (e-mail) to a business associate, paid a utility bill, read a newspaper from a distant city, or looked up a local movie schedule-all by using the Internet. Or maybe you researched a medical topic, booked a hotel reservation, chatted with a fellow Trekkie, or comparison-shopped for a car. The Internet is a communication .

A Brief History

A network is a group of connected communicating devices such as computers and printers. An internet (note the lowercase letter i) is two or more networks that can communicate with each other. The most notable internet is called the Internet (uppercase letter I), a collaboration of morethan hundreds of thousands of interconnected networks. Private individuals as well as various organizations such as government agencies, schools, research facilities, corporations, and libraries in more than 100 countries use the Internet. Millions of people are users. Yet this extraordinary communication system only came into being in 1969.

In the mid-1960s, mainframe computers in research organizations were standalone devices.Computers from different manufacturers were unable to communicate with one another. The Advanced Research Projects Agency (ARPA) in the Department of Defense (DoD) was interested in finding a way to connect computers so that the researchers they funded could share their findings, thereby reducing costs and eliminating duplication of effort.

In 1967, at an Association for Computing Machinery (ACM) meeting, ARPA presented its ideas for ARPANET, a small network of connected computers. The idea was that each host computer (not necessarily from the same manufacturer) would be attached to a specialized computer, called an inteiface message processor (IMP). The IMPs, in tum, would be connected to one another. Each IMP had to be able to communicate with other IMPs as well as with its own attached host. By 1969, ARPANET was a reality. Four nodes, at the University of California at Los Angeles (UCLA), the University of California at Santa Barbara (UCSB), Stanford Research Institute (SRI), and the University of Utah, were connected via the IMPs to form a network.

Software called the Network Control Protocol (NCP) provided communication between the hosts.

In 1972, Vint Cerf and Bob Kahn, both of whom were part of the core ARPANET group, collaborated on what they called the Internetting Projec1. Cerf and Kahn's landmark 1973 paper outlined the protocols to achieve end-to-end delivery of packets. This paper on Transmission Control Protocol (TCP) included concepts such as encapsulation, the datagram, and the functions of a gateway. Shortly thereafter, authorities made a decision to split TCP into two protocols Transmission Control Protocol (TCP) and Internetworking Protocol (lP). IP would handle datagram routing while TCP would be responsible for higher-level functions such as segmentation, reassembly, and error detection. The internetworking protocol became known as TCPIIP.

The Internet Today

The Internet has come a long way since the 1960s. The Internet today is not a simple hierarchicalstructure. It is made up of many wide- and local-area networks joined by connecting devices andswitching stations. It is difficult to give an accurate representation of the Internet because it iscontinually changing-new networks are being added, existing networks are adding addresses, andnetworks of defunct companies are being removed. Today most end users who want Internet connection use the services of Internet service providers (lSPs). There are international serviceproviders, national service providers, regional service providers, and local service providers. The Internet today is run by private companies, not the government. Figure 1.13 shows a conceptual (not geographic) view of the Internet.

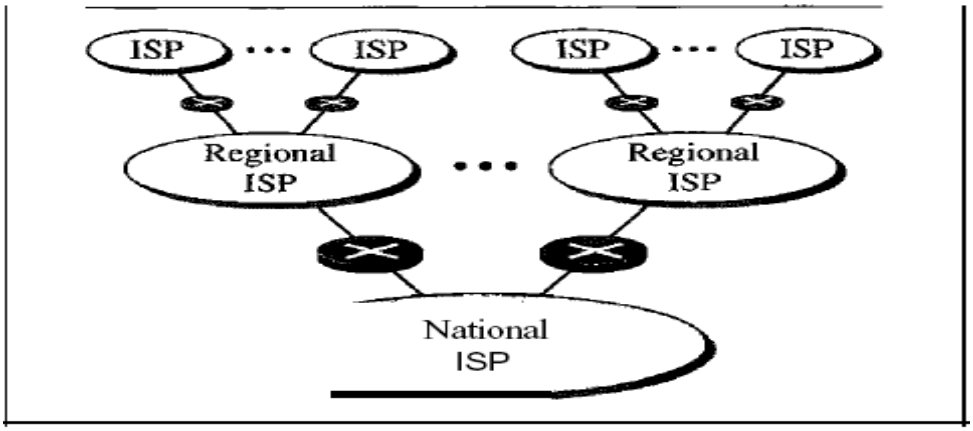

a. Structure of a national ISP

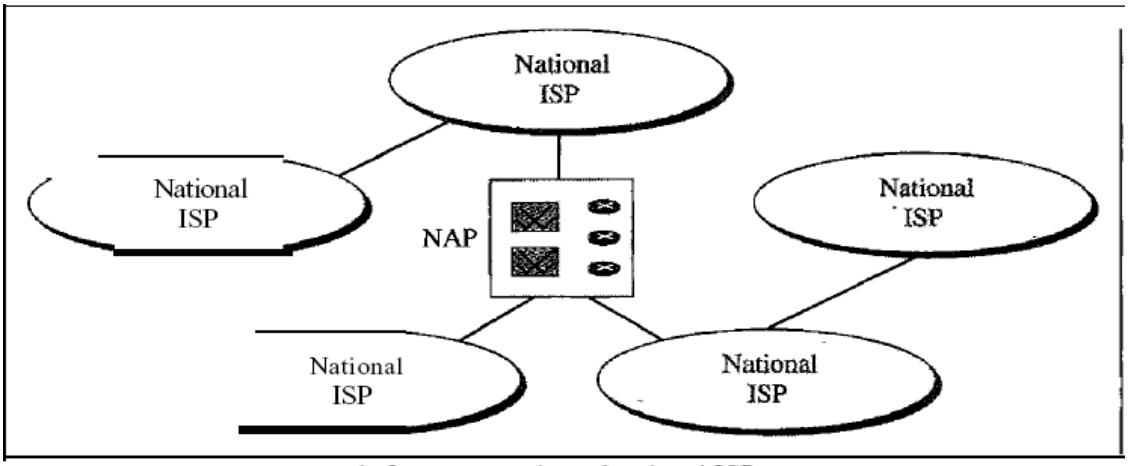

b. Interconnection of national ISPs

International Internet Service Providers

At the top of the hierarchy are the international service providers that connect nations together.

National Internet Service Providers

The national Internet service providers are backbone networks created and maintained by specialized companies. There are many national ISPs operating in North America; some of the most well known are SprintLink, PSINet, UUNet Technology, AGIS, and internet Mel. To provide connectivity between the end users, these backbone networks are connected by complex switching stations (normally run by a third party) called network access points (NAPs). Some national ISP networks are also connected to one another by private switching stations called peering points. These normally operate at a high data rate (up to 600 Mbps).

Regional Internet Service Providers

Regional internet service providers or regional ISPs are smaller ISPs that are connected to one or more national ISPs. They are at the third level of the hierarchy with a smaller data rate. Local Internet Service Providers

Local Internet service providers provide direct service to the end users. The local ISPs can be connected to regional ISPs or directly to national ISPs. Most end users are connected to the local ISPs. Note that in this sense, a local ISP can be a company that just provides Internet services, a corporation with a network that supplies services to its own employees, or a nonprofit organization, such as a college or a university, that runs its own network. Each of these local ISPs can be connected to a regional or national service provider.

PROTOCOLS AND STANDARDS

In computer networks, communication occurs between entities in different systems. An entity is anything capable of sending or receiving information. However, two entities cannot simply send bit streams to each other and expect to be understood. For communication to occur, the entities must agree on a protocol. A protocol is a set of rules that govern data communications. A protocol defines what is communicated, how it is communicated, and when it is communicated. The key elements of a protocol are syntax, semantics, and timing.

- Syntax. The term syntax refers to the structure or format of the data, meaning the order in which they are presented. For example, a simple protocol might expect the first 8 bits of datato be the address of the sender, the second 8 bits to be the address of the receiver, and the rest of the stream to be the message itself.
- Semantics. The word semantics refers to the meaning of each section of bits. How is a particular pattern to be interpreted, and what action is to be taken based on that interpretation? For example, does an address identify the route to be taken or the final destination of the message?
- Timing. The term timing refers to two characteristics when data should be sent and how fast they can be sent. For example, if a sender produces data at 100 Mbps but the receiver can process data at only 1 Mbps, the transmission will overload the receiver and some data will

be lost.

Standards

Standards are essential in creating and maintaining an open and competitive market for equipment manufacturers and in guaranteeing national and international interoperability of data and telecommunications technology and processes. Standards provide guidelines to manufacturers, vendors, government agencies, and other service providers to ensure the kind of interconnectivity necessary in today's marketplace and in international communications.

Data communication standards fall into two categories de facto (meaning "by fact" or "by convention") and de jure (meaning "by law" or "by regulation").

De facto. Standards that have not been approved by an organized body but have been adopted as standards through widespread use are de facto standards. De facto standards are often established originally by manufacturers who seek to define the functionality of a new product or technology. Those standards that have been legislated by an officially recognized body are de jure standards.

LAYERED TASKS

We use the concept of layers in our daily life. As an example, let us consider two friends who communicate through postal maiL The process of sending a letter to a friend would be complex if there were no services available from the post office. Below Figure shows the steps in this task.

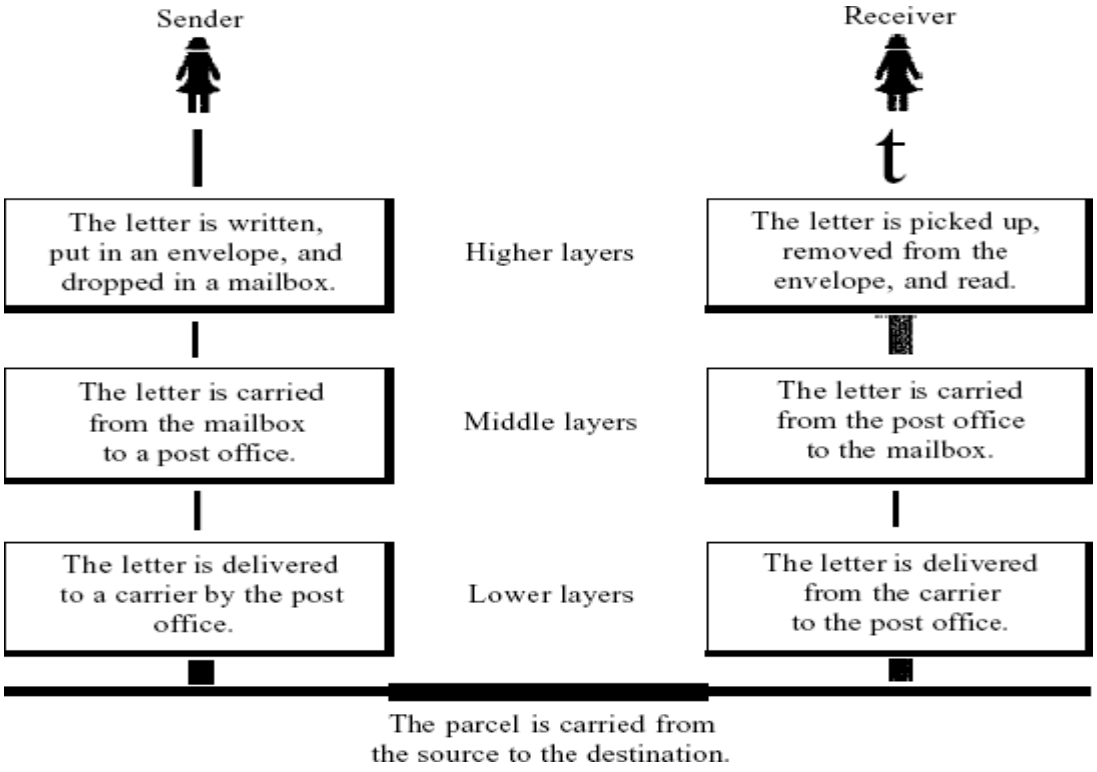

Figure 21 Sender, Receiver, and Carrier

In Figure we have a sender, a receiver, and a carrier that transports the letter. There is a hierarchyof tasks.

At the Sender Site

Let us first describe, in order, the activities that take place at the sender site.

Higher layer. The sender writes the letter, inserts the letter in an envelope, writes thesender and receiver addresses, and drops the letter in a mailbox.

Middle layer. The letter is picked up by a letter carrier and delivered to the post office.

Lower layer. The letter is sorted at the post office; a carrier transports the letter.

0n the Way The letter is then on its way to the recipient. On the way to the recipient's local post office, the letter may actually go through a central office. In addition, it may be transported by truck, train, airplane, boat, or a combination of things

At the Receiver Site

Lower layer. The carrier transports the letter to the post office.

Middle layer. The letter is sorted and delivered to the recipient's mailbox.

Higher layer. The receiver picks up the letter, opens the envelope, and reads it.

The OSI Reference Model

The OSI model (minus the physical medium) is shown in Fig. This model is based on a proposal developed by the International Standards Organization (ISO) as a first step toward international standardization of the protocols used in the various layers (Day and Zimmermann, 1983). It was revised in 1995(Day, 1995). The model is called the ISO-OSI (Open Systems Interconnection) Reference Model because it deals with connecting open systems—that is, systems that are open for communication with other systems.

The OSI model has seven layers. The principles that were applied to arrive at the seven layers can be briefly summarized as follows

- A layer should be created where a different abstraction is needed.
- Each layer should perform a well-defined function.
- The function of each layer should be chosen with an eye toward defining internationally standardized protocols.
- The layer boundaries should be chosen to minimize the information flow across the interfaces.
- The number of layers should be large enough that distinct functions need not be thrown together in the same layer out of necessity and small enough that the architecture does not become unwieldy.

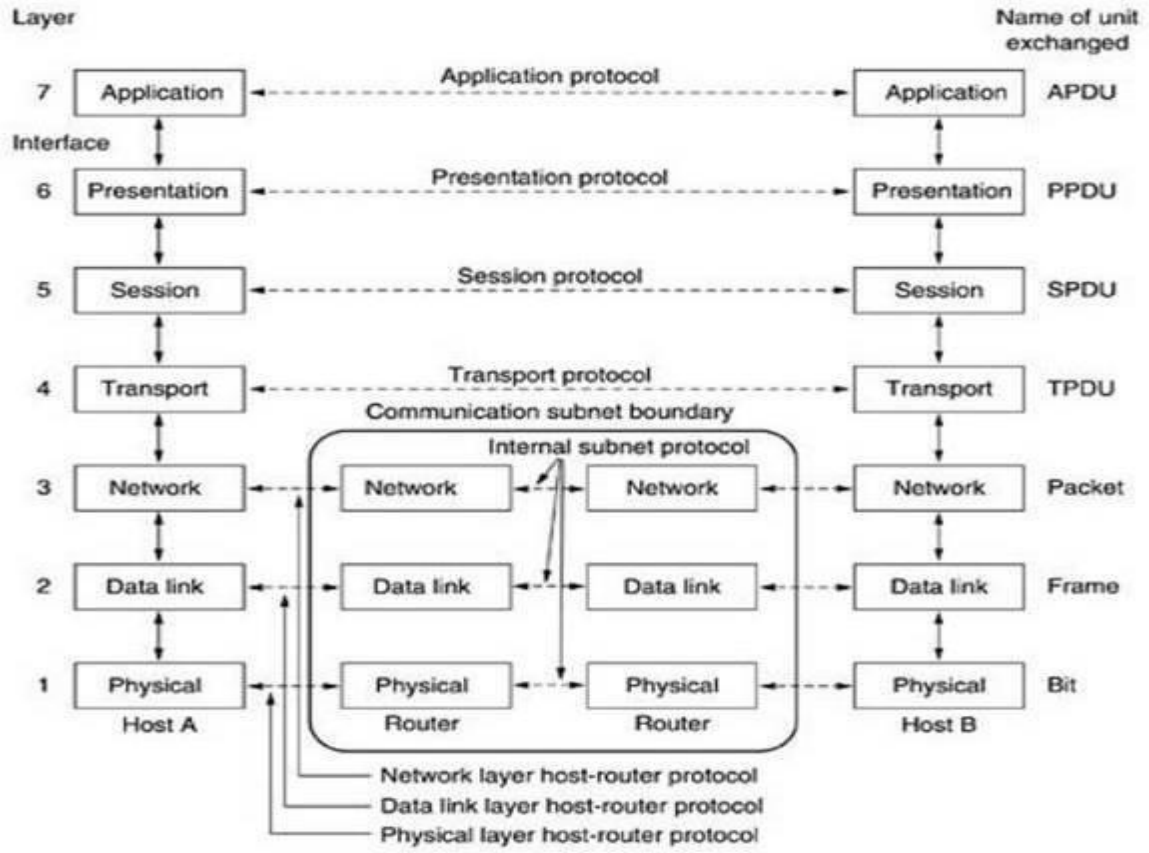

The Physical Layer

The physical layer is concerned with transmitting raw bits over a communication channel. The design issues have to do with making sure that when one side sends a 1 bit, it is received by the other side as a 1 bit, not as a 0 bit.

The Data Link Layer

The main task of the data link layer is to transform a raw transmission facility into a line that appears free of undetected transmission errors to the network layer. It accomplishes this task by having the sender break up the input data into data frames (typically a few hundred or a few thousand bytes) and transmits the frames sequentially. If the service is reliable, the receiver confirms correct receipt of each frame by sending back an acknowledgement frame.

Another issue that arises in the data link layer (and most of the higher layers as well) is how to keep a fast transmitter from drowning a slow receiver in data. Some traffic regulation mechanismis often needed to let the transmitter know how much buffer space the receiver has at the moment. Frequently, this flow regulation and the error handling are integrated.

The Network Layer

The network layer controls the operation of the subnet. A key design issue is determining how packets are routed from source to destination. Routes can be based on static tables that are "wired into" the network and rarely changed. They can also be determined at the start of each conversation, for example, a terminal session (e.g., a login to a remote machine). Finally, they

can be highly dynamic, being determined anew for each packet, to reflect the current network load.

If too many packets are present in the subnet at the same time, they will get in one another's way, forming bottlenecks. The control of such congestion also belongs to the network layer. More generally, the quality of service provided (delay, transit time, jitter, etc.) is also a network layer issue.

When a packet has to travel from one network to another to get to its destination, many problems can arise. The addressing used by the second network may be different from the first one. The second one may not accept the packet at all because it is too large. The protocols may differ, and so on. It is up to the network layer to overcome all these problems to allow heterogeneous networks to be interconnected. In broadcast networks, the routing problem is simple, so the network layer is often thin or even nonexistent.

The Transport Layer

The basic function of the transport layer is to accept data from above, split it up into smaller units if need be, pass these to the network layer, and ensure that the pieces all arrive correctly at the other end. Furthermore, all this must be done efficiently and in a way that isolates the upper layers from the inevitable changes in the hardware technology. The transport layer also determines what type of service to provide to the session layer, and, ultimately, to the users of the network. The most popular type of transport connection is an error-free point-to-point channel that delivers messages or bytes in the order in which they were sent. However, other possible kinds of transport service are the transporting of isolated messages, with no guarantee about the order of delivery, and the broadcasting of messages to multiple destinations. The type of service is determined when the connection is established.

The transport layer is a true end-to-end layer, all the way from the source to the destination. In other words, a program on the source machine carries on a conversation with a similar program on the destination machine, using the message headers and control messages. In the lower layer the protocols are between each machine and its immediate neighbours, and not between the ultimate source and destination machines, which may be separated by many routers.

The Session Layer

The session layer allows users on different machines to establish sessions between them. Sessions offer various services, including dialog control (keeping track of whose turn it is to transmit), token management (preventing two parties from attempting the same critical operation at the same time), and synchronization (check pointing long transmissions to allow them to continue from where they were after a crash).

The Presentation Layer

The presentation layer is concerned with the syntax and semantics of the information transmitted. In order to make it possible for computers with different data representations to communicate, the data structures to be exchanged can be defined in an abstract way, along with a standard encoding to be used "on the wire." The presentation layer manages these abstract data structures and allows higher-level data structures (e.g., banking records), to be defined and exchanged.

The Application Layer

The application layer contains a variety of protocols that are commonly needed by users. One widely-used application protocol is HTTP (Hypertext Transfer Protocol), which is the basis

for the World Wide Web. When a browser wants a Web page, it sends the name of the page it wants to the server using HTTP. The server then sends the page back. Other application protocols are used for file transfer, electronic mail, and network news.

The TCP/IP Reference Model

The TCP/IP reference model was developed prior to OSI model. The major design goals of thismodel were,

- To connect multiple networks together so that they appear as a single network.
- To survive after partial subnet hardware failures.
- To provide a flexible architecture.

Unlike OSI reference model, TCP/IP reference model has only 4 layers. They are,

- Host-to-Network Layer
- Internet Layer
- Transport Layer
- Application LayerApplication Layer Transport Layer Internet Layer
- Host-to-Network Layer

Host-to-Network Layer

The TCP/IP reference model does not really say much about what happens here, except to point out that the host has to connect to the network using some protocol so it can send IP packets to it.This protocol is not defined and varies from host to host and network to network.

Internet Layer

This layer, called the internet layer, is the linchpin that holds the whole architecture together. Its job is to permit hosts to inject packets into any network and have they travel independently to thedestination (potentially on a different network). They may even arrive in a different order than they were sent, in which case it is the job of higher layers to rearrange them, if in-order delivery is desired. Note that "internet" is used here in a generic sense, even though this layer is present in the Internet.

The internet layer defines an official packet format and protocol called IP (Internet Protocol). The job of the internet layer is to deliver IP packets where they are supposed to go. Packet routing is clearly the major issue here, as is avoiding congestion. For these reasons, it is reasonable to say that the TCP/IP internet layer is similar in functionality to the OSI network layer. Fig. shows this correspondence.

The Transport Layer

The layer above the internet layer in the TCP/IP model is now usually called the transport layer.It is designed to allow peer entities on the source and destination hosts to carry on a conversation, just as in the OSI transport layer. Two end-to-end transport protocols have been defined here. The first one, TCP (Transmission Control Protocol), is a reliable connection-oriented protocol that allows a byte stream originating on one machine to be delivered without error on any other machine in the internet. It fragments the incoming byte stream into discrete messages and passes each one on to the internet layer. At the destination, the receiving TCP

process reassembles the received messages into the output stream. TCP also handles flow control to make sure a fast sender cannot swamp a slow receiver with more messages than it can handle.

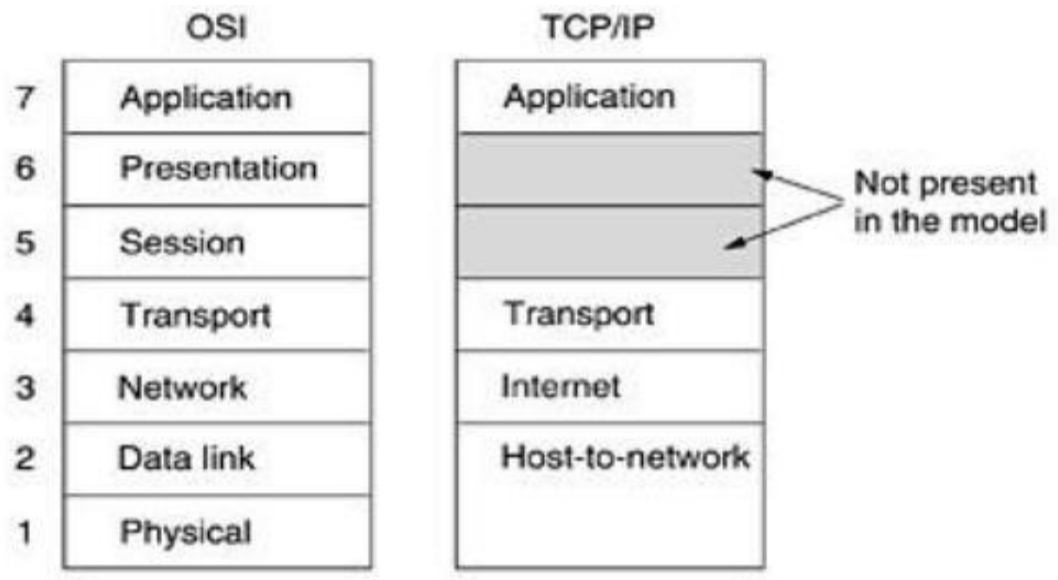

Figure 22 The TCP/IP reference model.

The second protocol in this layer, UDP (User Datagram Protocol), is an unreliable, connectionless protocol for applications that do not want TCP's sequencing or flow control and wish to provide their own. It is also widely used for one-shot, client-server-type request-reply queries and applications in which prompt delivery is more important than accurate delivery, such as transmitting speech or video. The relation of IP, TCP, and UDP is shown in Fig.2. Since the model was developed, IP has been implemented on many other networks.

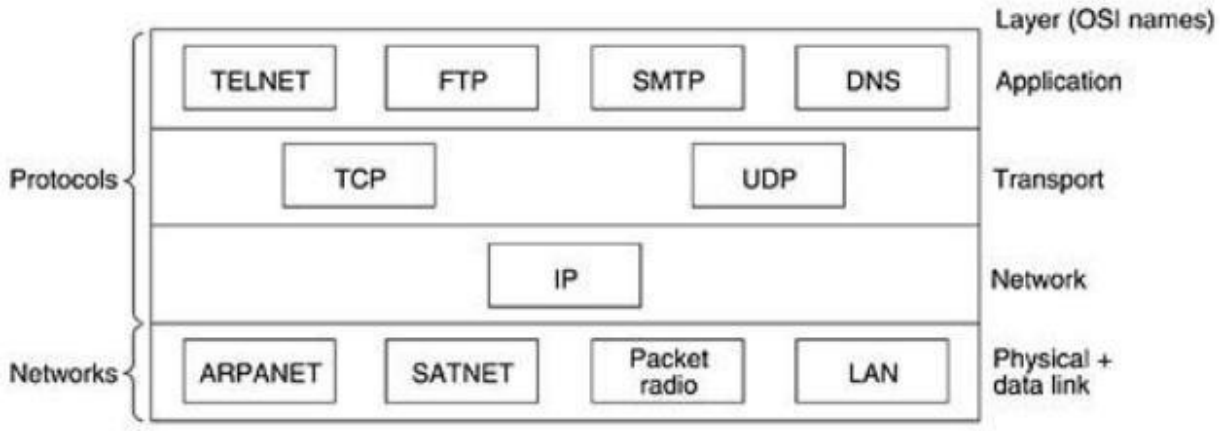

Figure 23 Protocols and networks in the TCP/IP model initially.

The Application Layer

The TCP/IP model does not have session or presentation layers. On top of the transport

layer is the application layer. It contains all the higher-level protocols. The early ones included virtual terminal (TELNET), file transfer (FTP), and electronic mail (SMTP), as shown in Fig.6.2. The virtual terminal protocol allows a user on one machine to log onto a distant machine and work there. The file transfer protocol provides a way to move data efficiently from one machine to another. Electronic mail was originally just a kind of file transfer, but later a specialized protocol (SMTP) was developed for it. Many other protocols have been added to these over the years the Domain Name System (DNS) for mapping host names onto their network addresses, NNTP, the protocol for moving USENET news articles around, and HTTP, the protocol for fetching pages on the World Wide Web, and many others.

Comparison of the OSI and TCP/IP Reference Models

The OSI and TCP/IP reference models have much in common. Both are based on the concept of a stack of independent protocols. Also, the functionality of the layers is roughly similar. For example, in both models the layers up through and including the transport layer are there to provide an end-to-end, network-independent transport service to processes wishing to communicate. These layers form the transport provider. Again in both models, the layers above transport are application-oriented users of the transport service. Despite these fundamental similarities, the two models also have many differences Three concepts are central to the OSI model

- Services.
- Interfaces.
- Protocols.

Probably the biggest contribution of the OSI model is to make the distinction between these three concepts explicit. Each layer performs some services for the layer above it. The service definition tells what the layer does, not how entities above it access it or how the layer works. It defines the layer's semantics.

A layer's interface tells the processes above it how to access it. It specifies what the parameters are and what results to expect. It, too, says nothing about how the layer works inside. Finally, the peer protocols used in a layer are the layer's own business

It can use any protocols it wants to, as long as it gets the job done (i.e., provides the offered services). It can also change them at will without affecting software in higher layers.

The TCP/IP model did not originally clearly distinguish between service, interface, and protocol, although people have tried to retrofit it after the fact to make it more OSI-like. For example, the only real services offered by the internet layer are send ip packet and receive ip packet.

As a consequence, the protocols in the OSI model are better hidden than in the TCP/IP model and can be replaced relatively easily as the technology changes. Being able to make such changes is one of the main purposes of having layered protocols in the first place. The OSI reference model was devised before the corresponding protocols were invented. This ordering means that the model was not biased toward one particular set of protocols, a fact that made it quite general. The downside of this ordering is that the designers did not have much experience with the subject and did not have a good idea of which functionality to put in which layer.

Another difference is in the area of connectionless versus connection-oriented communication. The OSI model supports both connectionless and connection-oriented communication in the network layer, but only connection-oriented communication in the

transport layer, where it counts (because the transport service is visible to the users). The TCP/IP model has only one mode in the network layer (connectionless) but supports both modes in the transport layer, giving the users a choice. This choice is especially important for simple request-response protocols.

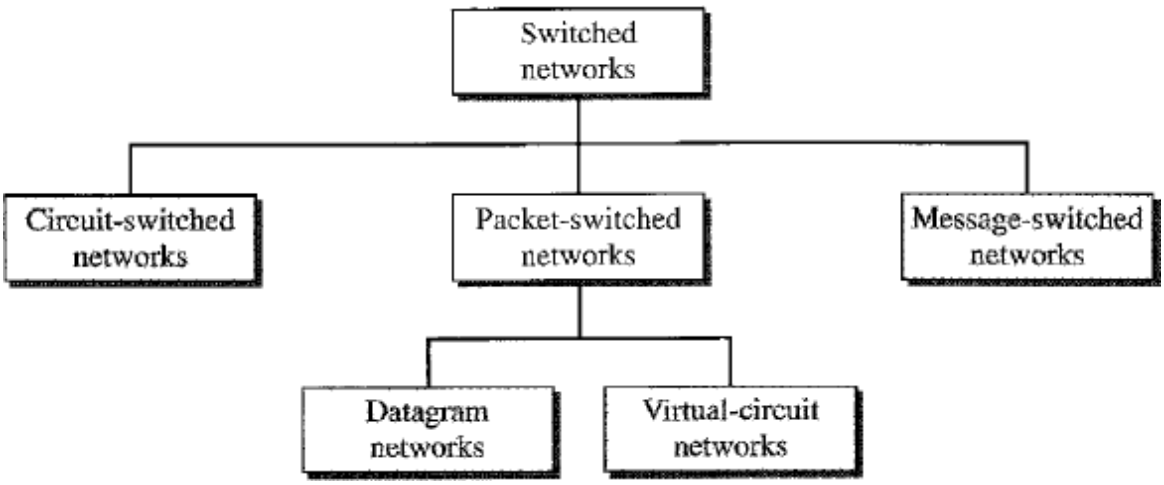

Figure 24 Taxonomy of switched network

CIRCUIT -SWITCHED NETWORKS

A circuit-switched network consists of a set of switches connected by physical links. A connection between two stations is a dedicated path made of one or more links. However, each connection uses only one dedicated channel on each link. Each link is normally divided into n channels by using FDM or TDM

Figure shows a trivial circuit-switched network with four switches and four links. Each link is divided into n (n is 3 in the figure) channels by using FDM or TDM.

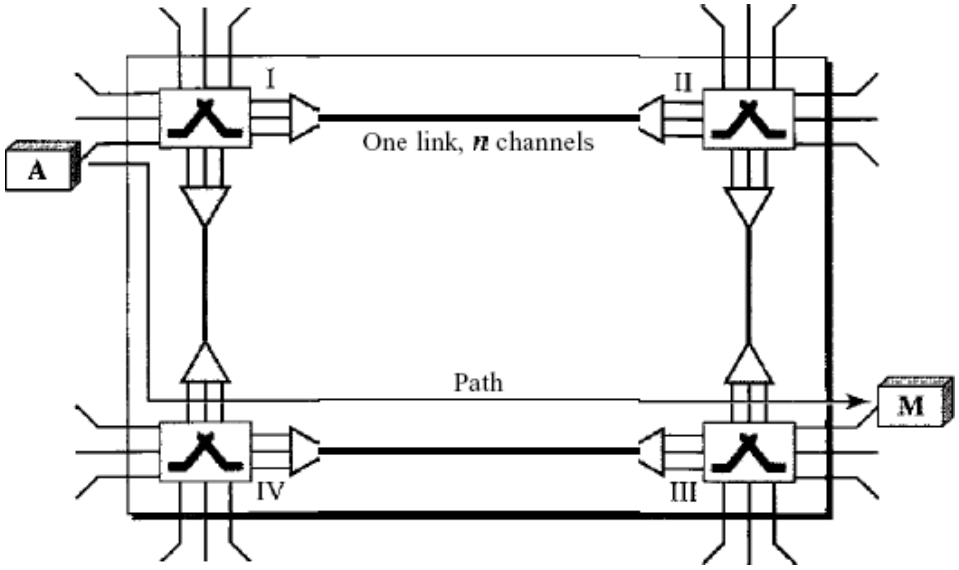

Figure 25 A trivial circuit-switched network

Three phases

The actual communication in a circuit-switched network requires three phases connection setup,data transfer, and connection teardown.

Setup Phase

Before the two parties (or multiple parties in a conference call) can communicate, a dedicated circuit (combination of channels in links) needs to be established. The end systems are normally connected through dedicated lines to the switches, so connection setup means creating dedicated channels between the switches. For example, in Figure, when system A needs to connect to system M, it sends a setup request that includes the address of system M, to switch I. Switch I finds a channel between itself and switch IV that can be dedicated for this purpose. Switch I then sends the request to switch IV, which finds a dedicated channel between itself and switch III. Switch III informs system M of system A's intention at this time.

In the next step to making a connection, an acknowledgment from system M needs to be sent in the opposite direction to system A. Only after system A receives this acknowledgment is the connection established. Note that end-to-end addressing is required for creating a connection between the two end systems. These can be, for example, the addresses of the computers assigned by the administrator in a TDM network, or telephone numbers in an FDM network.

Data Transfer Phase

After the establishment of the dedicated circuit (channels), the two parties can transfer data.

Teardown Phase

When one of the parties needs to disconnect, a signal is sent to each switch to release theresources.

Efficiency

It can be argued that circuit-switched networks are not as efficient as the other two types of networks because resources are allocated during the entire duration of the connection. These resources are unavailable to other connections. In a telephone network, people normally terminate the communication when they have finished their conversation. However, in computer networks, a computer can be connected to another computer even if there is no activity for a long time. In this case, allowing resources to be dedicated means that other connections are deprived.

Although a circuit-switched network normally has low efficiency, the delay in this type of network is minimal. During data transfer the data are not delayed at each switch; the resources are allocated for the duration of the connection. Figure 8.6 shows the idea of delay in a circuit-switched network when only two switches are involved. As Figure shows, there is no waiting time at each switch. The total delay is due to the time needed to create the connection, transfer data, and disconnect the circuit.

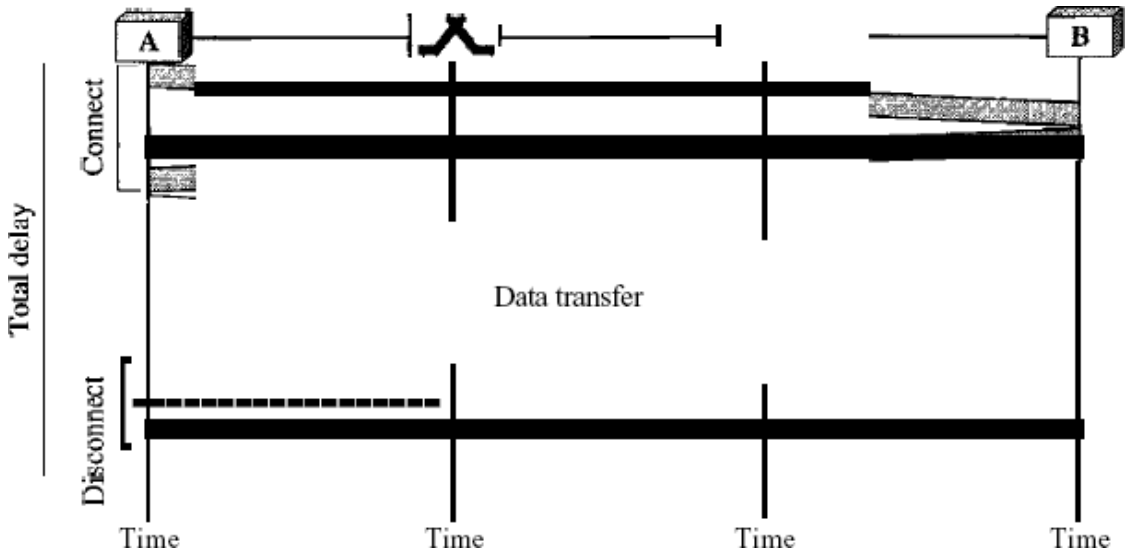

Figure 25 Delay in a circuit-switched network

The delay caused by the setup is the sum of four parts the propagation time of the source computer request (slope of the first gray box), the request signal transfer time (height of the first gray box), the propagation time of the acknowledgment from the destination computer (slope of the second gray box), and the signal transfer time of the acknowledgment (height of the second gray box). The delay due to data transfer is the sum of two parts the propagation time (slope of the colored box) and data transfer time (height of the colored box), which can be very long. The third box shows the time needed to tear down the circuit. We have shown the case in which the receiver requests disconnection, which creates the maximum delay.

DATAGRAM NETWORKS

In a datagram network, each packet is treated independently of all others. Even if a packet is part of a multipacket transmission, the network treats it as though it existed alone. Packets in this approach are referred to as datagrams.

Datagram switching is normally done at the network layer. We briefly discuss datagram networks here as a comparison with circuit-switched and virtual-circuit switched networks Figure shows how the datagram approach is used to deliver four packets from station A to station X. The switches in a datagram network are traditionally referred to as routers. That is why we use a different symbol for the switches in the figure.

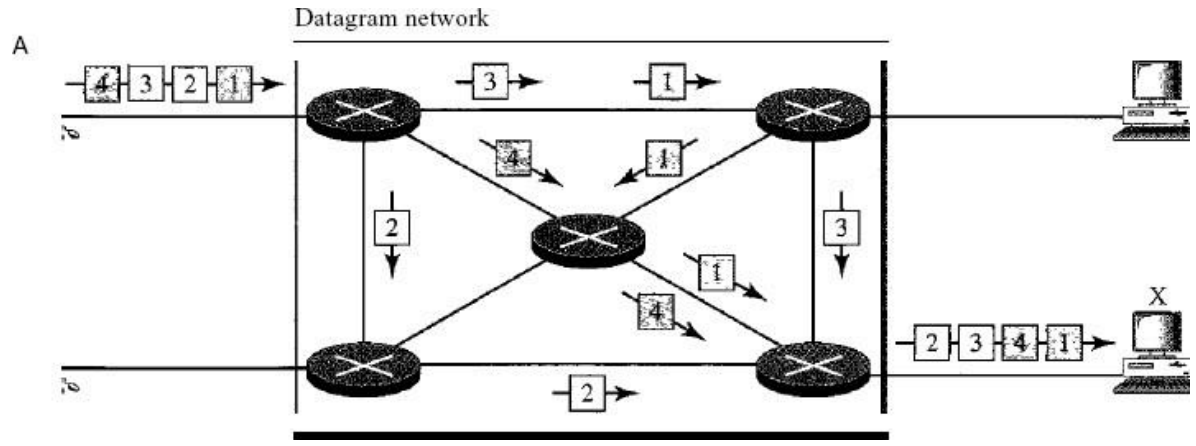

Fig 25 A datagram network with four switches (routers)

In this example, all four packets (or datagrams) belong to the same message, but may travel different paths to reach their destination. This is so because the links may be involved in carrying packets from other sources and do not have the necessary bandwidth available to carry all the packets from A to X. This approach can cause the datagrams of a transmission to arrive at their destination out of order with different delays between the packets. Packets may also be lost or dropped because of a lack of resources. In most protocols, it is the responsibility of an upper-layer protocol to reorder the datagrams or ask for lost datagrams before passing them on to the application.

The datagram networks are sometimes referred to as connectionless networks. The term connectionless here means that the switch (packet switch) does not keep information about the connection state. There are no setup or teardown phases. Each packet is treated the same by a switch regardless of its source or destination.

Routing Table

If there are no setup or teardown phases, how are the packets routed to their destinations in a datagram network? In this type of network, each switch (or packet switch) has a routing table which is based on the destination address. The routing tables are dynamic and are updated periodically. The destination addresses and the corresponding forwarding output ports are recorded in the tables. This is different from the table of a circuit switched network in which each entry is created when the setup phase is completed and deleted when the teardown phase is over. Figure shows the routing table for a switch.

Destination Address

Every packet in a datagram network carries a header that contains, among other information, the destination address of the packet. When the switch receives the packet, this destination address is examined; the routing table is consulted to find the corresponding port through which the packet should be forwarded. This address, unlike the address in a virtual-circuit-switched network, remains the same during the entire journey of the packet.

Efficiency

The efficiency of a datagram network is better than that of a circuit-switched network; resources are allocated only when there are packets to be transferred. If a source sends a packet and there is a delay of a few minutes before another packet can be sent, the resources can be reallocated during these minutes for other packets from other sources.

Delay

There may be greater delay in a datagram network than in a virtual-circuit network. Although there are no setup and teardown phases, each packet may experience a wait at a switch before it is forwarded. In addition, since not all packets in a message necessarily travel through the same switches, the delay is not uniform for the packets of a message.

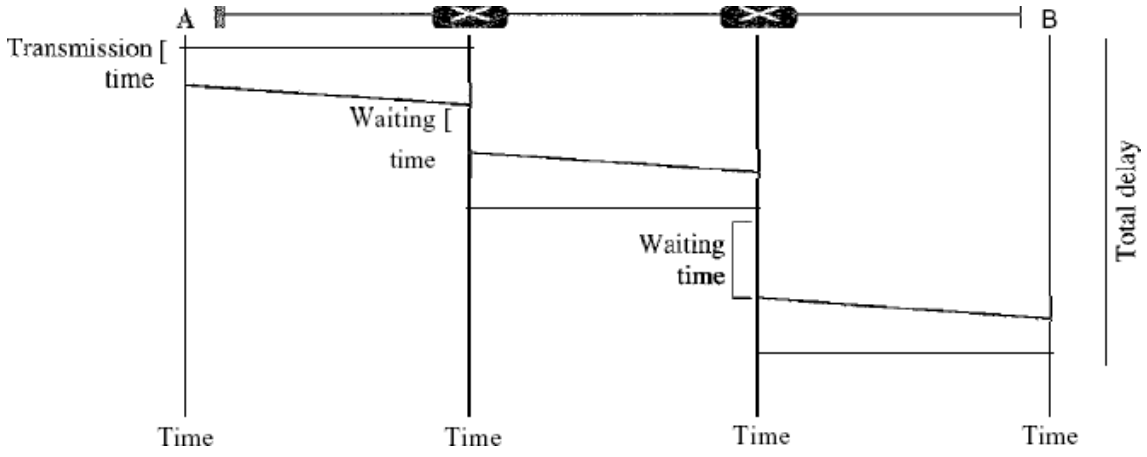

Figure 27 Delay in a datagram network

The packet travels through two switches. There are three transmission times (3T), three propagation delays (slopes 3't of the lines), and two waiting times (WI + w2)' We ignore the processing time in each switch. The total delay is

Total delay =3T + 3t + WI + W2

VIRTUAL-CIRCUIT NETWORKS

A virtual-circuit network is a cross between a circuit-switched network and a datagram network. It has some characteristics of both.

As in a circuit-switched network, there are setup and teardown phases in addition to the data transfer phase.

Resources can be allocated during the setup phase, as in a circuit-switched network, or on demand, as in a datagram network.

As in a datagram network, data are packetized and each packet carries an address in the header. However, the address in the header has local jurisdiction (it defines what should be the next switch and the channel on which the packet is being canied), not end-to-end jurisdiction. The reader may ask how the intermediate switches know where to send the packet if there is no final destination address carried by a packet. The answer will be clear when we discuss virtual-circuit identifiers in the next section.

As in a circuit-switched network, all packets follow the same path established during the connection.

A virtual-circuit network is normally implemented in the data link layer, while a circuit-switched network is implemented in the physical layer and a datagram network in the network layer. But this may change in the future. Figure is an example of a virtual-circuit network. The network has switches that allow traffic from sources to destinations. A source or destination can be a computer, packet switch, bridge, or any other device that connects other networks.

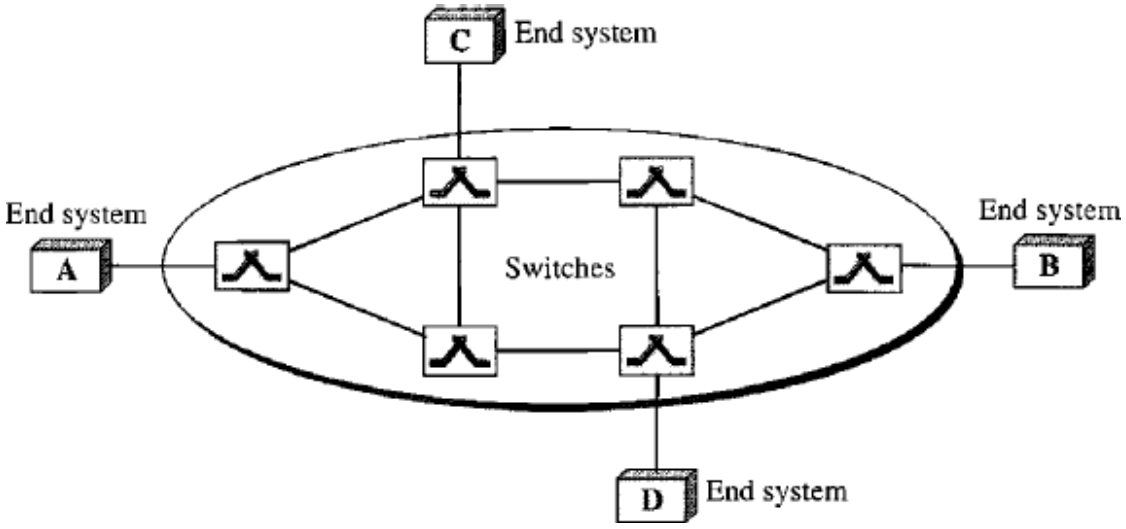

Figure 28 Virtual-circuit network

virtual-circuit network, two types of addressing are involved global and local (virtual-circuit identifier).

Global Addressing A source or a destination needs to have a global address-an address that can be unique in the scope of the network or internationally if the network is part of an international network. However, we will see that a global address in virtual-circuit networks is used only to create a virtual-circuit identifier, as discussed next.

Virtual-Circuit Identifier The identifier that is actually used for data transfer is called the virtual-circuit identifier (Vel). A vel, unlike a global address, is a small number that has only switch scope; it is used by a frame between two switches. When a frame arrives at a switch, it has a VCI; when it leaves, it has a different VCl. Figure 8.11 shows how the VCI in a data frame framechanges from one switch to another. Note that a VCI does not need to be a large number sinceeach switch can use its own unique set of VCls.

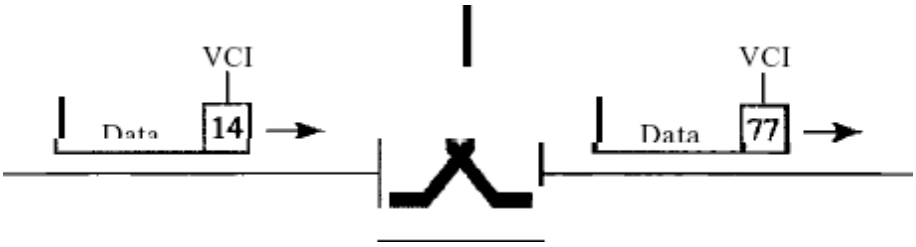

Figure 29 Virtual-circuit identifier

Three Phases

As in a circuit-switched network, a source and destination need to go through three phases in a virtual-circuit network setup, data transfer, and teardown. In the setup phase, the source and destination use their global addresses to help switches make table entries for the connection. In the teardown phase, the source and destination inform the switches to delete the corresponding entry. Data transfer occurs between these two phases. We first discuss the data transfer phase, which is more straightforward; we then talk about the setup and teardown phases.

Data Transfer Phase

To transfer a frame from a source to its destination, all switches need to have a table entry for this virtual circuit. The table, in its simplest form, has four columns. This means that the switch holds four pieces of information for each virtual circuit that is already set up. We show later how the switches make their table entries, but for the moment we assume that each switch has a table with entries for all active virtual circuits. Figure 2 shows such a switch and its corresponding table. And also shows a frame arriving at port 1 with a VCI of 14. When the frame arrives, the switch looks in its table to find port 1 and a VCI of 14. When it is found, the switch knows to change the VCI to 22 and send out the frame from port 3. Figure 3 shows how a frame from source A reaches destination B and how its VCI changes during the trip. Each switch changes theVCI and routes the frame. The data transfer phase is active until the source sends all its frames tothe destination. The procedure at the switch is the same for each frame of a message. The processcreates a virtual circuit, not a real circuit, between the source and destination.

Setup Phase

In the setup phase, a switch creates an entry for a virtual circuit. For example, suppose source A needs to create a virtual circuit to B. Two steps are required the setup request and the acknowledgment.

Figure 2 Switch and tables in a virtual-circuit network

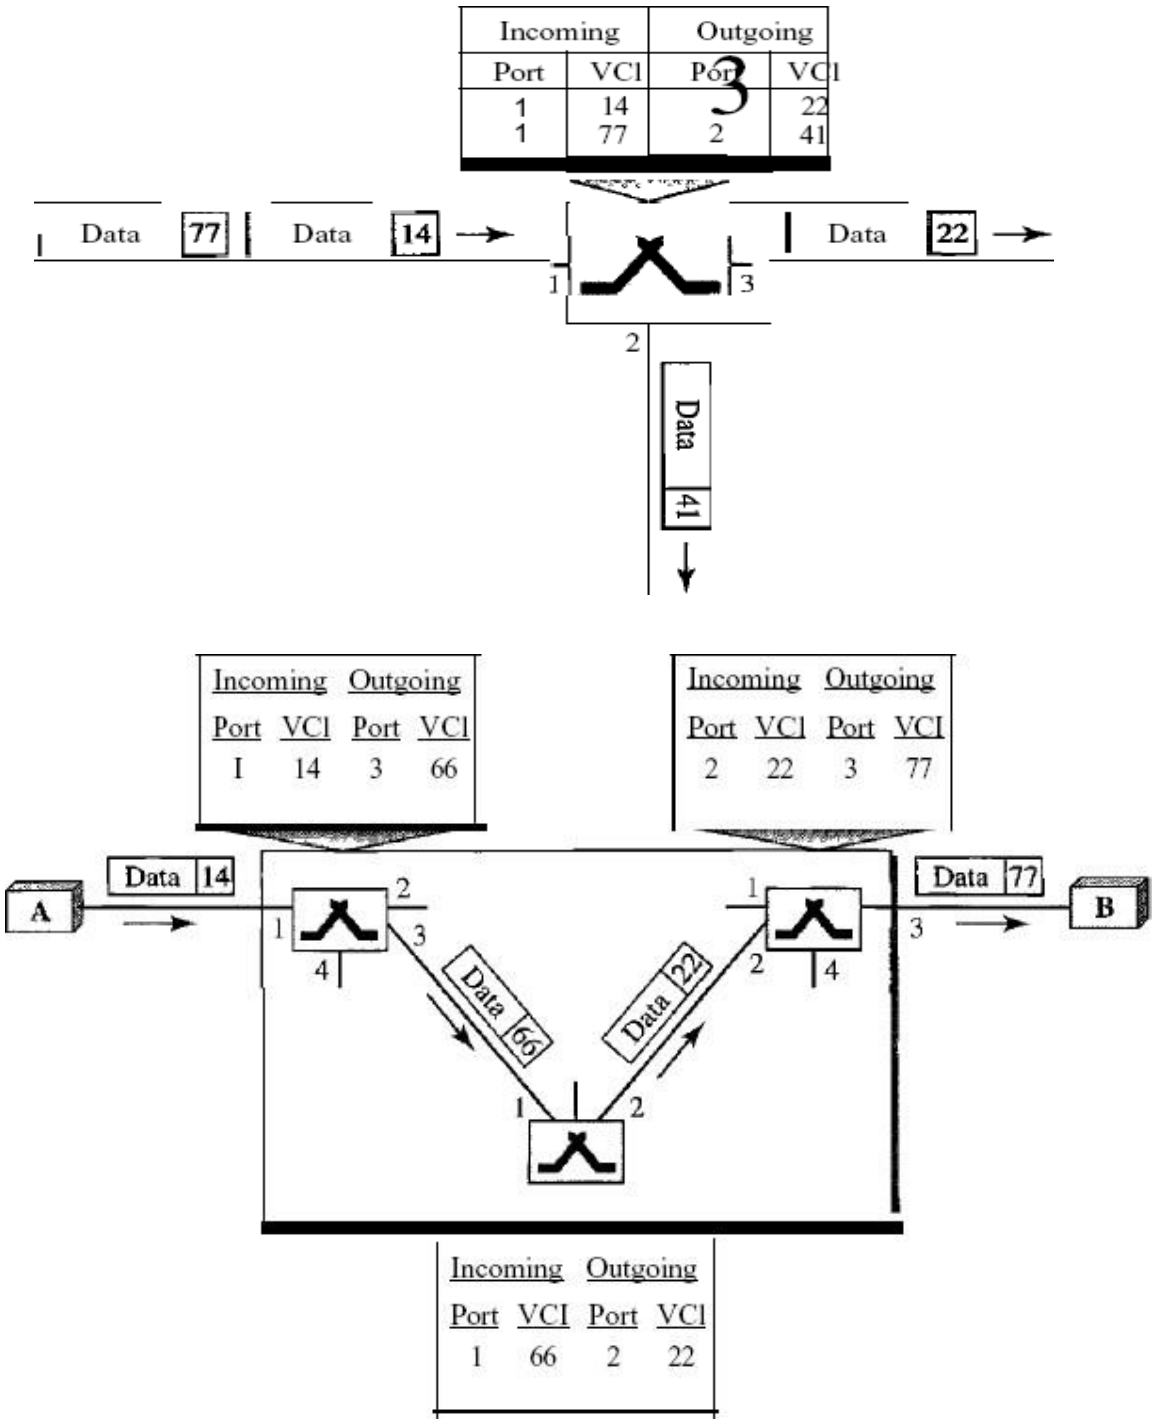

Figure 30 Source-to-destination data transfer in a virtual-circuit network

Setup Request

A setup request frame is sent from the source to the destination.Figure 4 shows the process.

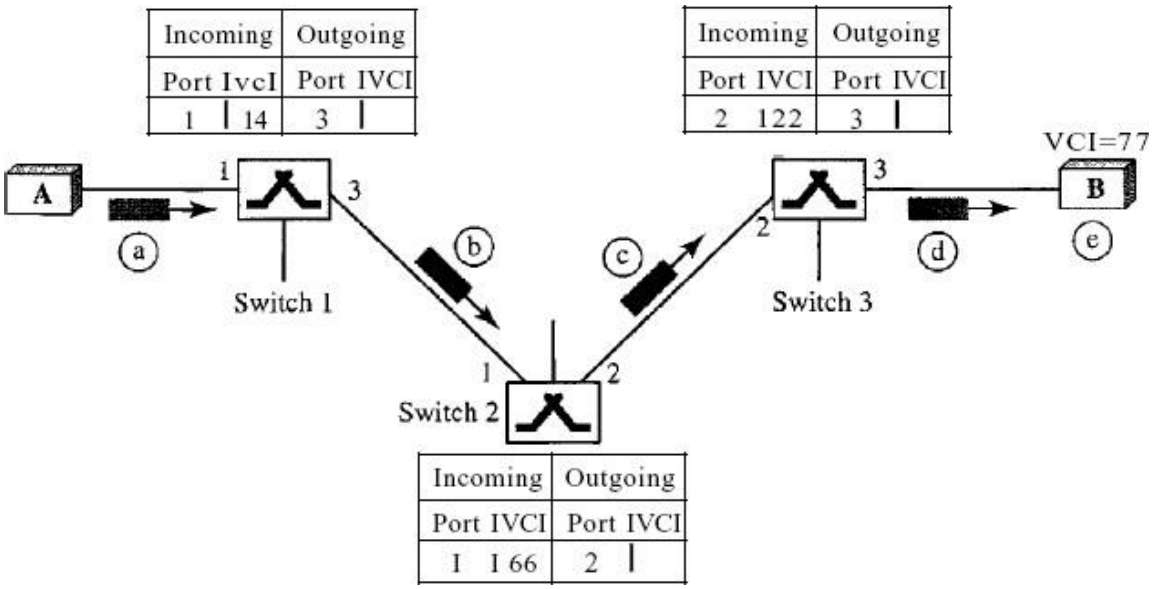

Figure 31 Setup request in a virtual-circuit network

Source A sends a setup frame to switch 1.

Switch 1 receives the setup request frame. It knows that a frame going from A to B goes out through port 3. How the switch has obtained this information is a point covered in future chapters. The switch, in the setup phase, acts as a packet switch; it has a routing table which is different from the switching table. For the moment, assume that it knows the output port. The switch creates an entry in its table for this virtual circuit, but it is only able to fill three of the four columns. The switch assigns the incoming port (1) and chooses an available incoming VCI (14) and the outgoing port (3). It does not yet know the outgoing VCI, which will be found during the acknowledgment step. The switch then forwards the frame through port 3 to switch 2 .Switch 2 receives the setup request frame. The same events happen here as at switch 1; three columns of the table are completed in this case, incoming port (l), incoming VCI (66), and outgoing port (2). Switch 3 receives the setup request frame. Again, three columns are completed incoming port (2), incoming VCI (22), and outgoing port (3).

- Destination B receives the setup frame, and if it is ready to receive frames from A, it assigns a VCI to the incoming frames that come from A, in this case 77. This VCI lets the destination know that the frames come from A, and not other sources.Acknowledgment A special frame, called the acknowledgment frame, completes the entries in the switching tables. Figure 8.15 shows the process.The destination sends an acknowledgment to switch 3. The acknowledgment carries the global source and destination addresses so the switch knows which entry in the table is to be completed.The frame also carries VCI 77, chosen by the destination as the incoming VCI for frames from A. Switch 3 uses this VCI to complete the outgoing VCI column for this entry. Note that 77 is the incoming VCI for destination B, but the outgoing VCI for switch 3.

- Switch 3 sends an acknowledgment to switch 2 that contains its incoming VCI in the table, chosen in the previous step. Switch 2 uses this as the outgoing VCI in the table.
- Switch 2 sends an acknowledgment to switch 1 that contains its incoming VCI in the table, chosen in the previous step. Switch 1 uses this as the outgoing VCI in the table.
- Finally switch 1 sends an acknowledgment to source A that contains its incoming VCI in the table, chosen in the previous step.
- The source uses this as the outgoing VCI for the data frames to be sent to destination B.

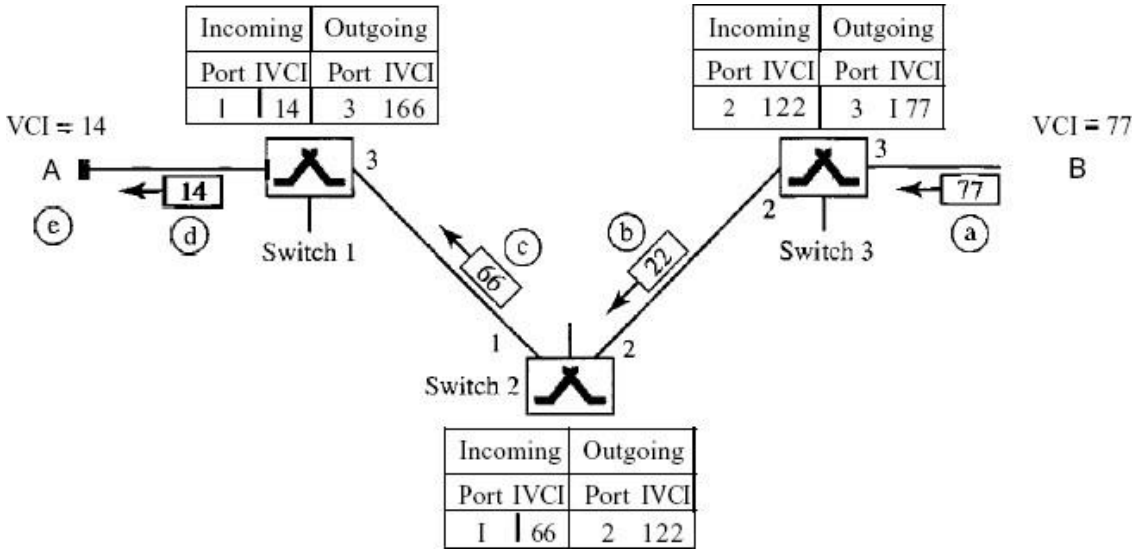

Figure 32 Setup acknowledgments in a virtual-circuit network

Teardown Phase

In this phase, source A, after sending all frames to B, sends a special frame called a teardown request. Destination B responds with a teardown confirmation frame. All switches delete the corresponding entry from their tables.

Efficiency

As we said before, resource reservation in a virtual-circuit network can be made during the setup or can be on demand during the data transfer phase. In the first case, the delay for each packet is the same; in the second case, each packet may encounter different delays. There is one big advantage in a virtual-circuit network even if resource allocation is on demand. The source can check the availability of the resources, without actually reserving it. Consider a family that

wants to dine at a restaurant. Although the restaurant may not accept reservations (allocation of the tables is on demand), the family can call and find out the waiting time. This can save the family time and effort.

Delay in Virtual-Circuit Networks

In a virtual-circuit network, there is a one-time delay for setup and a one-time delay for teardown. If resources are allocated during the setup phase, there is no wait time for individual packets. Below Figure shows the delay for a packet traveling through two switches in a virtual circuit network.

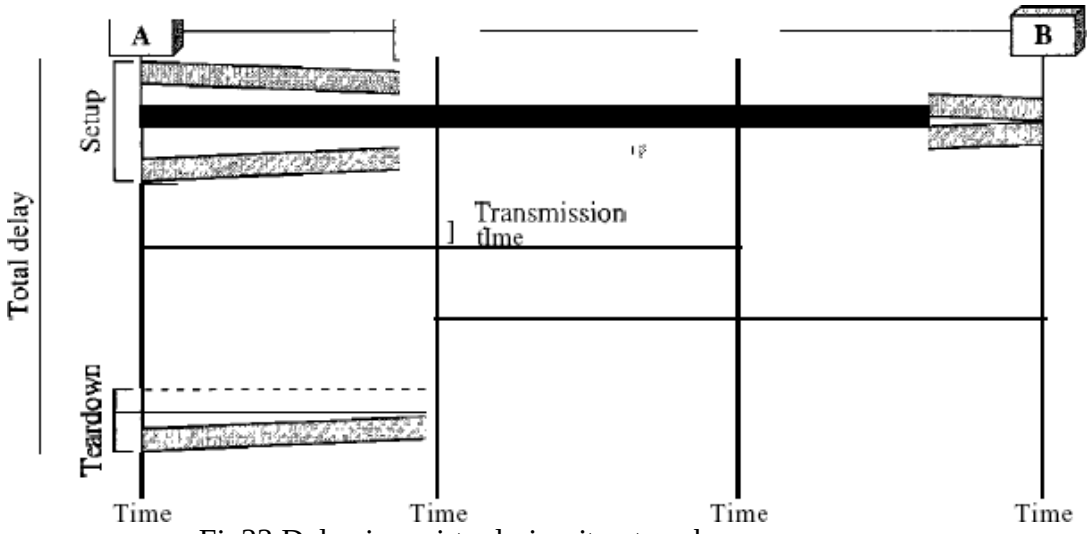

Fig33 Delay in a virtual-circuit network

The packet is traveling through two switches (routers). There are three transmission times (3T), three propagation times (3't), data transfer depicted by the sloping lines, a setup delay (which includes transmission and propagation in two directions), and a teardown delay (which includes transmission and propagation in one direction). We ignore the processing time in each switch. The total delay time is

Total delay = 3T+ 3't + setup delay + teardown delay

IEEE STANDARDS

In 1985, the Computer Society of the IEEE started a project, called Project 802, to set standards to enable intercommunication among equipment from a variety of manufacturers. Project 802 does not seek to replace any part of the OSI or the Internet model. Instead, it is a way of specifying functions of the physical layer and the data link layer of major LAN protocols.

The standard was adopted by the American National Standards Institute (ANSI). In 1987, the International Organization for Standardization (ISO) also approved it as an international standard under the designation ISO 8802. The relationship of the 802 Standard to the traditional OSI model is shown in Figure 1.

The IEEE has subdivided the data link layer into two sub layers logical link control (LLC) and media access control (MAC).

IEEE has also created several physical layer standards for different LAN protocols.

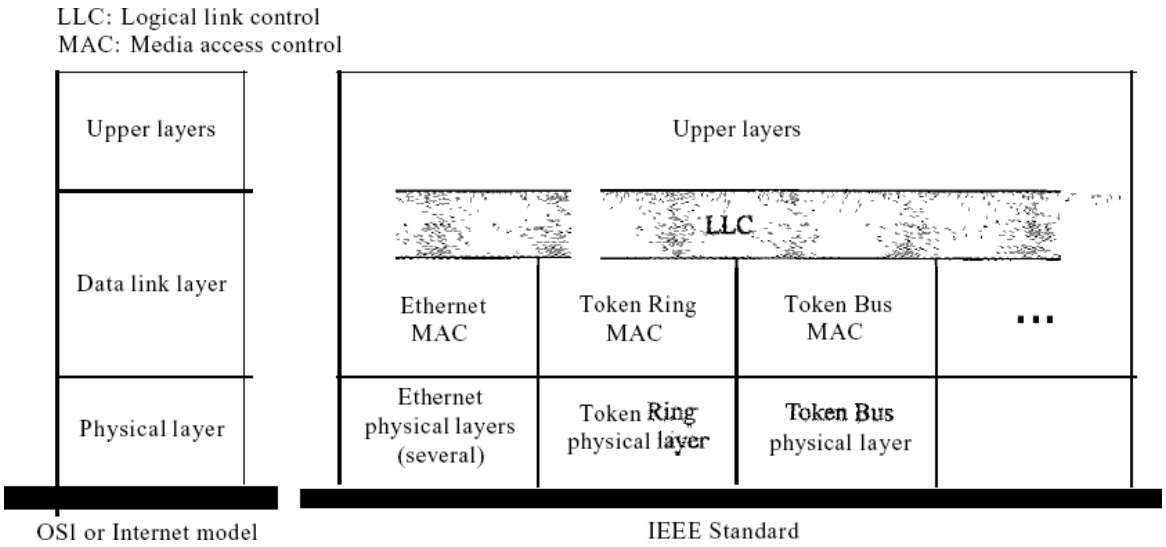

Figure 34 IEEE standardfor LANs

Data Link Layer

As we mentioned before, the data link layer in the IEEE standard is divided into two sublayersLLC and MAC.

Logical Link Control (LLC)

We said that data link control handles framing, flow control, and error control. In IEEE Project 802, flow control, error control,and part of the framing duties are collected into one sublayer called the logical link control. Framing is handled in both the LLC sublayer and the MAC sublayer.

The LLC provides one single data link control protocol for all IEEE LANs. In this way, the LLC is different from the media access control sublayer, which provides different protocols fordifferent LANs. A single LLC protocol can provide interconnectivity between different LANs

because it makes the MAC sublayer transparent. Figure 1 shows one single LLC protocol serving several MAC protocols. Framing LLC defines a protocol data unit (PDU) that is somewhat similar to that of HDLC. The header contains a control field like the one in HDLC; this field is used for flow and error control. The two other header fields define the upper-layer protocol at the source and destination that uses LLC. These fields are called the destination service access point (DSAP) and the source service access point (SSAP). The other fields defined in a typical data link control protocol such as HDLC are moved to the MAC sublayer. In other words, a frame defined in HDLC is divided into a PDU at the LLC sublayer and a frame at the MAC sublayer, as shown in Figure 35.

Need for LLC The purpose of the LLC is to provide flow and error control for the upper-layer protocols that actually demand these services. For example, if a LAN or several LANs are used in an isolated system, LLC may be needed to provide flow and error control for the application layer protocols. However, most upper-layer protocols such as IP, do not use the services of LLC.

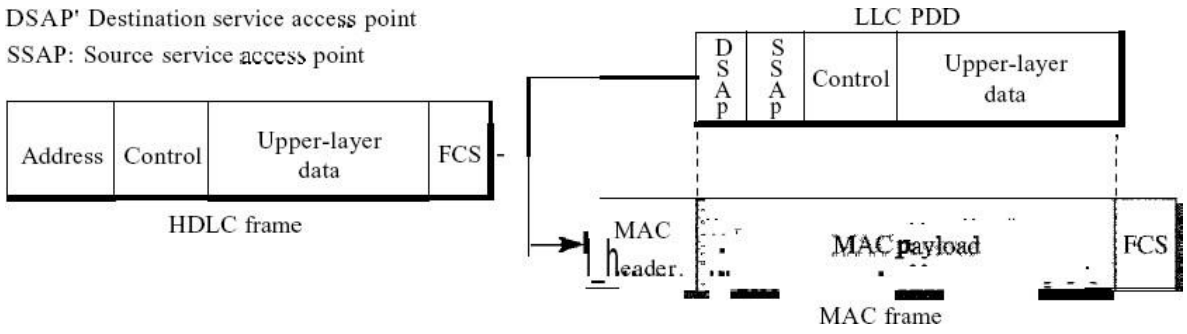

Figure 35 HDLCframe compared with LLC and MAC frames

Media Access Control (MAC)

IEEE Project 802 has created a sublayer called media access control that defines the specific access method for each LAN. For example, it defines CSMA/CD as the media access method for Ethernet LANs and the tokenpassing method for Token Ring and Token Bus LANs. As wediscussed in the previous section, part of the framing function is also handled by the MAC layer. In contrast to the LLC sublayer, the MAC sublayer contains a number of distinct modules; each defines the access method and the framing format specific to the

corresponding LAN protocol.

Physical Layer

The physical layer is dependent on the implementation and type of physical media used. IEEE defines detailed specifications for each LAN implementation. For example, although there is only one MAC sublayer for Standard Ethernet, there is a different physical layer specifications for each Ethernet implementations.

Standard Ethernet

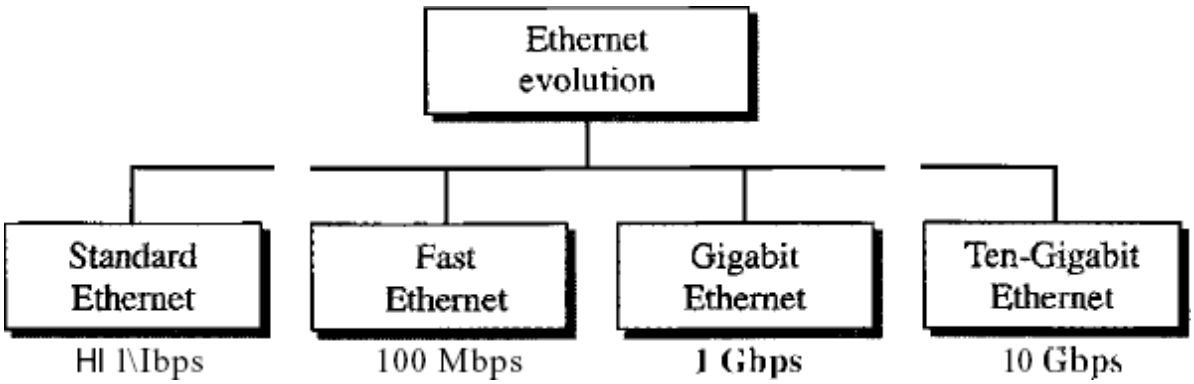

Figure 36 Ethernet evolution through four generations

MAC Sublayer

In Standard Ethernet, the MAC sublayer governs the operation of the access method. It also frames data received from the upper layer and passes them to the physical layer.

Frame Format

The Ethernet frame contains seven fields preamble, SFD, DA, SA, length or type of protocol data unit (PDU), upper-layer data, and the CRe. Ethernet does not provide any mechanism for acknowledging received frames, making it what is known as an unreliable medium.

Acknowledgments must be implemented at the higher layers. The format ofthe MAC frame isshown in Figure 37.

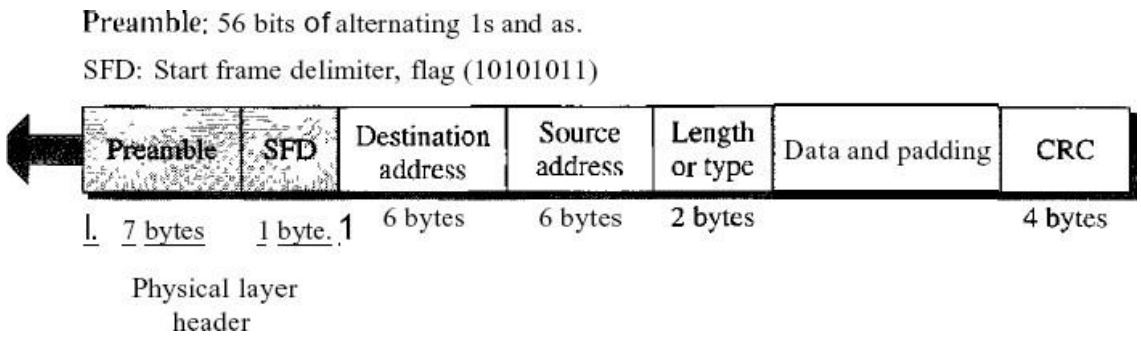

Figure 37 802.3 MACframe

- D Preamble. The first field of the 802.3 frame contains 7 bytes (56 bits) of alternating Os and Is that alerts the receiving system to the coming frame and enables it to synchronize its input timing. The pattern provides only an alert and a timing pulse. The 56-bit pattern allows the stations to miss some bits at the beginning of the frame. The preamble is actually added at the physical layer and is not (formally) part of the frame.
- D Start frame delimiter (SFD). The second field (l byte 10101011) signals the beginning of the frame. The SFD warns the station or stations that this is the last chance for synchronization. The last 2 bits is 11 and alerts the receiver that the next field is the destination address.
- Destination address (DA). The DA field is 6 bytes and contains the physical address of the destination station or stations to receive the packet.
- Source address (SA). The SA field is also 6 bytes and contains the physical address of the sender of the packet. We will discuss addressing shortly.
- Length or type. This field is defined as a type field or length field. The original Ethernet used this field as the type field to define the upper-layer protocol using the MAC frame. The IEEE standard used it as the length field to define the number of bytes in the data field. Both uses are common today.
- Data. This field carries data encapsulated from the upper-layer protocols. It is a minimumof 46 and a maximum of 1500 bytes.
- CRC. The last field contains error detection information, in this case a CRC-32

Frame Length

Ethernet has imposed restrictions on both the minimum and maximum lengths of a frame, as

shown in Figure 38.

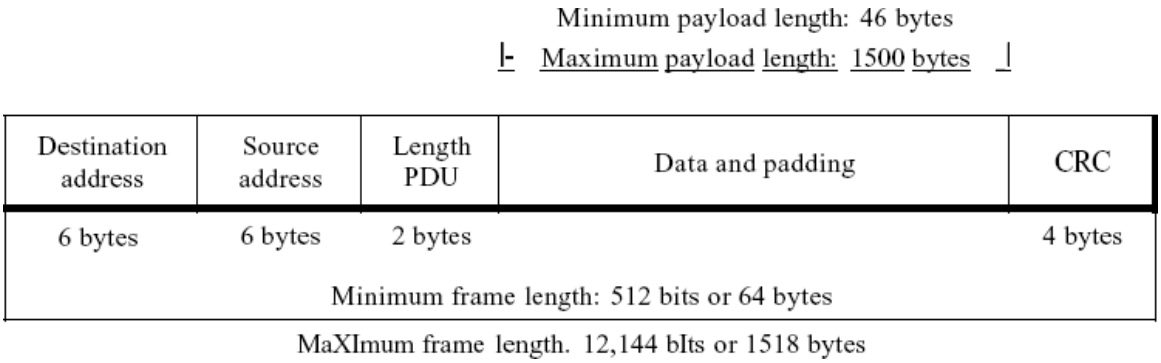

Figure 38 Minimum and maximum lengths

The minimum length restriction is required for the correct operation of CSMAlCD. An Ethernet frame needs to have a minimum length of 512 bits or 64 bytes. Part of this length is the header and the trailer. If we count 18 bytes of header and trailer (6 bytes of source address, 6 bytes of destination address, 2 bytes of length or type, and 4 bytes of CRC), then the minimum length of data from the upper layer is 64 - 18 = 46 bytes. If the upper-layer packet is less than 46 bytes, padding is added to make up the difference.

The standard defines the maximum length of a frame (without preamble and SFD field) as 1518 bytes. If we subtract the 18 bytes of header and trailer, the maximum length of the payload is 1500 bytes. The maximum length restriction has two historical reasons. First, memory was very expensive when Ethernet was designed a maximum length restriction helped to reduce the size of the buffer. Second, the maximum length restriction prevents one station from monopolizing the shared medium, blocking other stations that have data to send.

Frame length

Minimum 64 bytes (512 bits) Maximum 1518 bytes (12,144 bits)

Addressing

Each station on an Ethernet network (such as a PC, workstation, or printer) has its own networkinterface card (NIC). The NIC fits inside the station and provides the station with a 6-byte

physical address. As shown in Figure 6, the Ethernet address is 6 bytes (48 bits), nonnally written in hexadecimal notation, with a colon between the bytes.

0601 02012C4B

6 bytes =12 hex digits =48 bits

Unicast, Multicast, and Broadcast Addresses A source address is always a unicast address-the frame comes from only one station. The destination address, however, can be unicast, multicast, or broadcast. Figure 39 shows how to distinguish a unicast address from a multicast address. If the least significant bit of the first byte in a destination address is 0, the address is unicast; otherwise,it is multicast.

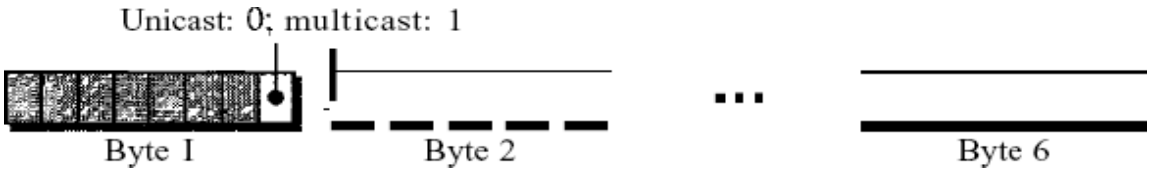

Figure 39 Unicast and multicast addresses

A unicast destination address defines only one recipient; the relationship between the sender and the receiver is one-to-one. A multicast destination address defines a group of addresses; the relationship between the sender and the receivers is one-to-many. The broadcast address is a special case of the multicast address; the recipients are all the stations on the LAN. A broadcast destination address is forty-eight.

Physical Layer

The Standard Ethernet defines several physical layer implementations; four of the most common,are shown in Figure 8.

Encoding and Decoding

All standard implementations use digital signaling (baseband) at 10 Mbps. At the sender, data areconverted to a digital signal using the Manchester scheme; at the receiver, the received signal is interpreted as Manchester and decoded into data. Manchester encoding is self-synchronous, providing a transition at each bit interval. Figure 9 shows the encoding scheme for Standard Ethernet.

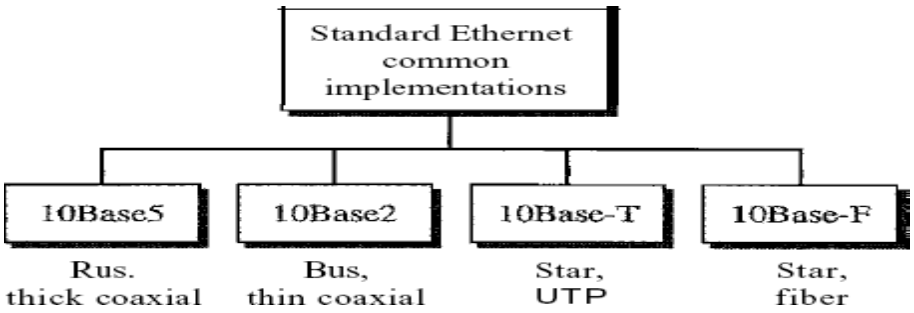

Figure 40 Categories ofStandard Ethernet

Fast Ethernet

Fast Ethernet was designed to compete with LAN protocols such as FDDI or Fiber Channel (or Fibre Channel, as it is sometimes spelled). IEEE created Fast Ethernet under the name 802.3u. Fast Ethernet is backward-compatible with Standard Ethernet, but it can transmit data 10 timesfaster at a rate of 100 Mbps.

The goals of Fast Ethernet can be summarized as follows

- Upgrade the data rate to 100 Mbps.
- Make it compatible with Standard Ethernet.
- Keep the same 48-bit address.
- Keep the same frame format.
- Keep the same minimum and maximum frame lengths.
- MAC Sublayer
- Phylsica Layer

IEEE 802.11

Architecture

The standard defines two kinds of services the basic service set (BSS) and the extended service set (ESS).

Basic Service Set

IEEE 802.11 defines the basic service set (BSS) as the building block of a wireless LAN. A basic service set is made of stationary or mobile wireless stations and an optional central base station, known as the access point (AP). Figure 9 shows two sets in this standard.

The BSS without anAP is a stand-alone network and cannot send data to other BSSs. It is called an ad hoc architecture. In this architecture, stations can form a network without the need of anAP; they can locate one another and agree to be part of a BSS. A BSS with an AP is sometimesreferred to as an infrastructure network.

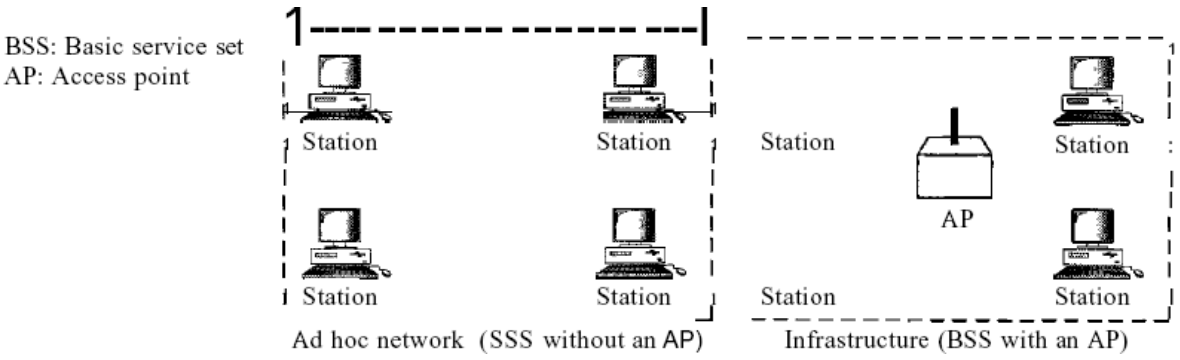

Figure. 41 Basic service sets (BSSs)

Extended Service Set

An extended service set (ESS) is made up of two or more BSSs with APs. In this case, the BSSs are connected through a distribution system, which is usually a wired LAN. The distribution system connects the APs in the BSSs. IEEE 802.11 does not restrict the distribution system; it can be any IEEE LAN such as an Ethernet. Note that the extended service set uses two types of stations mobile and stationary. The mobile stations are normal stations inside a BSS. The stationary stations are AP stations that are part of a wired LAN. Figure 42 shows an ESS.

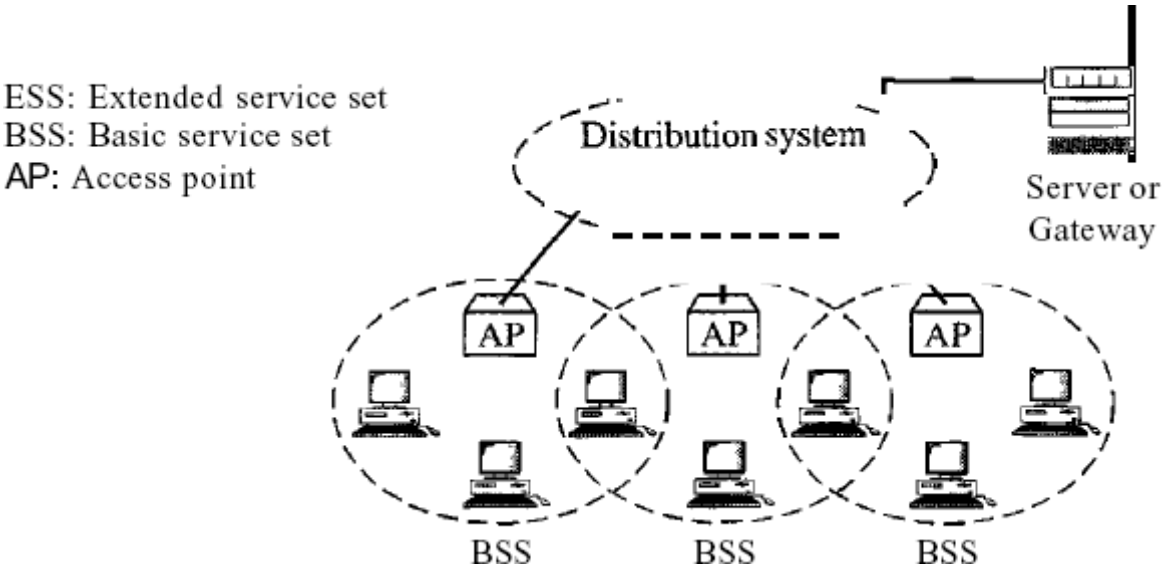

Figure 42 Extended service sets (ESSs)

Token Ring (IEEE 802.5)

Token Ring A Brief History

Originally, IBM developed Token Ring network in the 1970s. It is still IBM's primary local-area network (LAN) technology. The related IEEE 802.5 specification is almost identical to and completely compatible with IBM's Token Ring network. In fact, the IEEE 802.5 specification was modeled after IBM Token Ring, and on the same lines. Theterm Token Ring is generally used to refer to both IBM's Token Ring network and IEEE802.5 networks.

Introduction

Before going into the details of the Token Ring protocol, let's first discuss the motivation behind it. As already discussed, the medium access mechanism used by Ethernet (CSMA/CD) may results in collision. Nodes attempt to a number of times before they can actually transmit, and even when they start transmitting there are chances to encounter collisions and entire transmission need to be repeated. And all this become worse one the traffic is heavy i.e. all nodes have some data to transmit. Apart from this there is no way to predict either the occurrence of collision or delays produced by multiple stations attempting to capture the link at the same time. So all these problems with the Ethernet gives way to an alternate LAN technology, Token Ring.

Token Ring and IEEE802.5 are based on token passing MAC protocol with ring topology. They resolve the uncertainty by giving each station a turn on by one. Each node takes turns sending the data; each station may transmit data during its turn. The technique that coordinates this turn mechanism is called Token passing; as a Token is passed in the network and the station that gets the token can only transmit. As one node transmits at a time, there is no chance of collision. We shall discuss the detailed operation in next section.

Stations are connected by point-to-point links using repeaters. Mainly these links are of shielded twisted-pair cables. The repeaters function in two basic modes Listen mode, Transmit mode. A disadvantage of this topology is that it is vulnerable to link or station failure. But a few measures can be taken to take care of it.

Differences between Token Ring and IEEE 802.5

Both of these networks are basically compatible, although the specifications differ in some ways.

- IEEE 802.5 does not specify a topology, although virtually all IEEE 802.5 implementations are based on the star topology. While IBM's Token Ring network explicitly specifies a star, with all end stations attached to a device called a Multi- Station Access Unit (MSAU).
- IEEE 802.5 does not specify a media type, although IBM Token Ring networks use twisted-pair wire.
- There are few differences in routing information field size of the two.

Token Ring Operation

Token-passing networks move a small frame, called a token, around the network.

Possession of the token grants the right to transmit. If a node receiving the token has no information to send, it passes the token to the next end station. Each station can hold the token for a maximum period of time.

If a station possessing the token does have information to transmit, it seizes the token, alters 1 bit of the token (which turns the token into a start-of-frame sequence), appends the information that it wants to transmit, and sends this information to the next station on the ring. While the information frame is circling the ring, no token is on the network (unless the ring supports early token release), which means that other stations wanting to transmit must wait. Therefore, collisions cannot occur in Token Ring networks. If early token release is supported, a new token can be released immediately after a frame transmission is complete.

The information frame circulates around the ring until it reaches the intended destination station, which copies the information for further processing. The information frame makes a round trip and is finally removed when it reaches the sending station. The sending station can check the returning frame to see whether the frame was seen and subsequently copied by the destination station in error-free form. Then the sending station inserts a new free token on the ring, if it has finished transmission of its packets.

Unlike CSMA/CD networks (such as Ethernet), token-passing networks are deterministic, which means that it is possible to calculate the maximum time that will pass before any end station will be capable of transmitting. Token Ring networks are ideal for applications in which delay must be predictable and robust network operation is important.

Priority System

Token Ring networks use a sophisticated priority system that permits certain user-designated, high-priority stations to use the network more frequently. Token Ring frames have two fields that control priority the priority field and the reservation field.

Only stations with a priority equal to or higher than the priority value contained ina token can seize that token. After the token is seized and changed to an information frame, only stations with a priority value higher than that of the transmitting station can reserve the token for the next pass around the network. When the next token is generated, it includes the higher priority of the reserving station. Stations that raise a token's priority level must reinstate the previous priority after their transmission is complete.

Ring Maintenance

There are two error conditions that could cause the token ring to break down. One is the lost token in which case there is no token the ring, the other is the busy token that circulates endlessly. To overcome these problems, the IEEE 802 standard specifies that one of the stations be designated as ‘active monitor’. The monitor detects the lost condition using a timer by time-out mechanism and recovers by using a new free token. To detect a circulating busy token, the monitor sets a ‘monitor bit’ to one on any passing busy token. If it detects a busy token with the monitor bit already set, it implies that the sending station has failed to remove its packet and recovers by changing the busy tokento a free token. Other stations on the ring have the role of

passive monitor. The primary job of these stations is to detect failure of the active monitor and assume the role of active monitor. A contention-resolution is used to determine which station to take over.

Physical Layer

The Token Ring uses shielded twisted pair of wire to establish point-point links between the adjacent stations. The baseband signaling uses differential Manchester encoding. To overcome the problem of cable break or network failure, which brings the entire network down, one suggested technique, is to use wiring concentrator as shown in Fig43

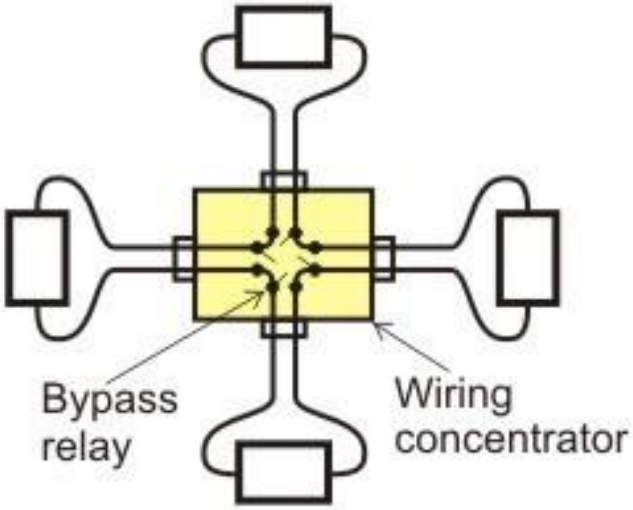

Figure 43 Star Connected Ring topology

It imposes the reliability in an elegant manner. Although logically the network remains as a ring, physically each station is connected to the wire center with two twisted pairs for 2-way communication. Inside the wire center, bypass relays are used to isolate a broken wire or a faulty station. This Topology is known as Star-Connected Ring.

Frame Format

Token Ring and IEEE 802.5 support two basic frame types tokens and data/command frames. Tokens are 3 bytes in length and consist of a start delimiter, an access control byte, and an end delimiter. Data/command frames vary in size, depending on the size of the Information field. Data frames carry information for upper-layer protocols, while command frames contain control information and have no data for upper-layer protocols.

Token Frame Fields

Start Delimiter	Access Control	Ending delimiter

Token Frame contains three fields, each of which is 1 byte in length

- Start delimiter (1 byte) Alerts each station of the arrival of a token (or data/command frame). This field includes signals that distinguish the byte from the rest of the frame by violating the encoding scheme used elsewhere in the frame.
- Access-control (1 byte) Contains the Priority field (the most significant 3 bits) and the Reservation field (the least significant 3 bits), as well as a token bit (used to differentiate a token from a data/command frame) and a monitor bit (used by the active monitor to determine whether a frame is circling the ring endlessly).
- End delimiter (1 byte) Signals the end of the token or data/command frame. This field also contains bits to indicate a damaged frame and identify the frame that is the last in a logical sequence.

Data/Command Frame Fields

Start Delimiter	Access Control	Frame Control	Destination address	Source address	Data	Frame check sequence	End Delimiter	Frame Status

Data/command frames have the same three fields as Token Frames, plus several others. The Data/command frame fields are described below

- Frame-control byte (1 byte)—Indicates whether the frame contains data or control information. In control frames, this byte specifies the type of control information.
- Destination and source addresses (2-6 bytes)—Consists of two 6-byte address fields that identify the destination and source station addresses.
- Data (up to 4500 bytes)—Indicates that the length of field is limited by the ringtoken holding time, which defines the maximum time a station can hold the token.
- Frame-check sequence (FCS- 4 byte)—Is filed by the source station with a calculated value dependent on the frame contents. The destination station recalculates the value to determine whether the frame was damaged in transit. If so, the frame is discarded.
- Frame Status (1 byte)—This is the terminating field of a command/data frame. The Frame Status field includes the address-recognized indicator and frame-copied indicator.

TEXT/ REFERENCE BOOKS

1. Gil Held, "Network Design: Principles and Applications (Best Practices)", Auerbach Publications, 1st edition, 2000.
2. Diane Tiare and Catherine Paquet, "Campus Network Design Fundamentals", Pearson Education, 1st edition, 2006.
3. Larry L. Peterson, Bruce S. Davie, "Computer Networks: A Systems Approach", Morgan Kaufmann Publishers Inc., 5th edition, 2012.
4. William Stallings, "Data and Computer Communications", Pearson Education, 8th edition, 2016.
5. James F. Kurose, Keith W. Ross, "Computer Networking - A Top-Down Approach Featuring the Internet", Pearson Education, 6th edition, 2012.

PART-A - 2 Mark Questions

S. No	Questions (2marks)	CO	Level
1	Illustrate how computer network is implemented in real time?	CO2	2
2	Summarize the list of requirements needed for LAN small business.	CO2	2
3	Compare the LAN ,WAN and MAN	CO2	5
4	Justify the answer why ring ad bus topology is implemented rather than other in LAN	CO2	5
5	Categorize the different types of Ethernet with differentiation	CO2	4
6	Outline the steps involved in integrating of voice over data network.	CO2	2
7	Justify your answer how ATM is evolved as ATM LANE	CO2	5
8	Brief the advantages of Ethernet over token ring	CO2	2
9.	Compare and contrast the various types of Etherne .	CO2	5
10.	Discuss about the Inter switching	CO2	2

PART-B - 10 Mark Questions

S. No	Questions (2marks)	CO	Level
1	Discuss in details about the transfer of information by token concept in token ring with necessary diagram and also explain its framing of packet details .Also justify how the token ring overcomes the token bus concept	CO2	2
2	Illustrate the answer how the Asynchronous transfer mode is evolved as ATM LANE with neat diagram ,	CO2	2
3	An organisation who need to combine the voice information to be integrated with the data network so the designing and establishing of voice and data on the LAN is to be discussed by an expert .Suggest an appropriate procedure for the voice integration . Also comment on ideas required in preparing cabling and other components.	CO2	5
4	Compare the details of different Ethernet with different parameter like bandwidth, distance, meter duplex and cable types with their applications	CO2	5
5	A company want to discuss about a implementation of LAN network for the business in a same organization in a block . Recommend a suitable procedure to them for establishing small business LAN connection	CO2	5

SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF COMPUTER SCIENCE AND ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE ENGINEERING

Computer Networks and Design – SCSA1502

UNIT III-VPNS, INTRANETS AND EXTRANETS-SCSA1502

UNIT III

VPNS, INTRANETS AND EXTRANETS

Virtual Network management and planning – VPNs for small businesses – Secure remote access in VPNs – IPSec VPNs – Integrating data centers with Intranets – Implementing and supporting Extranets.

3.1 VIRTUAL NETWORK MANAGEMENT AND PLANNING

A VIRTUAL NETWORK IS A DATA COMMUNICATIONS SYSTEM that provides access control and network configuration changes using software control. It functions like a traditional network butis built using switches.

The switched virtual network offers all the performance of the bridge with the value of the router. The constraints of physical networking are removed by the logical intelligence that structures and enforces policies of operation to ensure stability and security. Regardless of access technology or geographic location, any to any communications is the goal.

The switch could be considered a third generation internetworking device. First generation devices, or bridges, offered a high degree of performance throughput but relatively little value, because the bridge's limited decision intelligence resulted in broadcast storms that produced network instability. Routers, the second generation of internetworking devices, increased network reliability and offered great value with firewalling capabilities, but the tradeoff was in performance. When routers are used in combination with each other, bandwidth suffers, which is detrimental for delay sensitive applications such as multimedia.

3.1.1 THE BUSINESS CASE FOR VIRTUAL NETWORKING

Both the business manager and the technical manager should find interest in this new virtual networking scheme. The business manager is usually interested in cost of ownership issues. Numerous studies from organizations such as the Gartner Group and Forrester Research have found that only 20 percent of networking costs are associated with capital equipment acquisition. The other 80 percent of annual budgets are dedicated to items such as wide area networking charges, personnel, training, maintenance and vendor support, as

well as the traditional equipment moves, adds, and changes.

It is important for IS managers to remember that capital expenditure happens in year one, even though the equipment may be operating for another four years. Wide area network (WAN) charges can account for up to 40 percent of an organization's networking budget. For every dollar that the technical staff spends on new equipment, another four dollars is spent on the operation of that equipment. Therefore, focus should be on the cost of ownership issues, not necessarily the cost of the network devices.

3.1.2 Network Reliability

Business managers are also looking for increased reliability as the network plays a major role in the core operations of the organization. Networks have become a business tool to gain competitive advantage they are mission critical and, much like a utility, must provide a highly reliable and available means of communications. Every office today includes an electrical outlet, a phone jack, and a network connection. Electrical and phone service are generally regarded as stable utilities that can be relied on daily. Networks, however, do not always provide comparable levels of service.

3.1.3 Network Accountability

Managers also can benefit from the increased accountability that virtual networks are able to offer. Organizational networking budgets can range from hundreds of thousands of dollars to hundreds of millions per year. Accounting for the use of the network that consumes those funds is a critical issue. There is no better example than WAN access charges. Remote site connectivity can consume a great deal of the budget, and the questions of who, what, when, and where with regard to network use are impossible to determine. Most users consider the network to be free, and the tools to manage and account for its use are increasingly a requirement, not an option.

3.1.4 The Technology Case for Virtual Networking

The IS manager's needs for higher capacity, greater performance, and increased efficiency can be met through the deployment of switched virtual networks. Each user is offered dedicated bandwidth to the desktop with uplinks of increasing bandwidth to servers

or other enterprise networks. Rather than contending for bandwidth in shared access environments, all users are provided with their own private link. This degree of privacy allows for increased security because data are sent only to intended recipients, rather than seen by all.

The most attractive feature to the technical manager, however, may be the benefits gained through increased ease of operation and administration of virtual networks. A longstanding objective has been to deliver network services to users without continually having to reconfigure the devices that make up that network. Furthermore, many of the costs associated with moves, adds, and changes of users can be alleviated as the constraints of physical networking are removed. Regardless of user location, they can remain part of the same virtual network. Through the use of graphical tools, users are added and deleted from workgroups. In the same manner, policies of operation and security filters can be applied. In a sense, the virtual net work accomplishes the goal of managing the individual users and individual conversations, rather than the devices that make up the network.

3.1.5 Virtual Networking Defined

The ideal virtual network does not restrict access to a particular topology or protocol. A virtual network that can only support Ethernet users with Transmission Control Protocol/Internet Protocol (TCP/IP) applications is limited. The ultimate virtual network allows any to any connectivity between Ethernet, Token Ring, Fiber Distributed Data Interface (FDDI), Asynchronous Transfer Mode (ATM), Internet Protocol (IP), Internetwork Packet Exchange (IPX), AppleTalk, or Systems Network Architecture (SNA) networks. A single virtual network infrastructure under single management architecture is the goal.

Network management software becomes a key enabling requirement for the construction of switched virtual networks. The greatest challenge net work designer's face is the separation of the physical network connectivity from the logical connection services it can provide. Many of the design issues associated with networks can be attributed to the physical parameters of protocols and the routers used as the interconnection device. A challenge for any manager is to remain compatible with existing layer 3 protocols and routers and still preserve the investment in existing local area network (LAN) equipment to

the greatest extent possible.

3.1.5.1 Using Telephony as a Model

The principles of operation for switched virtual networks are concretely founded in the success of the global communications systems. Without doubt, the phone system is the world's largest and most reliable network. Built using advanced digital switches controlled by software, extensive accounting and management tools ensure the success of this highly effective means of communication. The connection oriented switch is the key. End to end connections across multiple switches and various transmission types ranging from copper to fiber optics to microwave to satellites allow millions of calls per day to be successfully completed, regardless of the type of phone or from where the user is calling. The telephony model is used throughout this chapter to help illustrate the workings of a virtual network.

3.1.5.2 FEATURES OF SWITCHING SOFTWARE

The software that runs on switches is just as important as the switches themselves. A salesperson from Lucent Technologies, Fujitsu, or Northern Telecom does not focus the potential customer on the hardware aspects of the telephone switches. On the contrary, the salesperson conveys the benefits of the call management software, accounting, and automatic call distributor (ACD) functions. Switched virtual networks should also be evaluated for their ability to deliver value because of the software features.

3.1.5.3 The Virtual Network Server

Network management software has traditionally been thought of as soft ware that passively reports the status and operation of devices in the net work. In the switched virtual network, the network management software takes on a new role as an active participant in operations as well as configuration and reporting. A new middleware component known as the virtual network server (VNS) enforces the policies of operation defined by the net work administrator through management software applications. The switches provide the data transport for the users of the network.

Directory Service. One of the software features in the VNS is the directory service. The

directory service allows the identification of a device by logical name, MAC address, network protocol address, and ATM address, along with the switch and port that the user is connected to within the virtual network domain. The directory listing could be populated manually or dynamically as addresses are discovered. To fully realize the benefits of switched virtual networking, automatic configuration is absolutely essential. The directory service allows end nodes to be located and identified.

Security Service. The VNS security service will be used during call setup phases to determine whether users or groups of users were allowed to connect to each other. On a user by user and conversation by conversation basis, the network manager would have control. This communications policy management is analogous to call management on a telephone private branch exchange (PBX) where 900 numbers, long distance, or international calls can be blocked. Users could be grouped together to form policy groups in which rules could be applied to individual users, groups, or even nested groups. Policies could be defined as open or secure, inclusive or exclusive.

A sample default policy can ensure that all communications are specifically defined to the VNS in order to be authorized. Policy groups can be manipulated either through drag and drop graphical user interfaces or programmatically through simple network management protocol (SNMP) commands.

Finally, and most important, the directory service can work in conjunction with the security service to ensure that policies follow the users as they move throughout the network. This feature alone could save time spend maintaining a router access list, as occurs when a user changes location in the traditional network. However, it is important to realize that switched virtual networks ease administrative chores, they do not eliminate them.

Connection Management Service. The VNS connection management service is used to define the path communications would take through the switch fabric. A site may be linked by a relatively high speed ATM link and a parallel but relatively low speed Ethernet link. Network connections with a defined high quality of service (QoS) could traverse the ATM link and lower QoS connections could traverse the Ethernet. This connection management service allows for the transparent rerouting of calls in the event of a network fault. Connection management could also provide ongoing net work monitoring in which individual user conversations could be tapped or traced for easy troubleshooting.

Bandwidth Service. The VNS bandwidth service is used during the call setup when a connection request is made. Video teleconferencing users may require a committed information rate (CIR) of 10 Mbps whereas the terminal emulation users may only require 1 Mbps. This is where ATM end stations and ATM switches negotiate the amount of bandwidth dedicated to a particular virtual circuit using user to network interface (UNI) signaling. Ethernet, Token Ring, and FDDI nodes do not recognize UNI signaling, but the switches they attach to could proxy the signal for the end station, thus allowing a single bandwidth manager for the entire network, not just the ATM portion.

Broadcast Service. The VNS broadcast service uses as its base the concept of the broadcast unknown server (BUS) that is part of the ATM Forum's LAN emulation draft standard. This is how broadcasts are flooded through the network to remain compatible with the operation of many of today's protocols and network operating systems. A degree of intelligence can be assigned to the VNS that would allow for broadcasts or multicasts based on protocol type or even policy group.

Virtual Routing Service. The VNS virtual routing service is one of the most critical components of a virtual network. Just as traditional networks required traditional routers for interconnection, virtual LANs will require virtual routers for internetworking between virtual LANs. In other words, routing is required, but routers may not be. Some protocols such as TCP/IP actually require a router for users on two different sub networks to speak with each other. In addition, most networks today are logically divided based on network layer protocol addresses with routers acting as the building block between segments.

The difference in operation between a virtual router and a traditional router goes back to the connection oriented vs. connectionless distinction. Routing allows for address resolution between the layer 3 protocol addresses and the layer 2 MAC address just as it happens through the address resolution protocol (ARP) process in TCP/IP networks. The VNS virtual routing service performs the address resolution function, but once the end station addresses are resolved, establishes a virtual connection between the two users. Two users separated by a traditional router would always have the router intervening on every single packet because the router would have resolved the protocol addresses to its own MAC address rather than the actual end station's MAC address. This VNS routing service

allows the network to route once for connection setup and switches all successive packets.

Accounting Service. The VNS accounting service is beneficial because it allows the creation of the network bill. Similar to the way a telephone bill is broken down, the accounting service details connection duration with date and time stamp along with bandwidth consumption details. This is most directly applicable in the WAN. For many managers, WAN usage is never really accounted for on an individual user basis, yet it can consume up to 40 percent of the operations budget.

As usage based WAN service options such as integrated services digital network (ISDN) gain popularity, accounting becomes that much more critical. Interexchange carriers (IXCs), competitive access providers, and the regional Bell operating companies (RBOCs) continue to deliver higher bandwidth links with usage based tariffs. In the future, they could install a 155 Mbps synchronous optical network (SONET) OC3 link and only charge for the actual bandwidth used. Unless managers have tools to control access to and account for usage of WAN links, WAN costs will continue to rise. This service lets IS managers know who is using the WAN.

3.1.6 Virtual Networks vs Virtual LANS

Throughout this discussion, words have been carefully chosen to describe the operation of switched virtual networks. Many of the current vendor offerings on the market have as their goal the construction of a switched virtual LAN. These virtual LANs are interconnected using a traditional router device. However, the router has been viewed as the performance bottleneck. Routers should be deployed when segmentation or separation is the need; switches should be used to deliver more bandwidth.

The virtual LAN (VLAN) concept is merely an interim step along the way to realizing the fully virtual network.

The ATM Forum's draft LAN emulation standard allows ATM devices to internetwork with traditional LAN networks such as Ethernet and Token Ring. However, it seems ironic that it essentially tries to make ATM net works operate like a traditional shared access LAN segment. Although it is required for near term deployment of ATM solutions into existing LAN architectures, its position as an end all solution is questionable. A more logical approach uses ATM as the model that LANs must emulate.

3.2 PUTTING VPNS TO WORK FOR SMALL BUSINESSES AND OFFICES

MODERN BUSINESS PROCESSES DEMAND TIGHT LINKS between mobile users, customers, and third parties on both a temporary basis (project based) and permanent basis. Virtual Private Networks (VPNs) can provide significant business benefits by overcoming the barriers to achieving widely available and secure communication. VPNs provide the appearance of a single network connecting corporate offices, telecommuters, customers, and even competitors, while using separate public and private networks. A company retains control of user access and the privacy and integrity of its data even though the data travel on the public Internet. VPNs can provide as much as 60 percent cost savings over private leased lines and significantly reduce telecommuter dialup charges.

VPNs and their many benefits, however, have traditionally been the domain of larger organizations. These huge companies enjoy access to the capital and scale necessary to build VPNs and have the technical staff to maintain them. They are able to use VPNs to enhance and sustain their competitive advantage over their smaller and less technically sophisticated competitors. In practical terms, the benefits of VPNs have been off limits to small and medium sized businesses. And, even larger organizations have had difficulty deploying VPNs in branch offices because they are often too small to justify onsite IT staff.

The barriers to creating and maintaining a VPN include the need to construct and maintain a secure physical infrastructure and administer a wide range of data communications services. The infrastructure challenges include setting up access equipment, firewalls, servers, telecommunications services, and maintaining connections to multiple Internet Service Providers (ISPs) at hundreds or even thousands of enterprise locations.

Administrative challenges include maintenance of servers, synchronization of software upgrades, replication of Web servers, and sophisticated policy management spanning the whole network. Services that must be supported include email, directory, internal and external Web, firewall, FTP, and access control.

Virtual Services Management (VSM) technology and secure VPN transports are making VPNs realistically deployable for smaller organizations and branch offices. VSM solves the service related headaches of multiple points of administration required when setting up multiple sites, users, devices, and Internet Service

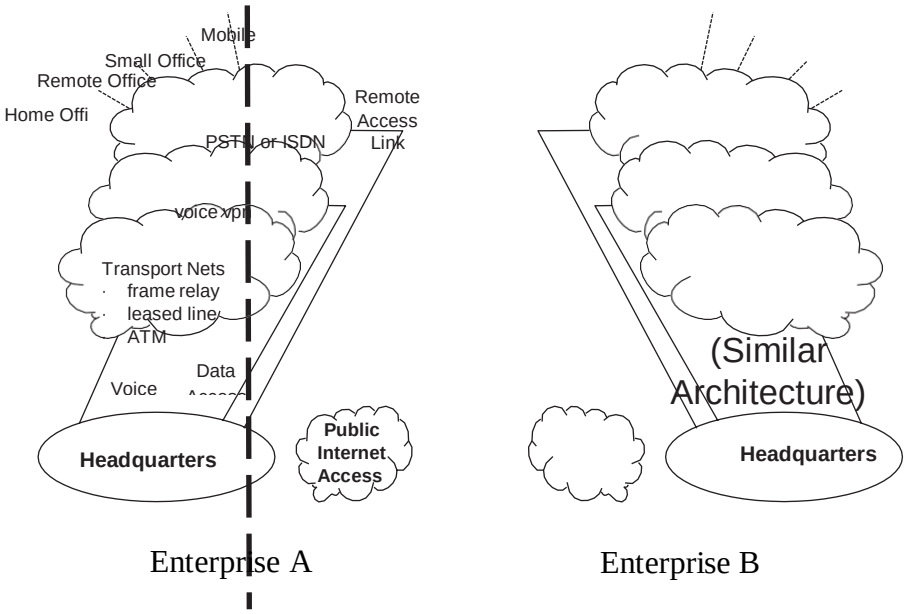

Today's enterprise networks.

3.2.1 Emergence of the Virtual Private Network

Today's private networks resemble the network in Figure. Basic connectivity is provided to a wide variety of locations, but the overhead costs are severe. The functionality includes the following services.

3.2.2 Remote Access

Remote access has matured from a "nice to have" option to a business critical requirement to support a mobile workforce and telecommuters.

For example, utility companies are increasing the productivity of their field service workers and eliminating the cost of maintaining distribution centers by applying remote access technologies. Line crews take their vehicles home with them and receive their day's work orders through either telephone or wireless dispatching systems. This setup eliminates the time it takes to report to the service center, pick up the service vehicle, and drive to the first job site. Remote access creates a win win situation for the company, worker, and customer. The utility company realizes increased worker productivity, reduced transportation costs, and reduced building and land costs. The worker eliminates commuting time and expense, while customers obtain faster, more responsive service, and lower rates.

Sales and marketing organizations are especially reliant on remote access capabilities.

The use of remote access capabilities and laptop computers enables salespeople to complete contracts and obtain real time technical sales support while being face to face with customers, meeting customer needs and resolving buyer objections through a single sales call and resulting in more successful sales and shorter sales cycles.

3.2.3 Intra corporate Core Connectivity

Business process reengineering programs and application of Enterprise Resource Planning (ERP), such as SAP, succeed by eliminating barriers to communications across departmental boundaries and by replacing slow paperwork procedures with shared electronic databases. These management practices and the associated computer software require reliable, high speed, and secure communications among all employees. The same high level of communications connectivity is required at all of the enter prise's establishments. This setup typically requires that small offices and branch offices be upgraded to the higher standards more commonplace at large headquarters locations. The payoff for successful ERP implementation is an order of magnitude reduction in cycle times, increased flexibility and responsiveness, and sharp reductions in IT overhead costs.

3.2.4 Closed User Groups with Partners, Customers, and Suppliers

Some of the most dramatic improvements in business processes are obtained by eliminating certain sub processes entirely. The supply chain is one business process where big improvements are being realized. For example, Boeing suppliers are required to participate in its supply network. This enables Boeing to eliminate stores and parts costs entirely by moving those functions back into the supplier's operation. Similar successes have been achieved in sales and marketing. In another example, Saturn customers can step through the entire sales process online. Saturn reduces selling costs and provides prospective customers with full and accurate examination of options and features, independent of high pressure sales people.

Highly technical sales organizations can create lock in relationships with their customers through creation of closed user groups. For example, semiconductor manufacturers provide online engineering design tools so that circuit designers can incorporate the manufacturer's chips directly into finished designs. Closed user groups not only assure product loyalty, they also provide value to circuit designers by reducing

cycle times.

3.2.5 Public Internet Access

Essentially all functional areas can benefit from public Internet access. Accounting organizations retrieve forms and advice from federal, state, and local revenue offices. Human resources organizations use the Monster Board for recruiting. Mechanical designers can peruse online parts catalogs and download CAD/CAM drawings directly into their blue prints. Energy marketers buy and sell natural gas through Internet based trading systems and retrieve weather data from government and private sources. Pension fund managers follow the financial markets and retrieve stock holder information from company Web pages. IT professionals stay ahead of industry developments and product releases by studying computer and software vendors' online product literature. The business benefit of most of this activity is faster and better informed decision making.

3.2.6 Internet Based Customer Interaction

Retail sales and service companies operate on thin operating margins. Their success depends on executing transactions rapidly and at low cost while giving the customer the appearance of custom tailored service this is sometimes referred to as mass customization. Industries such as air lines, utilities, banks, and brokerage, insurance, and mail-order retailers know that market segmentation, customer loyalty, and low transaction costs are the keys to their success (or survival). Of course, the more time customer service representatives spend with customers and the more they can learn about customers, the better the market segmentation and the customer relationship. Unfortunately, this tender loving care costs money and drives up transaction costs.

Well designed Internet based customer interaction systems resolve this dilemma by eliminating customer service staffing costs and simultaneously providing customers with many custom choices. Information provided by the customer during these online sessions flows directly to the enterprise's data warehouse and is used by data mining tools to further refine the market segmentation models. Brokerage and financial services firms are especially effective at using the Internet to drive down small lot trading fees and eliminate the cost of account representatives. For example, a trade of 100 shares that

once cost several hundred dollars can be done on the Internet for $10. As another example, airlines, including United Airlines, provide Web pages where customers can shop for the best price and schedule, and book their travel over the Internet.

3.2.7 Web Presence

The public Internet is rapidly replacing mass media, including television, radio, and print, as the vehicle for certain product and institutional advertising. While practically all businesses feel compelled to have a Web page, it is essential in many industry segments. Use of Web pages is firmly entrenched in the IT industry itself, financial services, education, and government services. The key item these enterprises share in common is a need for dissemination of large quantities of time sensitive information to millions of people. While these enterprises gain high value from rapid and cheap dissemination of information through Internet Web pages, they also face large risks. Incorrect or false information could destroy the public trust that was built up over decades. Slow information access or unreliable access could create an image of ineptitude or unresponsiveness, damaging institutional loyalty and trust. Failure to safeguard customer data and protect privacy could, at best, destroy trust and, at worst, cause financial ruin. Thus, a Web presence can be effective in reaching the mass market, but security and reliability must be assured.

3.2.8 Getting Real Business Value from Virtual Private Networks

The preceding section describes six ways data communications can be used to produce business value. However, today's data communications networks are failing to deliver the value, because they are too complex and costly. VPNs provide more efficient and secure data communications at a fraction of the cost of today's network architectures. In particular, VPNs reduce the administrative effort and costs of building and operating private networks. This is particularly true as customers, suppliers, and third parties are added to the network. Figure shows the emerging VPN architecture.

One difference between the VPN architecture and today's private net work architecture is that the VPN architecture is seamless. Users in each enterprise, regardless of whether their location is at headquarters or on a wireless link, obtain the same access and logical view of services, despite being served by a number of ISPs and through different physical media. Another difference between the private data communications network and the VPN is that business users never see the network complexity, and net

work administrators are freed from complex network engineering tasks.

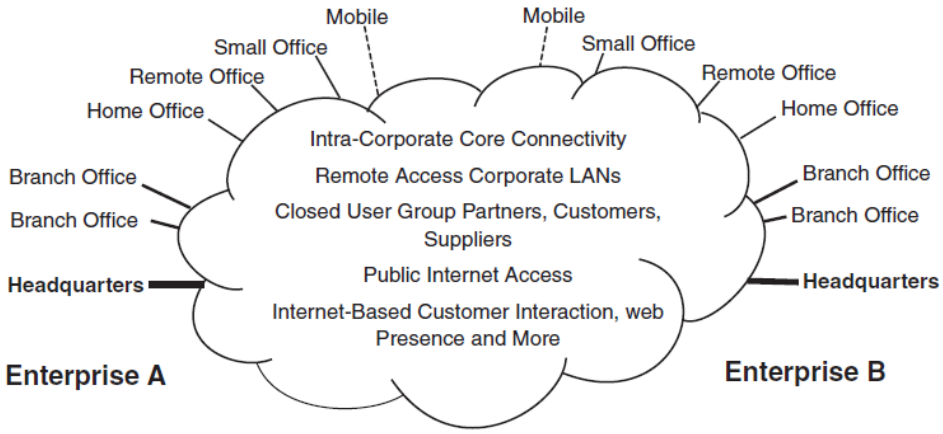

Figure 44 Emerging VPN architecture.

3.2.9 Virtual Service Management

Many of today's VPNs have focused only on providing a secure transport, the network "plumbing." But in practical terms, the benefits of VPN have been off limits to smaller businesses and organizations with limited IT staff and resources, because of the technical complexity of setting up and administering a VPN. Virtual Services Management (VSM) is critical to making a VPN easy to administer and manage across multiple locations and services.

The administrative challenge of creating and maintaining a VPN is formidable. A single enterprise often must accommodate headquarters, campuses, branch offices to home offices, and users who want to use a range of applications and services, and have specific accessing privileges and options. In addition, modern management practice requires many additional links to suppliers, customers, and third parties, as well as access to the public Internet with its 100 million computers.

Through a single point of administration anywhere on the network (local or remote), VSM technology simplifies the administrative burden of setting up multiple branch office email, Web, firewall, and other user services; multiple domain and user names; and coordination among multiple ISPs. It also simplifies the administrative burden through automatic synchronization of software upgrades, replication of Web servers, and sophisticated policy management. VSM overcomes the barrier to private network implementation and VPN that could previously be addressed by only a handful of the largest, more technically sophisticated enterprises.

VSM can help resellers by making it easy for them to add services without raising

the level of technical support they will need to provide. This can be done with service providers or as a standalone value added feature.

Similarly, service providers can take advantage of VSM and VPNs to provide a value added network feature to their customers. VPN services are typically provided on a monthly fee basis and often require customers to perform the network configuration and route determination for their VPN. Where the customer is doing much of the work already, customers often acquire the lines and build a VPN network using CPE products such as an all in one Internet system (described below). Many enterprises are finding if they partner with their service provider to produce a VPN solution, it can be a very effective way to take development costs out of the equation.

3.3 SECURE AND RELIABLE NETWORKING TRANSPORT

To provide secure and reliable transport across the network, three main issues must be resolved

1. Overall network security
2. Wide area network tunneling
3. Class of service and quality of service

Products and standards are in place to provide overall network security while emerging standards will soon resolve the other two issues.
Four functions are key to overall network security

1. Authentication - verify the identity of the user
2. Authorization - verify which services the user is allowed to access
3. Accounting - create an audit trail of the user's network activity
4. Encryption- protect data privacy

These four functions are typically provided by access control lists in routers that restrict access to data packets and network segments in both directions. Firewalls provide more sophisticated control of incoming and outgoing packets at the network's edge. Authentication and authorization is provided by services such as PAP or CHAP and by security servers. Proxy application servers and the network operating system provide additional network security. These necessary services and products are now widely deployed in ISPs and private networks.

Wide area network tunneling is a technique that establishes a secure network

connection across the public Internet. Trade press articles sometimes equate VPNs to tunneling. Our view is that tunneling, while an essential ingredient of the VPN solution is but one element of the VPN and that administrative and reliability issues are at least as important to successful VPN adoption. Major networking vendors have advanced proposed tunneling standards such as Point to Point Tunneling Protocol and L2F. Much marketplace confusion has resulted from these competing standards. Happily, it appears that a compromise approach called L2TP will resolve the differences between these competing standards and will soon emerge from the IETF standards setting process.

3.3.1 IMPLEMENTING THE VPN

The key to deploying a VPN is to give the appearance of a seamless net work with identical user services at all locations headquarters, branch offices, home offices, and those of partners, suppliers, and customers. One approach to VPN implementation for small and medium sized organizations is to deploy all in one Internet systems, sometimes called "Internet edge servers," at the network edge between each enterprise site and the local ISP. The all in one Internet system integrates Internet server, firewall, and networking functionality for organizations that want to take greater advantage of the Internet without adding a complex and costly assembly of boxes and IT staff.

VSM capabilities supported by the system can then provide single point administration of VPN services. Figure shows how the three versions of VSM technology in branch, remote, and extranet applications can be used.

A multi branch VPN can be used to connect a company's remotely located, LAN attached offices. An all in one Internet system will be required at each office, in this case. Class of Service policies, such as access privileges and priorities, can be applied as if the branch users were physically located at headquarters. Security can be implemented through the emerging industry standard IP Security (IPSec) protocol, which will provide DES encryption, authentication, and key management.

A remote VPN can enable mobile workers and telecommuters to dial into a local ISP to access corporate information and service, making it appears as if they were sitting at their desks in the main office. An all in one Internet system will be required at headquarters and Microsoft's Point to Point Protocol (PPTP), available with MS Windows clients; will be required on the remote user's desktop or laptop system. A Point to Point Protocol server in the system can authenticate the remote user, and then open an

encrypted path through which traffic flows as if through the LAN.

An extranet VPN opens a corporate network selectively to suppliers, customers, strategic business partners, and users having access to a limited set of information behind the corporate firewall. An extranet VPN implementation differs from branch and remote VPN implementations in that its use is likely to involve temporary virtual networks which may be set up for specific projects and dismantled as the project's end.

It is important that all necessary service management, security, and Quality of Service functions are combined in the system so that multiple systems can be administered as though they are on a single local network. The supported services should include all of the administrative, security, and reliability requirements of the VPN

Hardware costs can also be minimized because all the necessary administrative, security, and reliable transport functionalities are combined in a single unit. Administrative and operating expenses can be controlled through VSM, which permits management of all sites from a single point minimizing the need for costly data communications experts.

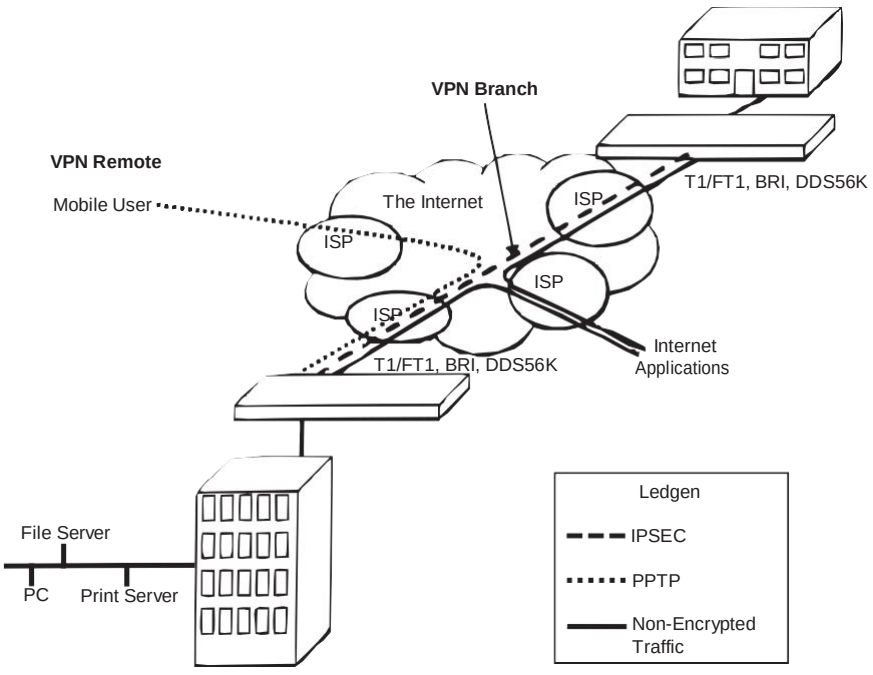

Figure 45 The three versions of VSM technology.

- IP router

- Web server
- firewall
- email
- file transfer (FTP)
- Domain Name Service (DNS)
- Dynamic Host Configuration Protocol (DHCP)
- remote management

3.3.2 SECURE REMOTE ACCESS OVER THE INTERNET

THE COMPONENTS AND RESOURCES OF ONE NETWORK OVER ANOTHER NETWORK are connected via a Virtual Private Network (VPN). As shown in Figure, VPNs accomplish this by allowing the user to tunnel through the Internet or another public network in a manner that lets the tunnel participants enjoy the same security and features formerly available only inprivate networks.

Using the routing infrastructure provided by a public internetwork (such as the Internet), VPNs allow telecommuters, remote employees like salespeople, or even branch offices to connect in a secure fashion to an enterprise server located at the edge of the enterprise local area network (LAN). The VPN is a point to point connection between the user's computer and an enterprise server from the user's perspective. It also appears as if the data is being sent over a dedicated private link because the nature of the intermediate internetwork is irrelevant to the user. As previously mentioned, while maintaining secure communications, VPN technology also allows an enterprise to connect to branch offices or to other enterprises (extranets) over a public internetwork (such as the Internet). The VPN connection across the Internet logically operates as a wide area net work (WAN) link between the sites. In both cases, the secure connection across the internetwork appears to the user as a private network communication (despite the fact that this communication occurs over a public internetwork); hence the name Virtual Private Network.

VPN technology is designed to address issues surrounding the current enterprise trend toward increased telecommuting, widely distributed global operations, and highly interdependent partner operations. Here, workers must be able to connect to central resources and communicate with each other. And, enterprises need to efficiently manage inventories for just in time production.

An enterprise must deploy a reliable and scalable remote access solution to provide employees with the ability to connect to enterprise computing resources regardless of their location. Enterprises typically choose one of the following

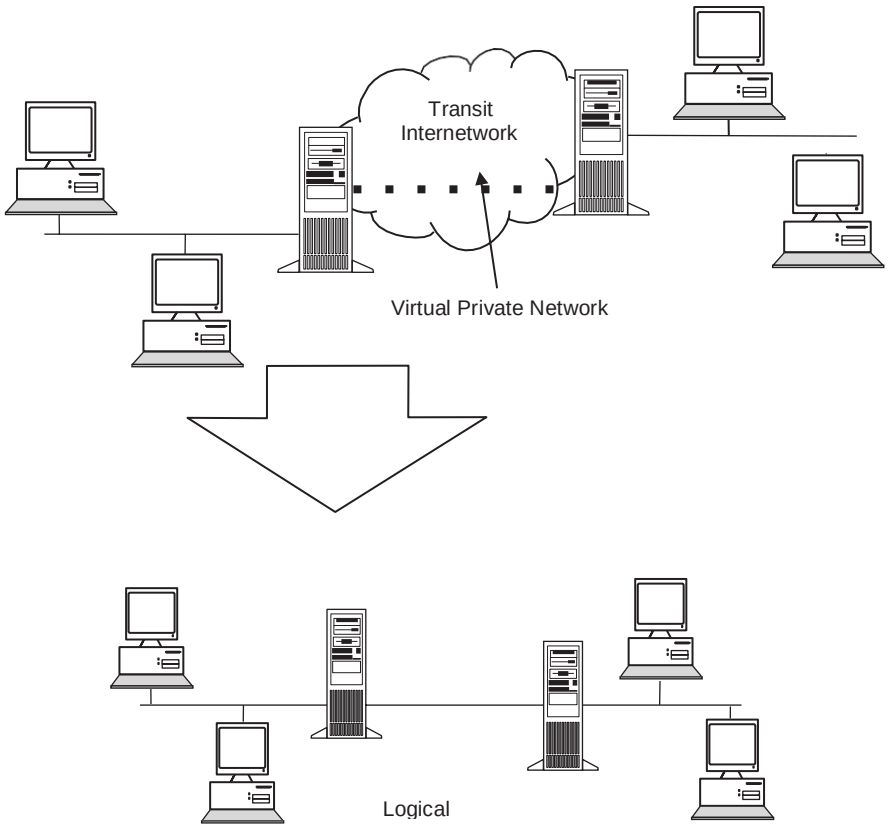

Figure 45 Virtual Private Network.

- an IT department driven solution, where an internal information systems department is charged with buying, installing, and maintaining enterprise modem pools and a private network infrastructure
- value added network (VAN) solutions, where an enterprise pays an outsourced enterprise to buy, install, and maintain modem pools and a Telco infrastructure

The optimum solution in terms of cost, reliability, scalability, flexible administration and management, and demand for connections is provided by neither of

these traditional solutions. Therefore, it makes sense to find a middle ground where the enterprise either supplements or replaces its current investments in modem pools and its private network infrastructure with a less expensive solution based on Internet technology. In this manner, the enterprise can focus on its core competencies with the assurance that accessibility will never be compromised, and that the most economical solution will be deployed. The availability of an Internet solution enables a few Internet connections (via Internet service providers, or ISPs) and deployment of several edge of network VPN server computers to serve the remote networking needs of thousands or even tens of thousands of remote clients and branch offices, as described next.

3.3.3 VPN Common Uses

The next few subsections of this chapter describe in more detail common VPN situations.

3.3.4 Secure Remote User Access over the Internet

While maintaining privacy of information, VPNs provide remote access to enterprise resources over the public Internet. A VPN that is used to connect a remote user to an enterprise intranet is shown in Figure. The user first calls a local ISP Network Access Server (NAS) phone number, rather than making a leased line; long-distance (or 1800) call to an enterprise or outsourced NAS. The VPN software creates a virtual private network between the dialup user and the enterprise VPN server across the Internet using the local connection to the ISP.

3.3.5 Connecting Networks over the Internet.

To connect local area networks at remote sites, there exist two methods for using VPNs using dedicated lines to connect a branch office to an enterprise LAN, or a dialup line to connect a branch office to an enterprise LAN.

Using Dedicated Lines to Connect a Branch Office to an Enterprise LAN. Both the branch office and the enterprise hub routers can use a local dedicated circuit and local ISP to connect to the Internet, rather than using an expensive long haul dedicated circuit between the branch office and the enterprise hub. The local ISP connections and the public Internet are used by the VPN software to create a virtual private network between the branch office router and the enterprise hub router.

Using a DialUp Line to Connect a Branch Office to an Enterprise LAN. The router at the branch office can call the local ISP, rather than having a router at the branch office make a leased line, long distance or (1800) call to an enterprise or outsourced NAS. Also, in order to create a VPN between the branch office router and the enterprise hub router across the Internet, the VPN soft ware uses the connection to the local ISP as shown in figure

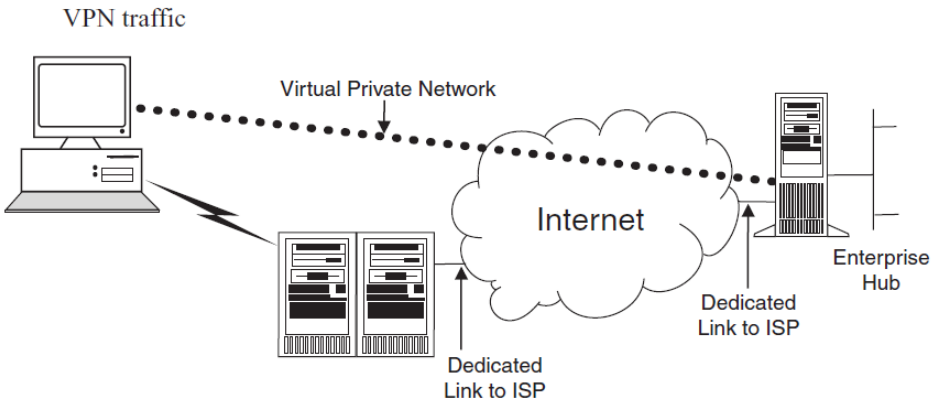

Figure 46 Using a VPN to connect a remote client to a private LAN.

The facilities that connect the branch office and enterprise offices to the Internet are local in both cases. To make a connection, both client/server, and server/server VPN cost savings are largely predicated on the use of a local access phone number. It is recommended that the enterprise hub router that acts as a VPN server be connected to a local ISP with a dedicated line. This VPN server must be listening 24 hours per day for incoming VPN traffic.

3.3.6 Connecting Computers over an Intranet

The departmental data is so sensitive that the department's LAN is physically disconnected from the rest of the enterprise internetwork in some enterprise internetworks. All of this creates information accessibility problems for those users not physically connected to the separate LAN, although the department's confidential information is protected.

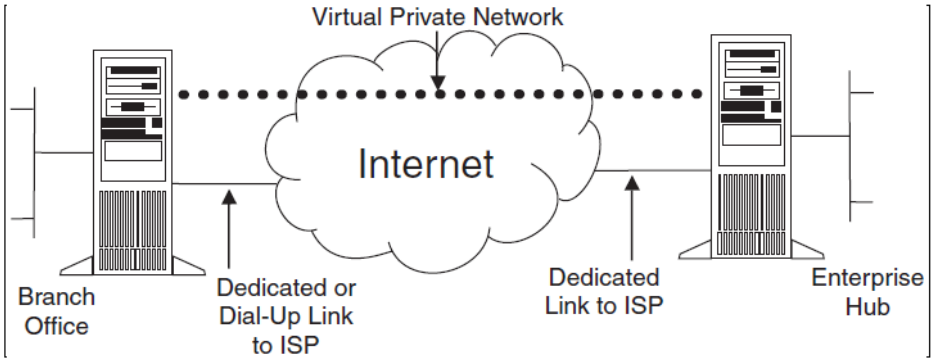

Figure 47 Using a VPN to connect two remote sites.

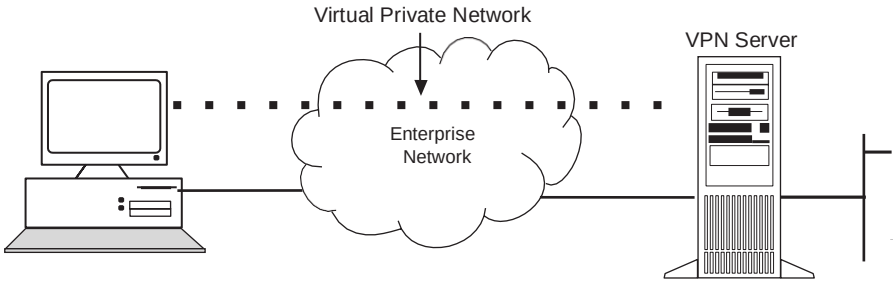

Secured or Hidden Network

Figure 48 Using a VPN to Connect to two computers on the same LAN.

VPNs allow the department's LAN to be separated by a VPN server but physically connected to the enterprise internetwork. One should note that the VPN server is not acting as a router between the enterprise internetwork and the department LAN. A router would interconnect the two networks, thus allowing everyone access to the sensitive LAN. The network administrator can ensure that only those users on the enterprise internetwork who have appropriate credentials (based on a need to know policy within the enterprise) can establish a VPN with the VPN server and gain access to the protected resources of the department by using a VPN. Additionally, all communication across the VPN can be encrypted for data confidentiality. Thus, the department LAN cannot be viewed by those users who do not have the proper credentials.

3.3.7 BASIC VPN REQUIREMENTS

Normally, an enterprise desires to facilitate controlled access to enterprise resources and information when deploying a remote networking solution. In order to

easily connect to enterprise local area network (LAN) resources, the solution must allow freedom for authorized remote clients. And, in order to share resources and information (LAN to LAN connections), the solution must also allow remote offices to connect to each other. Finally, as the data traverses the public Internet, the solution must ensure the privacy and integrity of data. Also, in the case of sensitive data traversing an enterprise internetwork, the same concerns apply. A VPN solution should therefore provide all of the following at a minimum

- Address management the solution must assign a client's address on the private net, and must ensure that private addresses are kept private
- Data encryption data carried on the public network must be rendered unreadable to unauthorized clients on the network
- Key management the solution must generate and refresh encryption keys for the client and server
- Multiprotocol support the solution must be able to handle common protocols used in the public network; these include Internet Protocol (IP), Internet Packet Exchange (IPX), etc.

- User authentication the solution must verify a user's identity and restrict VPN access to authorized users; in addition, the solution must provide audit and accounting records to show who accessed what information and when

Furthermore, all of these basic requirements are met by an Internet VPN solution based on the Point to Point Tunneling Protocol (PPTP) or Layer 2 Tunneling Protocol (L2TP). The solution also takes advantage of the broad availability of the worldwide Internet. Other solutions meet some of these requirements, but remain useful for specific situations, including the new IP Security Protocol (IPSec).

3.3.8 Point to Point Tunneling Protocol (PPTP)

PPTP is a Layer 2 protocol that encapsulates PPP frames in IP datagram for transmission over an IP internetwork, such as the Internet. PPTP can also be used in private LAN to LAN networking.

PPTP is documented in the draft RFC, "Point to Point Tunneling Proto col."[1] This draft was submitted to the IETF in June 1996 by the member enterprises of the PPTP Forum, including Microsoft Corporation, Ascend Communications, 3Com/Primary

Access, ECI Telematics, and U.S. Robotics (now 3Com).

The Point to Point Tunneling Protocol (PPTP) uses Generic Routing Encapsulation (GRE) encapsulated Point to Point Protocol (PPP) frames for tunneled data and a TCP connection for tunnel maintenance. The payloads of the encapsulated PPP frames can be compressed as well as encrypted. How a PPTP packet is assembled prior to transmission is shown in Figure. The illustration shows a dialup client creating a tunnel across an internetwork. The encapsulation for a dialup client (PPP device driver) is shown in the final frame layout.

3.3.9 Layer 2 Forwarding (L2F)

L2F (a technology proposed by Cisco Systems, Inc.) is a transmission protocol that allows dialup access servers to frame dialup traffic in PPP and transmit it over WAN links to an L2F server (a router). The L2F server then unwraps the packets and injects them into the network. Unlike PPTP and L2TP, L2F has no defined client.

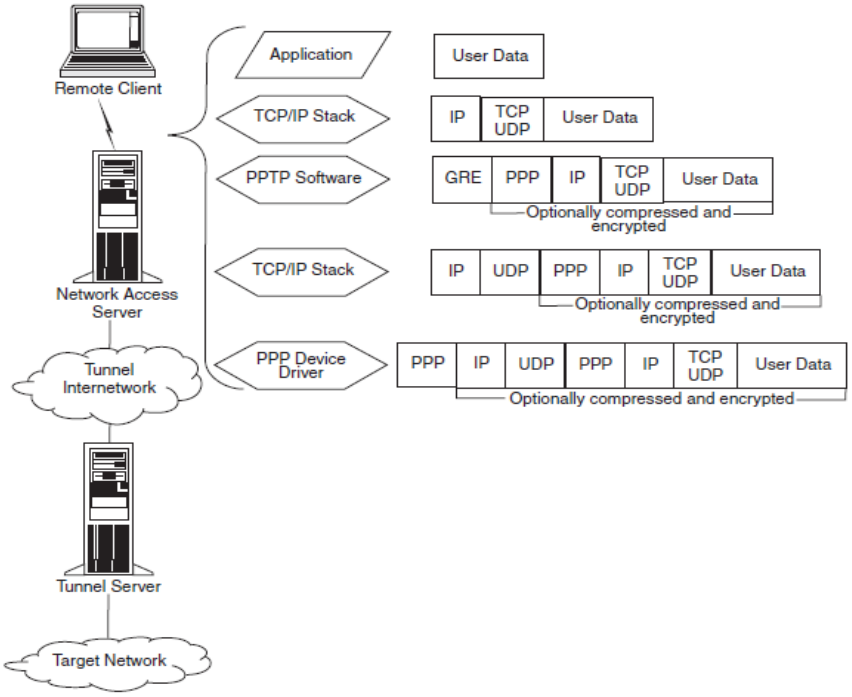

Figure 49 Construction of a PPTP packet.

3.3.10 Layer 2 Tunneling Protocol (L2TP)

A combination of PPTP and L2F makes up L2TP. In other words, the best features of PPTP and L2F are incorporated into L2TP.

L2TP is a network protocol that encapsulates PPP frames to be sent over Asynchronous Transfer Mode (ATM), IP, X.25, or frame relay networks. L2TP can be

used as a tunneling protocol over the Internet when configured to use IP as its datagram transport. Without an IP transport layer, L2TP can also be used directly over various WAN media (such as Frame Relay). L2TP is documented in the draft RFC, Layer 2 Tunneling Protocol "L2TP" (draftietfpppextl2tp09.txt). This draft document was submitted to the IETF.

For tunnel maintenance, L2TP over IP internetworks uses UDP and a series of L2TP messages. As the tunneled data, L2TP also uses UDP to send L2TPencapsulated PPP frames. The payloads of encapsulated PPP frames can be compressed as well as encrypted. How an L2TP packet is assembled prior to transmission is shown in Figure . A dialup client creating a tunnel across an internetwork is shown in the Figure. The encapsulation for a dialup client (PPP device driver) is shown in the final frame layout. L2TP over IP is assumed in the encapsulation.

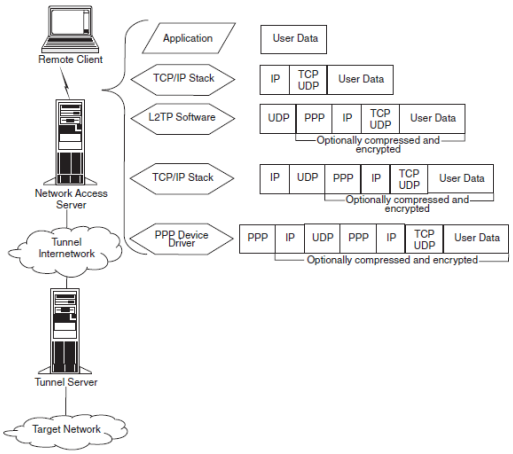

Figure 50 **Construction of an L2TP packet**

3.3.11 L2TP Compared to PPTP. PPP is used to provide an initial envelope for the data for both PPTP and L2TP. Then, it appends additional headers for transport through the internetwork. The two protocols are very similar. There are differences between PPTP and L2TP, however. For example,

- L2TP provides for header compression. When header compression is enabled, L2TP operates with four bytes of overhead, as compared to six bytes for PPTP.
- L2TP provides for tunnel authentication, while PPTP does not. However, when either protocol is used over IPSec, tunnel authentication is provided by IPSec so that Layer 2 tunnel authentication is not necessary.

- PPTP can only support a single tunnel between endpoints. L2TP allows for the use of multiple tunnels between endpoints. With L2TP, one can create different tunnels for different qualities of service.
- PPTP requires that the internetwork be an IP internetwork. L2TP requires only that the tunnel media provide packet oriented point to point connectivity. L2TP can be used over IP (using UDP), Frame Relay permanent virtual circuits (PVCs), X.25 virtual circuits (VCs), or ATM VCs.

3.4 INTERNET PROTOCOL SECURITY (IPSEC) TUNNEL MODE

The secured transfer of information across an IP internetwork is supported by IPSec (a Layer 3 protocol standard). Nevertheless, in the context of tunneling protocols, one aspect of IPSec is discussed here. IPSec defines the packet format for an IP over an IP tunnel mode (generally referred to as IPSec Tunnel Mode), in addition to its definition of encryption mechanisms for IP traffic. An IPSec tunnel consists of a tunnel server and tunnel client. These are both configured to use a negotiated encryption mechanism and IPSec tunneling.

For secure transfer across a private or public IP internetwork, IPSec Tunnel Mode uses the negotiated security method (if any) to encapsulate and encrypt entire IP packets. The encrypted payload is then encapsulated again with a plaintext IP header. It is then sent on the internetwork for delivery to the tunnel server. The tunnel server processes and discards the plain text IP header and then decrypts its contents to retrieve the original payload IP packet. Upon receipt of this datagram, the payload IP packet is then processed normally and routed to its destination on the target network. The following features and limitations are contained within the IPSec Tunnel Mode

- It is controlled by a security policy a set of filter matching rules. This security policy establishes the encryption and tunneling mechanisms available in order of preference and the authentication methods available, also in order of preference. As soon as there is traffic, the two machines perform mutual authentication, and then negotiate the encryption methods to be used. Thereafter, all traffic is encrypted using the negotiated encryption mechanism and then wrapped in a tunnel header.
- It functions at the bottom of the IP stack; therefore, applications and higher level protocols inherit its behavior.
- It supports IP traffic only.

- The remainder of this article discusses VPNs and the use of these technologies by enterprises to do secure remote access (e.g., by traveling employees and sales reps) over the Internet in greater detail.

3.4.1 EASY TO MANAGE AND USE

While squeezing the maximum possible from budget and support staffs, today's enterprises are asking their information technology groups (ITGs) to deliver an increasing array of communication and networking services. It appears that the situation is no different at Microsoft Corporation (Redmond, Washington). The Microsoft ITG needed to provide secure, Internet based remote access for its more than 35,000 mobile sales personnel, telecommuters, and consultants around the world.

Microsoft's ITG is currently using and deploying a custom Windows based remote dialup and virtual private networking (VPN) solution by using Windows based clients and enhanced Windows 2000® RAS (Remote Access Server) technology available in the Windows 2000 Option Pack (formerly named Windows NT 5.0). Users are given quick, easy, and low cost network access. Additional user services are provided with new Windows based network services from UUNet Technologies, Inc.[3]

3.4.2 Integrated RASVPN Clients

According to Microsoft, its ITG has learned that the widespread adoption and use of technology largely depends on how easy and transparent the experience is for the end user. Likewise, Microsoft's ITG has learned not to deploy technologies for which complexity results in an increased support burden on its limited support staff. Microsoft's ITG provided a single client interface with central management to simultaneously make the remote access solution easy to use and manage.

3.4.2.1 Single Client. A single client is used for both the direct dial up and virtual private network connections. Users utilize the same client interface for secure transparent access, whether dialing directly to the enterprise network or connecting via a VPN, by using Windows integrated dialup net working technology (DUN) and Microsoft Connection Manager. In fact, users do not need to concern themselves with which method is employed.

3.4.2.2 Central Management. Central management is used for remote dialup and VPN access phone numbers. According to Microsoft, its ITG has found that one of the most common support problems traveling users face is deter mining and managing local access phone numbers. This problem translates into one of the principal reasons for support calls to Microsoft's user support centers. Using the Connection Manager Administration Kit (CMAK) wizard (which is part of Microsoft's remote access solution), Microsoft's ITG preloads each client PC with an electronic phone book that includes every dialup remote access phone number for Microsoft's net work. The Windows solution also allows phone books to be centrally integrated and managed from a single remote location, and clients to be updated automatically.

3.4.3 WINDOWS COMMUNICATION PLATFORM

In order to provide a flexible and comprehensive network solution, the open extensibility of the Windows 2000 allows Microsoft's ITG to preserve its current hardware network investments while partnering with UUnet Technologies, Inc. According to Microsoft, the Windows platform enabled its ITG to integrate the best of breed network services and applications to best meet its client and network administration needs.

3.4.3.1 High Speed Internet Access on the Road

Microsoft employees can also connect to high speed Internet access by plugging into public IPORT[4] jacks in hotels, airports, cafes, and remote locations. The Microsoft ITG integrates the IPORT[5] pay per use Internet access features into its custom remote access solution. According to Microsoft, this high bandwidth, easily available connection helps Microsoft employees be more productive and have a better online experience while on the road.

3.4.4 Secure Internet Access and VPN

Microsoft's ITG, like its counterparts at every enterprise, must ensure that the edge of its network is secure while still providing all employees with the freedom needed to access information worldwide. Microsoft's ITG has also deployed Microsoft Proxy

Server to securely separate the LAN from the Internet to meet this need.

To ensure that no intruders compromise the edge of network, the Microsoft Proxy Server firewall capabilities protect Microsoft's network from unauthorized access from the Internet by providing network address translation and dynamic IPlevel filtering. Microsoft's ITG uses the powerful caching services in Microsoft Proxy Server to expedite the delivery of information at the same time.

The Proxy Server is able to service subsequent user requests of already requested information without having to generate additional network traffic by reusing relevant cached information. In addition, in order to operate at peak efficiency with the utmost security, ITG uses Microsoft Proxy Server to enable the Microsoft intranet and remote employees.

3.4.5 RAS Reporting and Internal Usage Chargeback (Billing)

Microsoft pays a substantial amount for remote access fees due to the need to maintain private leased lines and dedicated 800 numbers like many large enterprises with a multitude of branch offices and remote employees. In addition, according to Microsoft, the sheer number of LAN entry points and autonomy afforded its international divisions made centralized accounting and retail reporting for remote access use and roaming users important.

Microsoft's ITG is deploying a VPN solution — bolstered with centralized accounting and reporting of enterprise wide remote access and VPN use — by using Windows 2000, integrated user domain directory, and RADIUS services. As part of this solution, Microsoft is also deploying TRU RADIUS Accountant™ for Windows 2000 from TelcoResearch.[6]

Furthermore, Microsoft's ITG is also able to generate detailed reporting of remote access and VPN network use for internal cost accounting purposes while using familiar Windows 2000 management tools by using Telco Research's product. In addition, Microsoft's ITG is able to quickly and easily deploy a turnkey reporting solution built on the intrinsic communication services of Windows 2000 in this manner. According to Microsoft, while maintaining the flexibility to accommodate future change, they receive better security as a result, reduced implementation costs, and enhanced reporting to improve remote access management and charge back service.

3.4.6 VIP Services Economical Internet Access and VPN

By working with UUnet Technologies, Inc. (the largest Internet service provider in the world), the Microsoft ITG supplemented its private data network infrastructure and RAS with VPN services. Microsoft's VPN solution is integrated with the UUnet Radius Proxy servers through the Windows 2000 native support for RADIUS under this relationship.

Through the Windows 2000 Remote Access Service integrated RADIUS support, Microsoft's ITG made reliable and secure local access to UUnet Technologies IP network available to all Microsoft mobile employees. This resulted in the delivery of high quality VPN services over the UUnet Technologies, Inc. infrastructure at a reduced cost. The ITG conservatively estimates that this use of VPN service as an alternative to traditional remote access will save Microsoft more than $7 million per year in remote access fees alone. Additional savings are expected from the elimination of call requests for RAS phone numbers and greatly reduced remote access con figuration support.

The ITG utilized the integrated support for RADIUS based authentication available from the Windows Directory in Windows 2000. This allowed them to retain all existing authentication rights for both Internet and LAN access, avoiding change or redundant replication of directory, and provided for enhanced network security.

According to Microsoft, their ITG was able to instantly extend network access to its more than 50,000 employees in more than 100 countries through its relationship with UUnet Technologies. So that Microsoft employees can access information locally anywhere with reliability guarantees and the support of UUnet, UUnet Technologies' transcontinental backbone provides access throughout North America, Europe, and the Asia–Pacific region.

PLANNING FOR THE FUTURE

Finally, Microsoft's ITG wanted to ensure that its current investment in the remote access infrastructure would not only be able to meet today's needs, but also enable it to make the most of opportunities provided by the digital convergence of network aware applications in the near future. Evidence of an increased need for higher degrees of client/server network application integration is found in the momentum of Windows 2000 as a platform for IP telephony, media streaming technologies, and the migration to PBX systems based on Windows 2000.

The flexibility needed to economically address current and future needs of

Microsoft's ITG is provided through the use of Windows 2000 as the backbone of the remote access solution. Through partnerships with multiple service providers such as UUnet Technologies, the selection of a Windows based solution allows ITG the freedom to both centrally manage and incrementally extend the Microsoft direct dial and VPN infrastructure at a controlled pace and in an open manner.

In order to connect Microsoft subsidiaries, branch offices, and extranet partners securely to the enterprise network over private and public net works, Windows 2000 Routing, RAS, and VPN services — along with tight integration with Microsoft Proxy Server — are already enabling Microsoft's ITG to seamlessly extend its RAS–VPN infrastructure. Furthermore, to meet Microsoft's enterprise needs into the future, the broad application support enjoyed by the Windows communication platform ensures that ITG will continue to have access to a host of rich application services made available by developers and service providers, such as ATCOM, Inc., Telco Research, and UUnet Technologies, Inc.

3.5 IPSec VPNS

VPNS ARE MAKING A HUGE IMPACT ON THE WAY COMMUNICATIONS ARE VIEWED. They are also providing ample fodder for administrators and managers to have seemingly endless discussions about various applications. On one side are the possible money savings, and the other are implementation issues. There are several areas of serious concern

- performance
- interoperability
- scalability
- flexibility

3.5.1 Performance

Performance of data flow is typically the most common concern, and IPSec is very processor intensive. The performance costs of IPSec are the encryption being performed, integrity checking, packet handling based on policies, and forwarding, all of which become apparent in the form of latency and reduced throughput. IPSec VPNs over the Internet increase the latency in the communication that conspires with the processing costs to discourage VPN as a solution for transport sensitive applications. Process time

for authentication, key management, and integrity verification will produce delay issues with SA establishment, authentication, and IPSec SA maintenance. Each of these results in poor initialization response and, ultimately, disgruntled users.

The application of existing hardware encryption technology to IPSec vendor products has allowed these solutions to be considered more closely by prospective clients wishing to seize the monetary savings associated with the technology. The creation of a key and its subsequent use in the encryption process can be offloaded onto a dedicated processor that is designed specifically for these operations. Until the application of hard ware encryption for IPSec, all data was managed through software computation that was also responsible for many other operations that may be running on the gateway.

Hardware encryption has released IPSec VPN technology into the realm of viable communication solutions. Unfortunately, the client operating system participating in a VPN is still responsible for the IPSec process. Publicly available mobile systems that provide hardware based encryption for IPSec communications are becoming available, but are sometime away from being standard issue for remote users.

3.5.2 Interoperability

Interoperability is a current issue that will soon become antiquated as vendors recognize the need to become fully IPSec compliant — or consumers will not implement their product based simply on its incompatibility. Shared secret and ISAKMP key management protocol are typically allowing multivendor interoperability. As Certificate Authorities and the technology that supports them become fully adopted technology, they will only add to the cross platform integration. However, complex and large VPNs will not be manageable using different vendor products in the near future. Given the complexity, recentness of the IPSec standard and the various interpretations of that standard, time to complete interoperability seem great.

3.5.3 Scalability

Scalability is obtained by the addition of equipment and bandwidth. Some vendors have created products focused on remote access for roaming users, while others have concentrated on network to network connectivity without much attention to remote users.

The current ability to scale the solution will be directly related to the service required. The standard supporting the technology allows for great flexibility in the addition of services. It will be more common to find limitations in equipment configurations than in the standard as it pertains to growth capabilities. Scalability ushers in a wave of varying issues, including

- authentication
- management
- performance

Authentication can be provided by a number of processes, although the primary focus has been on RADIUS, Certificates, and forms of two factor authentication. Each of these can be applied to several supporting databases. RADIUS is supported by nearly every common authenticating system from Microsoft Windows NT to Net Ware's NDS. Authentication, when implemented properly, should not become a scalability issue for many implementations, because the goal is to integrate the process with existing or planned enterprise authenticating services.

A more interesting aspect of IPSec vendor implementations and the scalability issues that might arise is management. As detailed earlier, certain implementations do not scale, due to the shear physics of shared secrets and manual key management. In the event of the addition of equipment or increased bandwidth to support remote applications, the management will need to take multiplicity into consideration. Currently, VPN management of remote users and networks leaves a great deal to be desired. As vendors and organizations become more acquainted with what can be accomplished, sophisticated management capabilities will become increasingly available.

Performance is an obvious issue when considering the increase of an implementation. Typically, performance is the driving reason, followed by support for increased numbers. Both of these issues are volatile and inter related with the hardware technology driving the implementation. Performance capabilities can be controlled by the limitation of supported SAs on a particular system a direct limitation in scalability. A type of requested encryption might not be available on the encryption processor currently available. Forcing the calculation of encryption onto the operating system ultimately limits the performance. A limitation may resonate in the form of added equipment to accomplish the link between the IPSec equipment and the authenticating database. When users authenticate, the granularity of control over the capabilities of that user may be

directly related to the form of authentication. The desired form of authentication may have limitations in various environments due to restrictions in various types of authenticating databases. Upgrade issues, service pack variations, user limitations, and protocol requirements also combine to limit growth of the solution.

3.5.4 THE MARKET FOR VPN

Several distinct qualities of VPN are driving the investigation by many organizations to implement VPN as a business interchange technology. VPNs attempt to resolve a variety of current technological limitations that represent themselves as costs in equipment and support or solutions where none had existed prior. Three areas that can be improved by VPNs are

- remote user access and remote office connectivity
- extranet partner connectivity
- internal departmental security

3.5.4.1 Remote Access

Providing remote users access via a dialup connection can become a costly service for any organization to provide. Organizations must consider costs for

- telephone lines
- terminating equipment
- long distance
- calling card
- 800/877 number support

Telephone connections must be increased to support the number of proposed simultaneous users that will be dialing in for connectivity to the net work. Another cost that is rolled up into the telephone line charge is the possible need for equipment to allow the addition of telephone lines to an existing system. Terminating equipment, such as modem pools, can become expenses that are immediate savings once VPN is utilized. Long distance charges, calling cards that are supplied to roaming users, and toll free lines require initial capital and continuous financial support. In reality, an organization employing conventional remote access services is nothing more than a service provider for their employees. Taking this into consideration, many organizations tend to overlook

the use of the Internet connection by the remote users. As the number of simultaneous users access the network, the more bandwidth is utilized for the existing Internet service.

The cost savings are realized by redirecting funds, originally to support telephone communications, in an Internet service provider (ISP) and its ability to support a greater area of access points and technology. This allows an organization to eliminate support for all direct connectivity and focus on a single connection and technology for all data exchange — ultimately saving money. With the company access point becoming a single point of entry, access controls, authenticating mechanisms, security policies, and system redundancy is focused and common among all types of access regardless of the originator's communication technology.

The advent of high speed Internet connectivity by means of cable modems and ADSL (Asynchronous Digital Subscriber Line) is an example of how VPN becomes an enabler to facilitate the need for high speed, individual remote access where none existed before. Existing remote access technologies are generally limited to 128K ISDN (Integrated Services Digital Network), or more typically, 56K modem access. Given the inherent properties of the Internet and IPSec functioning at the network layer, the communication technology utilized to access the Internet only needs to be supported at the immediate connection point to establish an IP session with the ISP. Using the Internet as a backbone for encrypted communications allows for equal IP functionality with increased performance and security over conventional remote access technology.

Currently, cable modem and ADSL services are expanding from the home user market into the business industry for remote office support. A typical remote office will have a small frame relay connection to the home office. Any Internet traffic from the remote office is usually forwarded to the home office's Internet connection, where access controls can be centrally managed and Internet connection costs are eliminated at the remote office. However, as the number of remote offices and the distances increase, so does the financial investment. Each frame relay connection, PVC (Permanent Virtual Circuit), has costs associated with it. Committed Information Rate (CIR), port speed (e.g., 128K), and sometimes a connection fee add to the overall investment. A PVC is required for any connection; so as remote offices demand direct communication to their peers, a PVC will need to be added to support this decentralized communication. Currently within the United States, the cost of frame relay is very low and typically outweighs the cost of an ISP and Internet connectivity. As the distance increases and moves beyond the United States, the costs can increase exponentially and will typically

call for more than one telecommunications vendor. With VPN technology, a local connection to the Internet can be established. Adding connectivity to peers is accomplished by con figuration modifications; this allows the customer to control communications without the inclusion of the carrier in the transformation.

The current stability of remote, tier three and lower ISPs is an unknown variable. The arguable service associated with multiple and international ISP connectivity has become the Achilles' heel for VPN acceptance for business critical and time critical services. As the reach of tier one and tier two ISPs increases, they will be able to provide contiguous connectivity over the Internet to remote locations using an arsenal of available technologies.

3.5.4.2 Extranet Access

The single, most advantageous characteristic of VPNs is to provide protected and controlled communication with partnering organizations. Years ago, prior to VPN becoming a catchword, corporations were beginning to feel the need for dedicated Internet access. The dedicated access is becoming utilized for business purposes, whereas before it was viewed asa service for employees and research requirements.

The Internet provides the ultimate bridge between networks that was relatively nonexistent before VPN technology. Preceding VPNs, a corporation needing to access a partner's site was typically provided a frame relay connection to a common frame relay cloud where all the partners claimed access. Other options were ISDN and dial on demand routing. As this requirement grows, several limitations begin to surface. Security issues, partner support, controlling access, disallowing unwanted interchange between partners, and connectivity support for partners without supported access technologies all conspire to expose the huge advantages of VPNs over the Internet. Utilizing VPNs, an organization can maintain a high granularity of control over the connectivity per partner or per user on a partner network.

Internal Protection

As firewalls became more predominant as protection against the Internet, they were increasingly being utilized for internal segmentation of departmental entities. The need for protecting vital departments within an organization originally spawned this concept of using firewalls internally. As the number of departments increase, the management, complexity, and cost of the firewalls increase as well. Also, any attacker

with access to the protected network can easily obtain sensitive information due to the fact that the firewall applies only perimeter security.

VLANs (Virtual Local Area Networks) with access control lists became a minimized replacement for conventional firewalls. However, the same security issue remained, in that the perimeter security was controlled and left the internal network open for attack.

As IPSec became accepted as a viable secure communication technology and applied in MAC environments, it also became the replacement for other protection technologies. Combined with strategically placed fire walls, VPN over internal networks allows secure connectivity between hosts. IPSec encryption, authentication, and access control provide protection for data between departments and within a department.

3.5.5 CONSIDERATION FOR VPN IMPLEMENTATION

The benefits of VPN technology can be realized in varying degrees depending on the application and the requirements it has been applied to. Considering the incredible growth in technology, the advantages will only increase. Nevertheless, the understandable concerns with performance, reliability, scalability, and implementation issues must be investigated.

3.5.5.1 System Requirements

The first step is determining the foreseeable amount of traffic and its patterns to ascertain the adjacent system requirements or augmentations. In the event that existing equipment is providing all or a portion of the service the VPN is replacing, the costs can be compared to discover initial savings in the framework of money, performance, or functionality.

3.5.5.2 Security Policy

It will be necessary to determine if the VPN technology and how it is planned to be implemented meets the current security policy. In case the security policy does not address the area of remote access, or in the event a policy or remote access does not exist, a policy must address the security requirements of the organization and its relationship with the service provided by VPN technology.

3.5.5.3 Application Performance

As previously discussed, performance is the primary reason VPN technology is not the solution for many organizations. It will be necessary to determine the speed at which an application can execute the essential processes. This is related to the type of data within the VPN. Live traffic or user sessions are incredibly sensitive to any latency in the communication. Pilot tests and load simulation should be considered strongly prior to large scale VPN deployment or replacement of exiting services and equipment.

Data replication or transient activity that is not associated with human or application time sensitivity is a candidate for VPN connectivity. The application's resistance to latency must be measured to determine the minimum requirements for the VPN. This is not to convey that VPNs are only good for replication traffic and cannot support user applications. It is necessary to determine the application needs and verify the requirements to properly gauge the performance provisioning of the VPN. The performance "window" will allow the proper selection of equipment to meet the needs of the proposed solution; otherwise, the equipment and application may present poor results compared to the expected or planned results. Or, more importantly, the acquired equipment is underworked or does not scale in the direction needed for a particular organization's growth path. Each of these results in poor investment realization and make it much more difficult to persuade management to use VPN again.

3.5.5.4 Training

User and administrator training are an important part of the implementation process. It is necessary to evaluate a vendor's product from the point of the users, as well as evaluating the other attributes of the product. In the event the user experience is poor, it will reach management and ultimately weigh heavily on the administrators and security practitioners. It is necessary to understand the user intervention that is required in the everyday process of application use. Comprehending the user knowledge requirements will allow for the creation of a training curriculum that best represents what the users are required to accomplish to operate the VPN as per the security policy.

3.5.6 FUTURE OF IPSec VPNs

Like it or not, VPN is here to stay. IP version 6 (IPv6) has the IPSec entrenched in its very foundation; and as the Internet grows, Ipv6 will become more prevalent. The

current technological direction of typical net works will become the next goals for IPSec; specifically, Quality of Service (QoS). ATM was practically invented to accommodate the vast array of communication technologies at high speeds; but to do it efficiently, it must control who gets in and out of the network.

Ethernet Type of Service (ToS) (802.1p) allows for three bits of data in the frame to be used to add ToS information and then be mapped into ATM cells. IP version 4, currently applied, has support for a ToS field in the IP Header similar to Ethernet 802.1p; it provides three bits for extended information. Currently, techniques are being applied to map QoS information from one medium to another. This is very exciting for service organizations that will be able sell end to end QoS. As the IPSec standard grows and current TCP/IP applications and networks begin to support the existing IP ToS field, IPSec will quickly conform to the requirements.

The IETF and other participants, in the form of RFCs, are continually addressing the issues that currently exist with IPSec. Packet sizes are typically increased due to the added headers and sometimes trailer information associated with IPSec. The result is increased possibility of packet fragmentation. IPSec addresses fragmentation and packet loss; the over head of these processes are the largest concern.

IPSec can only be applied to the TCP/IP protocol. Therefore, multiprotocol networks and environments that employ IPX/SPX, NetBEUI, and others will not take direct advantage of the IPSec VPN. To allow non TCP/IP protocols to communicate over an IPSec VPN, an IP gateway must be implemented to encapsulate the original protocol into an IP packet and then be forwarded to the IPSec gateway. IP gateways have been in use for some time and are proven technology. For several organizations that cannot eliminate non TCP/IP protocols and wish to implement IPSec as the VPN of choice, a protocol gateway is imminent.

As is obvious, performance is crucial to IPSec VPN capabilities and cost. As encryption algorithms become increasingly sophisticated and hard ware support for those algorithms become readily available, this currentlimitation will be surpassed.

Another perceived limitation of IPSec is the encryption export and import restrictions of encryption. There are countries that the United States places restrictions on to hinder the ability of those countries to encrypt possibly harmful information into the United States. In 1996, the International Traffic in Arms Regulation (ITAR) governing the export of cryptography was reconditioned. Responsibility for cryptography exports was transferred to the Department of Commerce from the

Department of State. However, the Department of Justice is now part of the export review process. In addition, the National Security Agency (NSA) remains the final arbiter of whether to grant encryption products export licenses.

The NSA staff is assigned to the Commerce Department and many other federal agencies that deal with encryption policy and standards. This includes the State Department, Justice Department, National Institute for Standards and Technology (NIST), and the Federal Communications Commission. As one can imagine, the laws governing the export of encryption are complicated and are under constant revision. Several countries are completely denied access to encrypted communications to the United States; other countries have limitations due to government relationships and political posture. The current list of (as of this writing) embargoed countries include

- Syria
- Iran
- Iraq
- North Korea
- Libya
- Cuba
- Sudan
- Serbia

As one reads the list of countries, it is easy to determine why the United States is reluctant to allow encrypted communications with these countries. Past wars, conflict of interests, and terrorism are the primary ingredients to become exiled by the United States.

Similar rosters exist for other countries that have the United States listed as "unfriendly," due to their perception of communication with the United States.

As one can certainly see, the concept of encryption export and import laws is vague, complex, and constantly in litigation. In the event a VPN is required for international communication, it will be necessary to obtain the latest information available to properly implement the communication as per the current laws.

3.6 INTEGRATING DATA CENTERS WITH INTRANETS

NEARLY ALL ENTERPRISES THAT HAVE MAINFRAMES or large, networked AS/400s now have an intranet. Most, in addition, already have a presence on

the Internet in the form of a home page, and many are actively exploring the possibilities of using the Internet for electronic commerce, customer support, and as an ultra cost effective means of global remote access. In parallel, intranet to intranet communication via extranets is being viewed as the means of streamlining and expediting enterprise transactions. Very few enterprises at present have tightly integrated their intra nets with their data centers. This is despite the fact that up to 70 percent of the vital data, and many of the mission critical applications required by these enterprises, are still likely to reside on their mainframes or AS/400s. That is akin to baking an apple pie with no apple filling.

Integrating an intranet with a data center is not simply a matter of implementing TCP/IP on a mainframe or AS/400 along with a Web server. Many of the host resident, mission critical applications still required were developed, typically 15 years ago, such that they only work in Systems Network Architecture mode. The nearest that one can come to making these applications TCP/IP compatible is to use them in conjunction with a host resident or "Off Board" tn3270(E) (or tn5250, in the case of AS/400s) server which will perform standards based SNA to TCP/IP protocol conversion. Otherwise, the applications will have to be rewritten to work in TCP/IP mode. This is not feasible since the cost and effort of doing so for the $20 trillion installed base of SNA mission critical applications would make all the tribulations associated with the Y2K challenge appear trivial.

While some of the data center resident data could be accessed using an Open Database Connectivity type scheme, this is certainly not true for all of the data center resources. Some data, especially if stored on "flat files" or non relational databases (such as IBM's still widely used Information Management System), can only be accessed via SNA applications. In other instances, the data make sense only when combined with the "business logic" embedded within an SNA mission critical application. In addition to these crucial SNA applications, there is inevitably a large installed base of SNA only "legacy" devices such as IBM 4700 Financial Systems, automated teller machines, and control units that still need to be supported. Thus, there is a need for explicit SNA related technologies in order to get the most from your host intranet.

The good news is that highly proven and stable technology from more than 40 credible vendors including IBM, Cisco, Attachmate, Open Connect Systems, Wall Data, Eicon, Novell, WRQ, Farabi, Client/Server Technology, Sterling Software, Blue Lobster,

etc., is now readily available to facilitate data center to intranet integration in a seamless and synergistic manner. Enterprises around the world such as GM, FedEx, Ohio State University, Royal Jordanian Airlines, Nestles, The Chickering Group, National Van Lines, the State of Idaho, Al Rajhi Banking & Investment Corp. (Saudi Ara bia's largest bank), and Gazprom (a $30 billion natural gas company in Rus sia) are already gainfully using this intranet to data center integration technology on a daily basis for business critical production use. Al Rajhi Bank, for example, uses browser based access to SNA to provide home banking, while GM, National Van Lines, Royal Jordanian Airlines, and The Chickering Group use it to permit agents to access applications or databases resident on mainframes or AS/400s over the Internet.

3.6.1 INTRANET TO DATA CENTER INTEGRATION TECHNOLOGIES

To be viable, integration technologies need to be able to accommodate an extremely broad and disparate population of client equipment and functionality including PCs, UNIX workstations, coax attached 3270/5250 terminals, printers, minicomputers, SNA applications that communicate program to program using LU 6.2 or LULU Session Type 0based protocols, SNA only devices, SNALAN gateways (e.g., NetWare for SAA), and legacy control units. The PCs, workstations, and printers may work in either SNA or TCP/IP mode. Consequently, you will need SNA Access technologies to deal with TCP/IP clients, particularly PCs and workstations, and SNA Transport technologies to deal with SNA only clients. The most pertinent technologies are

- SNA Access technologies that permit non SNA clients to gain access to SNA applications
- ip3270/ip5250 — the use of existing PC/workstation SNA emulators (e.g., Attachmate EXTRA! Personal Client) and existing SNALAN gateways (e.g., Microsoft's SNA server) with proprietary encapsulation schemes for conveying a 3270/5250 data stream within TCP/IP
- tn3270(E)/tn5250 — IETF standard that enables TCP/IP clients (e.g., Attachmate EXTRA! Personal Client) to access SNA applications via tn3270(E) (e.g., IBM 2216) or tn5250 servers
- Browser based Access with 3270/5250toHTML Conversion — thin client solution where a server resident SNAWeb gateway performs 3270/5250 data stream to HTML conversion replete with some amount of user interface rejuvenation so that

SNA applications can be accessed directly from a browser.

- Browser invoked Java or ActiveX applets — dynamically download able applets, which can optionally be cached on a PC/workstation hard disk, that provide 3270/5250 emulation either directly or in conjunction with an intermediate SNA Web gateway
- Browser invoked applets as "4" above, but with user interface reju venation
- Application specific web to data center gateways, e.g., IBM's CICS Web Interface or Interlink's ActiveCICX
- Programmatic (or Middleware) Servers, e.g., IBM's MQSeries, Blue Stone's Sapphire/Web, or Blue Lobster's Stingray SDK SNA end to end transport
- Data Link Switching — ubiquitous, standardsbased encapsulation scheme performed by bridge/routers that permits any kind of SNA/APPN traffic, independent of session type, to be transported end toend across a TCP/IP WAN. Desktop DLSw (DDLSw) is also available where SNA traffic can be encapsulated within TPC/IP at the source PC
- High Performance RoutingoverIP — alternative to the DLSw champi oned by IBM, whereby SNAoriented routing is performed across IP
- AnyNet — IBM protocol conversion technology, integrated within IBM server software including Comm. Server/NT and OS/390 as well as within some SNA/3270 emulation packages, that converts SNA message units into corresponding TCP/IP packets

The three transport technologies ensure that the still large installed base of SNA devices and control units are able to communicate with main frame or AS/400resident SNA/APPN applications across an intranet using SNA on an end to end basis. Of the three, standards based DLSw, which is available on nearly all major brands of bridge/routers, is by far the most widely used and the most strategic. AnyNet, in marked contrast, is not available on bridge/routers or within SNA devices such as 3174s, 4700s, etc. Consequently, it cannot be used easily as a universal scheme for supporting any and all SNA devices and control units as can DLSw. Thus, AnyNet is not as strategic or useful as DLSw. High Performance Routing (HPR) is IBM's follow on architecture to APPN and SNA. HPRoverIP, now available on IBM 2216 and CS/NT, has irrefutable advantages over DLSw it can support native, data center to data center SNA/APPN routing over TCP/IP; SNA LU 6.2 Class of Service (COS)based path selection; and traffic

prioritization. If and when this technology is more readily available, corporations that require SNA/APPN routing to obtain optimum traffic routing in multi data center networks, or those that have LU 6.2based applications that rely on COS, may want to consider HPRoverIP as an alternative to DLSw.

DLSw's ability to support any and all types of SNA/APPN traffic effortlessly could be easily abused when trying to integrate intranets with data centers. DLSw could be used all by itself to realize the integration by grafting the existing SNA/APPN network, totally unchanged, onto the intranet through the extensive deployment of DLSw all around the periphery of the intranet. This brute force, "no SNA reengineering whatsoever" approach has been used in the past to integrate SNA networks into TCP/IP networks. With this type of DLSw only network you would find SNALAN gateways being used downstream of the intranet, and then DLSw being used to transport the SNA output of these gateways across the intranet. While such net works indubitably work, there are other strategic techniques such as a 3270toHTML and applet based 3270/5250 emulation that should typically be used in conjunction with DLSw to achieve the necessary integration. Figure summarizes how the various SNA Transport and SNA Access integration techniques can be gainfully synthesized to integrate data centers with intranets.

3.7 IMPLEMENTING AND SUPPORTINGEXTRANETS

EXTRANETS HAVE BEEN AROUND as long as the first rudimentary LAN to LAN networks began connecting two different business entities together to form WANs. In its basic form, an extranet is the interconnection of two previous separate LANs or WANs with origins from different business entities. This term emerged to differentiate between the previous definitions of external "Internet" connection and a company's internal intranet. Figure depicts an extranet as a Venn diagram, where the intersection of two (or more) nets forms the extranet. The network in this intersection was previously part of the "intranet" and has now been made accessible to external parties.

Under this design, one of the simplest definitions comes from R.H. Baker,[1] "An extranet is an intranet that is open to selective access by outside parties." The critical security concept of the extranet is the new net work area that was previously excluded from external access now being made available to some external party or group. The critical security issue evolves from the potential vulnerability of allowing more than the

intended party, or allowing more access than was intended originally for the extra net. These critical areas will be addressed in this article, from basic extranet setup to more complex methods and some of the ongoing support issues.

The rapid adoption of the extranet will change how a business looks at its security practices, as the old paradigm of a hard outer security shell for a business LAN environment has now been disassembled or breached with a hole to support the need for extranets. In many cases, the age-old firewall will remain in place, but it will have to be modified to allow this "hole" for the extranet to enable access to some degree to internal resources that have now been deemed part of the extranet.

Recognizing the growth of extranets as a common part of doing business today is important, and therefore the business enterprise must be ready with architectures, policy, and approaches to handle the introduction of extranets into its environment. A few of the considerations are the requirements versus security balance, policy considerations, risk assessments, and implementation and maintenance costs.

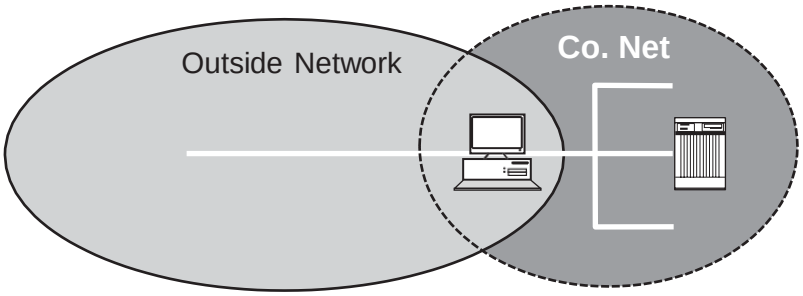

Figure 51 Extranet Venn diagram.

Recognizing the growth of extranets as a common part of doing business today is important, and therefore the business enterprise must be ready with architectures, policy, and approaches to handle the introduction of extranets into its environment. A few of the considerations are the requirements versus security balance, policy considerations, risk assessments, and implementation and maintenance costs.

From requirements versus security balance standpoint, the issue is the initial claim by business that extranets are an immediate need and absolutely must be established "if we are to remain competitive." But from a security standpoint, such a drastic change to the environment, which may not have had any form of an extranet in place, may well be throwing their financial data assets out the door with the first implementation of an extranet. Therefore, care must be taken from a security perspective and put in balance with the claimed business need for an extranet implementation.

One of the first areas of review and (possibly) update is the inner company's security policy. This policy most likely was not written with extra nets in mind and thus may need modification if a common security philosophy is to be established regarding how a company can securely implement extranets. However, the policy review does not stop with one company's review of its own policy, but also includes connecting the company or companies on the outside. In the case of strategic business relationships that will be ongoing, it is important that both parties fully understand each other's responsibilities for the extranet, what traffic they will and will not pass over the joined link — what degree of access, and by whom, will occur over this link.

Part of any company's policy on extranets must include an initial requirement for a security risk assessment, the main question being what additional levels of risk or network vulnerability will be introduced with the implementation of the proposed extranet? As well as vulnerability assessment, a performance assessment should be conducted to assist in the design of the extranet to ensure that the proposed architecture not only addresses the security risk, but that it also will meet performance expectations. Some of the questions to be asked in a combined security and performance assessment should be

- data classification/value of data
- data location(s) in the network
- internal users' access requirements to extranet components (internal access design)
- data accessibility by time of day (for estimating support costs)
- protocol, access services used to enter extranet (network design implications)
- degree of exposure by transmission mechanism (Internet, private net, wireless transmission)
- end user environment (dialup, Internet)
- number of users, total/expectation for concurrent users access (line sizing)
- growth rate of user base (for estimating administrative costs)
- CONUS (continental U.S.), international access (encryption implications)

The risk and performance assessment would, of course, be followed by a risk mitigation plan, which comes in the form of selecting an acceptable extranet architecture and identifying the costs. The cost aspect of this plan is, of course, one of the critical drivers in the business decision to implement an extranet. Is the cost of implementing and maintaining the extranet (in a secure manner) less than the benefit gained by putting

the extranet in place? This cost must include the costs associated with implementing it securely; otherwise, the full costs will not be realistically reflected.

Finally, the company implementing the extranet must have a clear set of architectures that best mitigate the identified vulnerabilities, at the least cost, without introducing an unacceptable degree of risk into its computing environment. The following section reviews various extranet architectures, each with differing costs and degrees of risk to the environment.

3.7.1 EXTRANET ARCHITECTURES

3.7.1.1 Router Based Extranet Architecture

The earliest extranet implementations were created with network routers that have the capability to be programmed with rudimentary "access control lists" or rules. These rules were implemented based solely on TCP/IP addresses. A rule could be written to allow External User A access to a given computer B, where B may have been previously unreachable due to some form of private enterprise network firewall (and in the early days, this firewall may have been a router also). Figure depicts this very basic extranet. A more realistic rule can be written where all computers in an "outside network" are allowed to access computer B in a company network, thus forming an extranet. This is depicted in Figure.

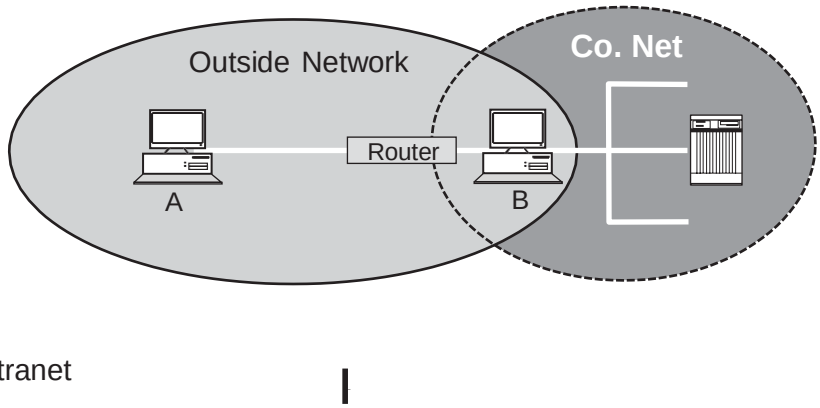

Figure 52 Basic extranet with router

As network security architectures matured, routers as the sole network access control device were replaced by more specific security mechanisms. Routers were originally intended as network devices and not as security mechanisms and lost

functionality as more and more security rules were placed in them. Additionally, the security rules that were put into them were based on TCP/IP addresses, which were found to be subject to spoofing/masquerading and thus deemed ineffective in positively identifying the real external device being granted access. Therefore, routers alone do not provide an entirely secure extranet implementation; but when used in conjunction with one of the following extranet architectures, routers can be a component to add some degree of security, but only when used in conjunction with other network security devices.

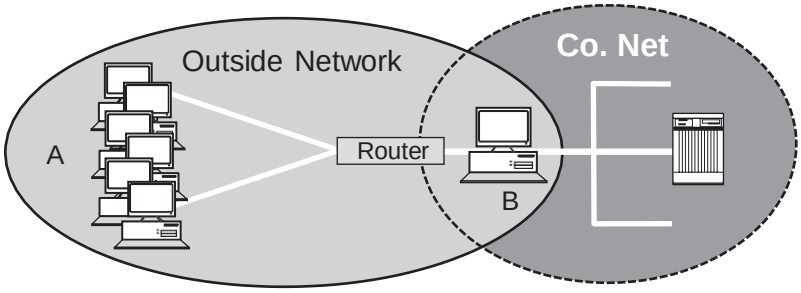

Figure 53 More realistic extranet

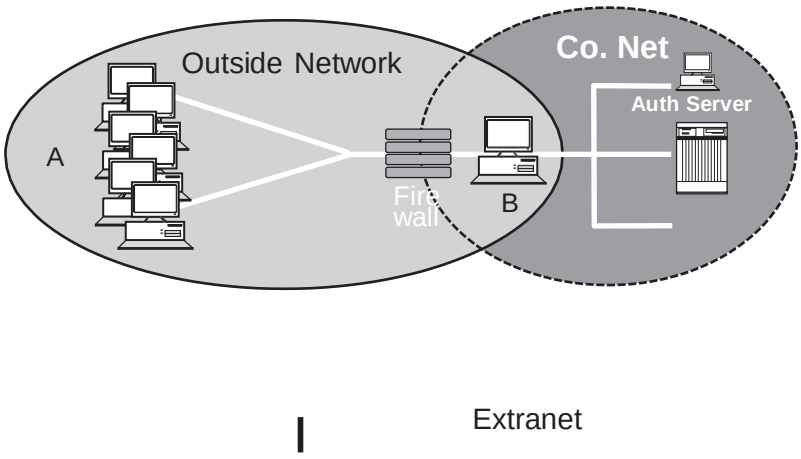

Figure 54 Extranet using an application layer gateway firewall.

Application Gateway Firewalls

As network security architectures matured, the introduction of application layer

gateway firewalls, software tools on dedicated machines, usually dual homed (two network interfaces, one internal, one external), became the more accepted external protection tool. These software tools have the ability to not only perform router type functions with access control rules, but also provide user authentication services on a per user basis. This user authentication can take the form of an internal user authentication list, or an external authentication call to token based authentication services, such as the ACE Secure ID™ system. Figure depicts this type of architecture setup to support an extranet using an a application layer gateway firewall to enable authenticated users inward access to an enterprise in a controlled manner.

In addition to supporting access control by IP address and user, some gateways have the further capability to restrict access by specific TCP/IP service port, such as Port 80, HTTP, so the extranet users can only access the internal resource on the specific application port and not expose the internal machine to any greater vulnerability than necessary.

Follow on application layer gateway implementations have since emerged to provide varying additional degrees of extranet connectivity and security. One such method is the implementation of a proxy mechanism from an outside network to a portion of an internal company network. Normally, a proxy performs control and address translation for access from an intranet to the external Internet. These types of proxies normally reside on the firewall, and all user access to the Internet is directed through the proxy. The proxy has the ability to exert access control over who in the intranet is allowed external access, as well as where they can go on the Internet. The proxy also provides address translation such that the access packet going

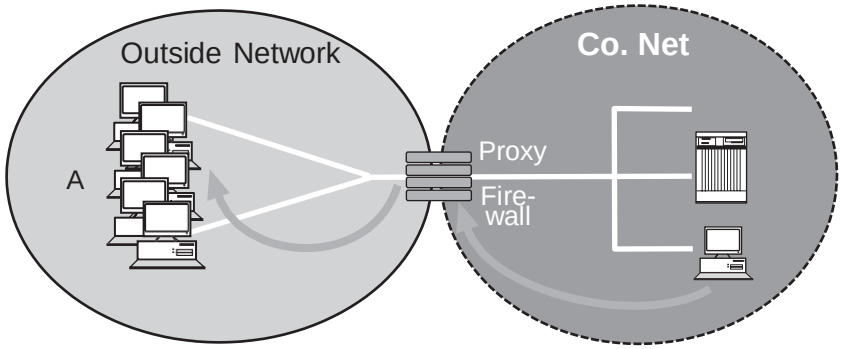

Figure 54 Outbound proxy architecture.

to the Internet is stripped of the user's original internal address, and only the external

gateway address of the enterprise is seen on the packet as it traverses the Internet. Figure depicts these proxy functions.

The proxy provides both security and network address functions, although the entire process can be used in its reverse to provide an extranet architecture because of its ability to provide access rules over who can use the proxy, where these proxy users are allowed to go, and what resources they can access. Figure depicts a reverse proxy extranet architecture.

Today, most proxies are set up for HTTP or HTTPS access, although application layer gateway proxies exist for most popular Internet access services (Telnet, FTP, SQL, etc.). One of the major issues with proxy servers, however, is the amount of cycle time or machine overhead it takes to manage many concurrent proxy sessions through a single gateway. With highly scalable hardware and optimized proxy software, it can be carried

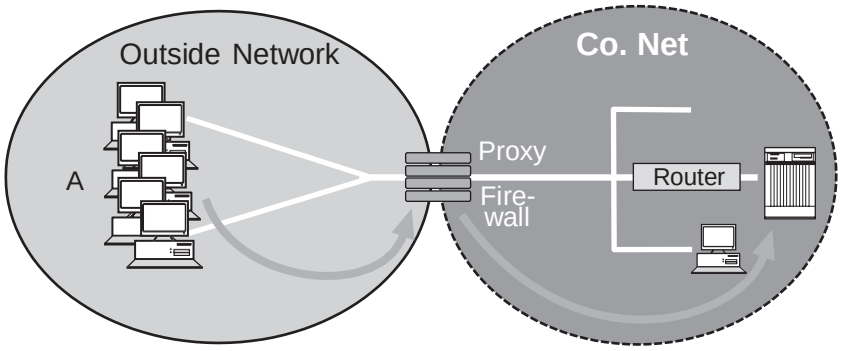

Figure 54 Reverse proxy extranet architecture

to potentially handle high user demands, but the system architecture must be specifically designed for high loads to be able to meet user response expectations while still providing the security of an authenticated proxy architecture. On the inward proxy depicted in Figure, the proxy can be configured to only allow access to a single internal resource on a given TCP/IP port. Further protection can be added to this reverse proxy architecture by putting the target internal resource behind a router with specific access control rules, limiting the portion on the company intranet that inbound proxies can reach, which can ensure limited access on the intra net; should the internal machine ever be compromised, it cannot be used as a “jumping off point” into the rest of company

intranet.

A somewhat hybrid architecture extranet, where some firewall controls are put in place but the external user is not granted direct inward access to an enterprise's internal domain, has been evolving and put in place as a more popular extranet implementation. In this architecture, the external user is granted access to an external resource (something outside of the enterprise firewall), but still on the property of the enterprise. Then, this external resource is granted access to one or more internal resources through the enterprise firewall. This architecture is based on minimizing the full external access to the intranet, but still makes intranet based data available to external users. The most popular implementation is to place an authenticating Web server outside the firewall and program it to make the data queries to an internal resource on the enterprise intranet, over a specific port and via a specific firewall rule, allowing only that one external resource to have access to the one internal resource, thus reducing the external exposure of the intranet. Figure depicts this type of extranet.

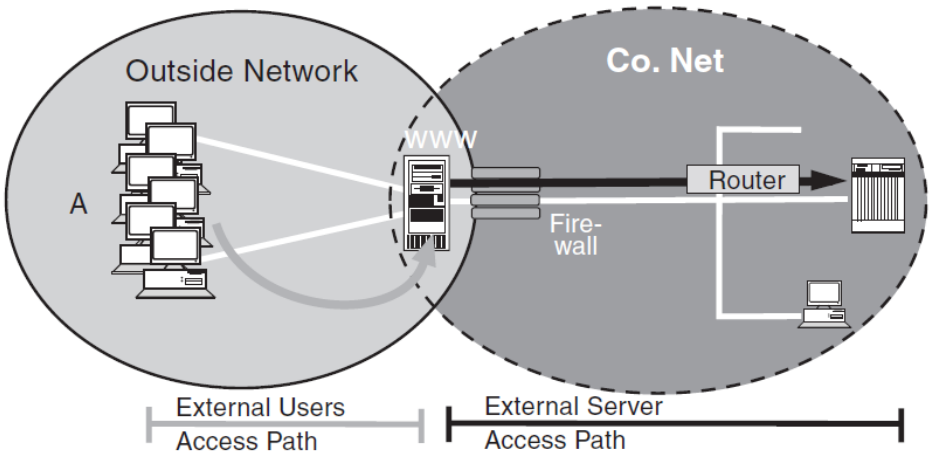

Figure 55 Extranet with authenticating Web server.

Issues with this type of architecture include reliance on a single user interface that can be safely placed outside the enterprise firewall, which makes it vulnerable to attack. Additionally, there is the issue of whether tight enough access rules can be placed on the access method between the external user interface resource (the Web server, in this example) and the internal resources that it needs access to on the protected enterprise intra net. If these two issues can be safely addressed, then this form of extranet can be

very useful for an enterprise extranet with a high volume or varied user base and a large intranet based data repository.

The user front end has been deployed as a Web server, usually SSL enabled to ensure data integrity and protection by encrypting the data as it passes over an external SSL link. Access to this external server is also associated with some form of user authentication, either a static ID or password over the SSL link, and more recently with client digital certificates where each individual accessing the SSL enabled site is issued his own unique digital certificate from an acknowledged certificate authority, thereby validating his identity. Each client maintains its own digital certificate, with the Web server having some record of the public key portion of the client's digital certificate, either directly in the Web server internally, or accessible from a standalone directory server (usually LDAP reachable).

The most recent entrant in the extranet architecture arena is the Virtual Private Network (VPN). This architecture is based on a software tunnel established between some external entity, either client or external net work, and a gateway VPN server. Figure depicts both types of VPN architectures. External Network A has a VPN server at its border which encrypts all traffic targeted for Company Network C; this is a gateway to gateway VPN. Or, External Client B may have client VPN software on his workstation which would enable him to establish a single VPN tunnel from his workstation over the external network to Company C's VPN server.

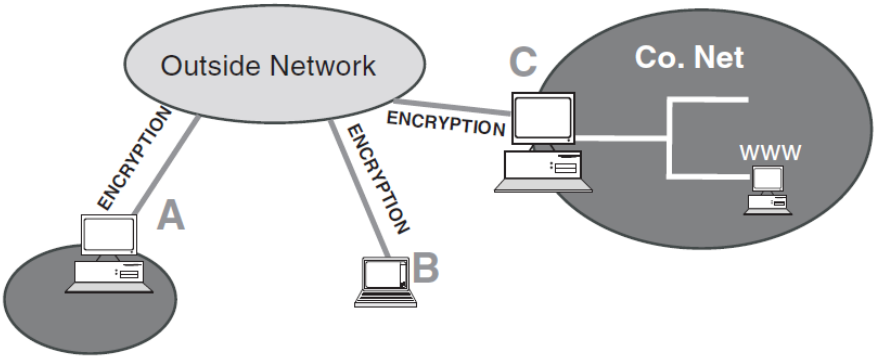

Figure 56 VPN architectures

Although both server to server VPN and client to server VPN architectures are offered in the industry today, it is this author's experience that the more popular extranet architect is the client to server VPN architecture, as it offers the most flexibility

for the most diverse audience of external users. This flexibility does add to the complexity of the implementation, as it can involve a potentially large number of external desktops, all with differing configurations. The benefits of VPNs include the ability to safely traverse external public networks with some assurance of data integrity and authentication as part of the VPN implementation. This architecture shows the most promise to meet the needs of extranets and cost savings for a world hungry for connectivity over public/external networks, although it still has some growing pains to go through to reach full product maturity.

An emerging standard for VPNs is coming out of the ITEF IPSec implementation, which draws a roadmap for the next generation TCP/IP security protocol. Under this protocol, standards are being drafted that will enable differing devices to securely communicate under an agreed upon security protocol, including key exchange for encryption and standardized authentication. Today, there are IPSec compliant products on the market; however, the standard is still evolving and tests are being conducted to evaluate differing vendor compatibilities with each under the IPSec standard. One of the leading initiatives to evaluate this compliance is the Automotive Network Exchange (ANX) test, which is intended to establish a large extranet environment between the core automotive manufacturers and their vendors.

In the meantime, there are a wide variety of VPN product vendors on the market some touting IPSec compliance and others, with proprietary implementations with IPSec in their future product roadmaps, choosing to wait until the standard stabilizes. The recommendation is to either select a vendor offering IPSec if it has some degree of maturity within its own product line, or one that is planning on adopting the standard; IPSec appears to be a viable standard once it fully matures.

Regardless of what VPN solution is being considered for implementing secure extranets, a few technical considerations must be understood and planned for before selecting and implementing a VPN extranet architecture.

Scalability. Similar to proxy servers, VPN servers incur a fair amount of processing overhead that consumes processing resources as high levels of concurrent VPN sessions pass through a single server. It is important to attempt to estimate one's projected user base and current access to appropriately size a VPN server. Some servers are on established lower level processors for smaller environments and should not be implemented where high concurrent access rates are expected, although there is some benefit to physical load balancing spreading the access among multiple servers.

However, there is also concern about implementing too many servers to manage easily. A balance between installing a single large server and creating a single point of failure versus implementing many smaller servers creates an administrative nightmare.

Multi homed Intranets and Address Translation. In large intranet environments, many operate under a split DNS (domain naming structure) where intranet addresses are not "advertised" to the external networks, and external addresses are kept external so as not to flood the internal net work. Additionally, many larger intranet environments have multiple gate ways to external networks. If one of the gateways is established with a VPN gateway and an external client makes a connection to the internal intranet, it is important that the tunnel comes in through the appropriate VPN gate way, but also that the return traffic goes back out through that same gate way so that it gets re encrypted and properly returned to the external VPN client. Figure depicts the correct traffic patterns for a multi homed intranet with a single VPN gateway and an external VPN client.

VPN Based Access Control. Many forms of gateway VPN servers offer the ability to restrict user access to a company intranet based on access groupings. This is especially important when intranets are being established for a diverse set of external users and it is important to minimize user access to the intranet. This type of access control is, of course, critical in establishing secure extranets, which further highlights the importanceof understanding VPN access control capabilities.

User Authentication. Multiple options exist for user authentication, although the recommended option is to select a high level authentication method (e.g., onetime passwords) or a time synchronized password method. Under the IPSec standard, client side digital certificates are evolving as a standard for high level authentication. Unfortunately, initial implementations of client side digital certificates for user authentication are entirely software based, eliminating the second factor authentication, the something the user physically has" in their possession. The return to true two factor authentication under digital certificates will not really occur until physical smart cards become part of the authentication architecture. (Smart cards are credit card type tokens that have a physically embedded chip which can be read electronically and written to, either with a portion of the client's digital certificate or the encryption algorithm used to unlock the digital certificate.)

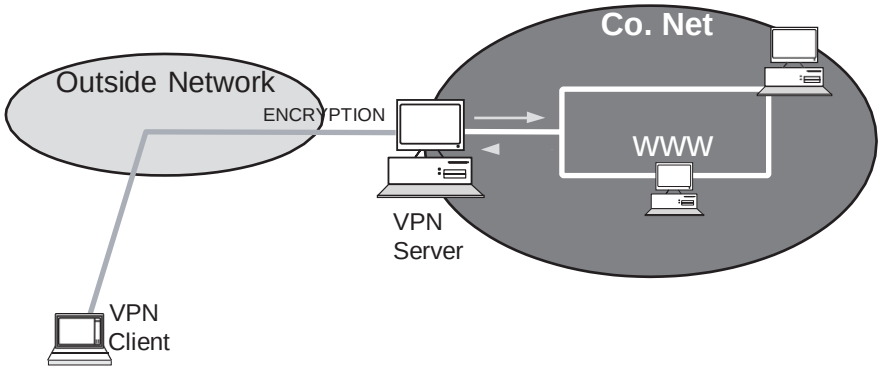

Figure 57 Traffic patterns for multi homed intranet with a single VPN gateway and an external VPN client.

IPSec Interoperability. Ultimately, the IPSec standard will stabilize, and all vendors following the established standard will allow different vendors' VPN products to interoperate. Under this environment, a company can implement a vendor's VPN server, and their acknowledged clients can purchase and use an IPSec compliant client to gain access to the company intranet once they are authorized.

TEXT/ REFERENCE BOOKS

1. Gil Held, "Network Design: Principles and Applications (Best Practices)", Auerbach Publications, 1st edition, 2000.
2. Diane Tiare and Catherine Paquet, "Campus Network Design Fundamentals", Pearson Education, 1st edition, 2006.
3. Larry L. Peterson, Bruce S. Davie, "Computer Networks: A Systems Approach", Morgan Kaufmann Publishers Inc., 5th edition, 2012.
4. William Stallings, "Data and Computer Communications", Pearson Education, 8th edition, 2016.
5. James F. Kurose, Keith W. Ross, "Computer Networking - A Top-Down Approach Featuring the Internet", Pearson Education, 6th edition, 2012.

PART-A - 2 Mark Questions

S. No	Questions (2marks)	CO	Level
1	VPN uses a public network to transmit private datagrams. Illustrate how Data Confidentiality is maintained in VPN	CO3	2
2	Summarize the need for Encryption in VPN's.	CO3	2
3	Interpret the need for Tunnelling in VPN's	CO3	5
4	Justify the need for VPN network Management	CO3	5
5	Categorize the different protocols involved in VPN	CO3	4
6	Illustrate the different requirements to install a perfect VPN network?	CO3	2
7	Interpret how Data Encapsulation helps in improving the security of VPN's?	CO3	5
8	In today's scenario we need Intelligent Networks. In this regard Appraise on the working concept of Virtual Private Networks.	CO3	5
9	Comment on the various characteristics of a Remote and Site-to-site Virtual Private Networks	CO3	4
10	Today's business scenario is multi-polarised. Justify how VPN helps in reducing the investment cost and increases the data throughput efficiency with suitable explanations	CO3	2

PART-B - 10 Mark Questions

S. No	Questions (2marks)	CO	Level
1	Data Integrity and User Authentication is a major bottle neck in Virtual Private networks (VPN). As a Network design expert suggest how the above said security issues are addressed in a VPN with necessary explanations and diagrams	CO1	5
2	A large corporate house has branch offices located in different parts of the world. Illustrate how the remote offices can be connected over a public network by the use of virtual tunnels with necessary examples and diagrams	CO1	2
3	A Network engineer is tasked with the creation of a VPN for a small business house with 500 nodes spread across different geographical locations. Help him by suggesting an appropriate VPN based Encryption and Encapsulation procedures. Also comment on the pros and cons of the algorithm used in this process.	CO1	5
4	As a computer network expert your tasked with the identification of a suitable procedure to protect your data over the Internet. Suggest a suitable IP security and firewall mechanism which can be implemented on the virtual network, your planned to create for the company.	CO1	5
5	Recommend a suitable Intranet based Virtual Private network with necessary diagram. Also provide its pros and cons.	CO1	5
6	As a computer design network expert explain the different implementing & supporting procedures involved in Extranet based VPN architecture.	CO1	5

SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF COMPUTER SCIENCE AND ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE ENGINEERING

Computer Networks and Design – SCSA1502

UNIT IV -NETWORKING TOOLS AND TECHNIQUES-SCSA1502

UNIT IV
NETWORKING TOOLS AND TECHNIQUES

Simulation method for designing multimedia networks – Determining remote bridge and router delays – Network baselining as a planning tool.

INTRODUCTION

A Simulation Method for the Design of Multimedia Networks introduces three basic phases of the simulation method used for multimedia network design. In the preparation phase, the author describes how to define goals in measurable terms. Next, in the baseline phase, information concerning data capture and the validation of such data are presented. The third phase of the design process is referred to by the author as the delta phase. During this phase, changes are applied to the baselined network, after which the results are analyzed and summarized through information presented in the three phases. This chapter provides a methodology one can use to simulate a wide variety of networks. The second chapter in this section focuses attention on techniques one can use to determine a number of quantitative metrics involved in client/server computing via bridges and routers. In Gilbert Held's chapter, titled Determining Remote Bridge and Router Delays, the author uses queuing theory to illustrate various performance metrics associated with interconnection geographically separated LANs via remote bridges or routers. Although the author illustrates how to project queuing delays, of far more importance is the presentation of how one can use queuing theory to answer a classic network design problem. That problem is the appropriate selection of a WAN operating rate to interconnect two geographically separated LANs.

THE MODELING PROCESS

Networks that are complicated enough to have multimedia segments are frequently too complex and too large for simple rule-of-thumb calculations. Simulation is the best means of collecting performance data on networks that are in the planning and design stages.

Complexity requires a great deal of organization in data collection, model validation, and analysis of results. A defined process increases confidence in the results by increasing organization and thereby managing complexity.This chapter specifically draws to the reader's attention the following aspects of simulation

- important data that will be required
- means for obtaining this data
- suggestions for modeling the data
- possible interpretations of results
- how the modeling procedure might be segmented into tractable units

Simulation results are highly dependent on the quality of data being input. The large amount of detailed data being handled increases the mar- gin for error. By paying attention to the process, the network manager can maintain the data in an organized fashion, track the validation of the model, and increase the confidence in the results.

There is no such thing as being too organized. The corollary is that if there is any confusion about the results, the procedure should be inter- rupted and investigated until there is no longer any confusion. Bad data and bugs are the worst enemies of accurate

simulations.

The process described in this chapter contains the following three basic phases

- Phase I Preparation
- Phase II Baseline
- Phase III Delta

The tasks for each phase are outlined briefly here, then discussed in detail

- Phase I tasks include
  - Goals. These must be stated in measurable and clear terms.
  - Data collection. The topology and traffic of the existing network is captured.
- Phase II tasks include
  - Capture. The collected data is captured in the model.
  - Validation. The captured model is validated.
- Phase III tasks include
  - Delta. Changes are applied to the baselined network.
  - Analysis. The results are analyzed and summarized.

PHASE I PREPARATION

This phase includes the definition of goals and the collection of topology and traffic data for the baseline network. The first task is identifying the goals.

Identifying Goals

A simulation should have clearly defined goals. For the hypothetical case discussed in this chapter, there are two principal goals The first is to develop a validated baseline model of the network in its current configura- tion; the second is to model the introduction of an asynchronous transfer mode (ATM) backbone.

A brief summary of the modeling strategy entails the following actions

1. Decide if modeling is appropriate.
2. Determine simulation goals.
3. Describe the network in one or two slides.
4. Combine each goal and its network description into a series of sce- narios, each with a simple, testable model description and a clearly defined goal.
5. For each model description, define the data to be collected, the results expected, and how the model will be validated.
6. Combine these individual documents into a simulation notebook.
7. After the individual models have been validated, repeat the process by combining the models into more complex models and validating each in a stepwise, iterative fashion.

Performance Metrics

In order to manage a network, its performance must be measurable and network goals must be specified in measurable metrics. Some of the more important metrics include

- queue buildup

- end-to-end delay
- throughput
- jitter
- goodput

Models must be instrumented according to the data collection require ments, by placing probes or turning on data collection routines in certain modules. The way that data collection is performed is unique to each simulation tool, but most tools will allow the collection of much of this infomation.

Data Collection

Once metrics are defined, the actual characteristics of the network as it exists currently must be collected. This is the first step in baselining the network.

First, a topology data collection sheet is created. A sample is shown in table Using the topology data collection spreadsheet, each hardware

Example of a topology data collection sheet.

Network Type	**Node Description**	**ID**		
Data	10Base-T hub	DH1		
Link Name	**Link Type**	**Speed**	**From**	**To**
WAN	Coaxial	DS-3	DH1	AT&T
Backbone	Fiber	SONET	DH1	DR1

device in the network and the links that connect them are documented. The number and rate of each link are identified.

This is also a good time to note all network costs. The cost of a link con- sists of the local component as well as the long-distance charges; some simulation tools are integrated with tariff data bases to some extent.

Most simulation tool manufacturers have or are planning interfaces to network management tools, which will simplify information collection on the data portions of the network. Tools have different levels of integration with network management. Some tools can collect topology information; others can also collect traffic information. There are three categories of traffic voice, video, and data.

Traffic collection is the more difficult activity, especially for data and voice. Once again, a very careful, systematic approach yields the best results.

Some useful numbers for calculating propagation delay times in various media are

Medium	**Propagation Delay**
Coaxial cable	4 μs/km
Fiber	5 μs/km

Voice Network Information. To collect the topology and traffic information for the voice portion of the network, here are some recommendations as to how the information might be represented in the model.

Topology. All segments of the voice network have to be described, including trunks, PBXs, and additional analog lines used by fax and data equipment for each location that is serviced by the network. Links to remote users also should be documented by listing the closest points of presence (POP) of any services that remote users will be calling. Quality of service and bandwidth required for each link should be noted.

Traffic. Billing information for any voice lines should be helpful in pro- viding usage patterns for voice links. Most PBXs collect this information

Any simple network management protocol (SNMP)-managed devices in the voice network may deliver usage information to the management tool.

This information may be expressed in several forms, but for the simula- tion a probability distribution function is needed to drive the traffic sources in the models. Because voice traffic exhibits a high degree of ran- domness, it is frequently viewed as a Poisson distribution. Depending on the modeling tool, there should be several distributions for voice traffic

- the distribution of addresses (i.e., who is calling whom)
- the distribution of the length of the message (i.e., call holding time)
- the distribution of the number of calls (i.e., call attempts)
- the desired quality of service

Video Network Information. To collect the topology and traffic informa- tion for the video portion of the network, here are some recommendations as to how it might be represented in the model.

Topology. All video segments should be described, in the same fashion as the voice network, noting areas of overlap.

Traffic. Again, billing information may be of great assistance when determining the usage pattern for video links. Many video codecs (e.g., PictureTel's) use two switched 56K-bps lines. It can be safely assumed that current usage for such a link is at least 2x56K-bps multiplied by the holding time of the call in this case. The manufacturer of the video codec equip- ment should be able to provide a more accurate idea of the traffic the device generates. The holding times of the video sessions may be derived from billing or from equipment checkout logs.

For each video source there should be the following information

- maximum bit rate generated by the codec
- holding time for each session
- number of sessions for the sample period
- address distribution of sessions (i.e., who is calling whom)

PHASE II BASELINE MODEL POPULATION AND VALIDATION

After capturing the data required to construct the baseline, model design can begin. The goal now is to transform the collected information into a valid baseline.

Simple small steps, gradually increasing the complexity as each step has been validated, should be used. This way, data collection efforts can be val- idated,

establishing confidence in the tool and modeling methodology.

Guidelines for Building Models

Creating Subnets. Preliminary steps should be modeling a portion of the network that the network manager understands very well — for instance, the simple case of determining the loading of the video portion of the Net. Although it may not prove to be a very interesting model, it will give confi- dence in the use of the tool. Later, smaller separate models for portions of each type of traffic in the network (e.g., voice, video, and data) can be built. By keeping the problems simple and only gradually adding complexity, the overall quality of the work is greatly improved.

The following paragraphs provide more detailed, step-by-step guide- lines. Vendors of network design tools provide considerable support for their products that can be of additional assistance.

Step 1 Tool Use and Data Collection Validation. During this step, the goal isto learn how to use the tool and to validate the data collection techniques.

It is likely that some data requirements of the problem or model will have been overlooked. A small representative of the network should be modeled, preferably using portions that are fairly well understood validating knowledge of the tool and the manager's ability to capture raw data in a model, much work will be saved. Putting off gaining this intu- itive understanding of how the modeling tool represents the network's components only postpones difficulties. Later, when the model has more data in it and more processes running, it will not be possible to see what happens with the very simple cases. The idea is to conduct simple experi- ments until the tool is completely understood.

Validating Subnets

Once a subnet is built, it must be validated. The process of validation requires running the model and comparing the results against data col- lected for the real subnet. Validation should be in three steps

- Topology. The topology of the subnetwork or network should be checked. For example, are all of the links connected to the correct devices? Are they of the correct type and bandwidth? Is each traffic source connected correctly?
- Routing. Are the router tables set up correctly? For the path of a packet traveling through the network, does it go where it should (and on the reverse path as well)?
- Load. Run the simulation with only one source turned on at a time. Is the correct amount of traffic sent to each destination? Is the delay close to what the network really experiences?

Step 2 Beginning to Validate the Data Network. At this point subnets (which later will be collected into a larger model) should be created. Start with subnets that can operate independently. Rather than spending time encoding large amounts of routing or other information for larger future net- works, it is better to model a small subnet correctly. In doing so, tools and models will be built up for future use in other areas.

Integrating and Validating Subnets. When the subnets have been vali- dated, they must be integrated and validated in a stepwise fashion. An example might be considering replacing a collapsed Ethernet backbone with a distributed ATM backbone that also carries video and contiguous portions of the voice network. First the data is integrated,

then the video, and finally the voice.

Once again, there are three tips to successfully building an accurate sim- ulation of a network

1. Never proceed if there are any doubts about a result. Stop immediately and investigate the problem until it is resolved. If the problem is not resolved, numerous bugs will creep in. This practice cannot be over-emphasized.
2. Be organized. There is a lot of data to compile. The spreadsheet can be the biggest asset.
3. Understand the network. To this end, a good network analyzer is invalu- able. If the expense of buying one cannot be justified, many leasing agencies and sometimes even the manufacturers rent them out.

PHASE III ALTERATION OF BASELINE TO ACQUIRE DATA

When a baseline with which to compare is completed and validated, alterations can be introduced. The alterations should be introduced with the same care that the baseline was constructed.

An example might be to compare the performance of both 100Base-T Ethernet and ATM backbone links. (Both should be able to handle the offered data load easily, but video traffic may be too much for the Ethernet backbone.) Loss and delay experienced by each traffic source should be measured, as well as the throughput on the backbone, to determine if these meet the required quality of service.

There are limits as to how much accuracy can be achieved by modeling. Efforts should focus on those areas that yield the highest return. Focusing on describing the offered load accurately will yield a more accurate simulation.

The goal is to have a model that yields valid information to guide decisions when doing experimentation on the network.

Determining Remote Bridge and Router Delays

THE USE OF QUEUING THEORY to determine the delays associated with remote bridges and routers. In addition, it investi- gates the effects of modifying the operating rate of the WAN links — in par-ticular, the effects of various communications circuit operating rates on equipment delays. There is a point beyond which increasing the operating rate of a communications circuit has an insignificant effect on equipment and network performance.

WAITING LINE ANALYSIS

Queuing theory, the formal term for waiting line analysis, can be traced to the work of A.K. Erlang, a Danish mathematician. His pioneering work spanned several areas of mathematics, including the dimensioning or siz- ing of trunk lines to accommodate long-distance calls between telephone company exchanges. This chapter bypasses Erlang's sizing work to con- centrate on the analysis of waiting lines.

BASIC COMPONENTS

Exhibit illustrates the basic components of a simple waiting line sys- tem. The input process can be considered the arrival of people, objects, or frames of data. The service facility performs some predefined operation on arrivals, such as collecting tolls from passengers in cars arriving at a toll booth or the conversion of a LAN data frame

into a synchronous data link connection (SDLC) frame by a bridge or router for transmission over a WAN transmission facility. If the arrival rate temporarily exceeds the service rate of the service facility, a waiting line known as a queue forms. If a waiting line never exists, the server is idle or there is too much service capacity.

Basic components of a simple waiting line system.

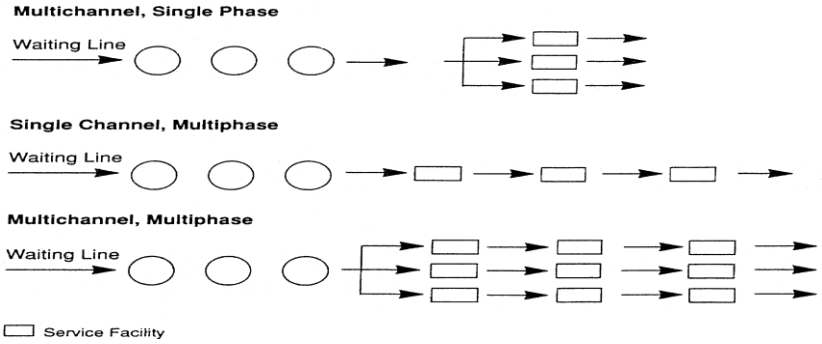

Figure **Other types of waiting line systems**

The waiting line system illustrated in Exhibit is more formally known as a single-channel, single-phase waiting line system — single chan-nel because there is one waiting line, and single phase because the process performed by the service facility occurs at one location. One toll booth on a highway or a single-port bridge connected to a LAN are two examples of single-channel, single-phase waiting line systems. Figurre illustrates three additional types of waiting line systems. On multichannel systems, arrivals are serviced by more than one service facil- ity, which results in multiple paths or channels to those service facilities. On multiphase systems, arriving entities are processed by multiple service facilities.

example of a multiphase service facility is a toll road in which driv- ers of automobiles are serviced by several series of toll booths; for exam- ple, a turnpike that has toll plazas every few miles. Another example of a multiphase system is the routing of data through a series of bridges and routers. The computations associated with multiphase systems can become quite complex, and because most networks can be analyzed on a point-to-point basis as a single-phase system, this chapter restricts its examination of queuing models to single-phase systems.

Network Baselining as a Planning Tool

BASELINING PROVIDES A MECHANISM FOR DETERMINING THE LEVEL OF UTILIZATION OF A NETWORK, including its computational and transmission facili- ties. As such, it plays a central role in a network manager's capacity plan- ning effort

because the baseline shows whether or not there is sufficient capacity available, as well as providing a foundation for future network measurements that can be compared to the baseline to indicate the direc- tion of network utilization. Thus, the network baselining effort represents the first major step in the capacity planning effort.

In addition, baselining enables network managers and administrators to identify and respond to network capacity requirements before they become an issue, in effect providing a mechanism to head off network- related problems.

BASELINING TOOLS AND TECHNIQUES

There are a variety of network baseline tools and techniques that can be used to facilitte an organization's capacity planning effort. The actual tech- niques employed are commonly based on the type of tool used. This chap- ter focuses on a number of commercially available network baselining tools, and discusses appropriate techniques concerning their use.

SimpleView

SimpleView is an easy to use and relatively inexpensive Simple Network Management Protocol (SNMP) management platform from Triticom, Inc., of Eden Prairie, Minnesota. Through the use of SimpleView, users can retrieve statistical information maintained by Remote Monitoring (RMON) network probes. SimpleView supports a Management Information Base (MIB) walk capability, shown in the MIB Walk window, that lets a user click on an MIB group to select the group starting point, or double-click on the group to explode its elements, enabling a specific element from the group to be selected for retrieval

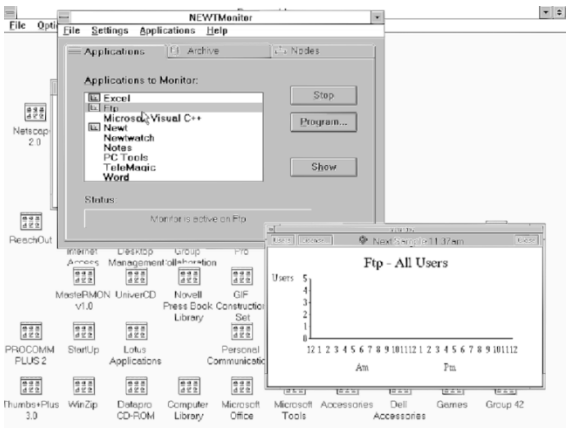

The NetManage NEWTMonitor program provides the capabil- ity to examine the activity on a network based upon certain types of predifined applications.

NEWT

NetManage of Cupertino, California, well known for its Chameleon suite of Internet applications, also markets a program called NEWT that can be used to monitor the use of desktop applications as well as to provide sta- tistics on network activity associated with individual users. Exhibit 45-1 illustrates the use of NEWTMonitor on the author's computer to monitor the number of simultaneous FTP sessions occurring over a period of time. Doing so can be extremely important, especially when used in conjunction with normal RMON traffic statistics that do not look beyond the data link layer. NEWTMonitor enables the use of specific types of TCP/IP applica- tions. In comparison, if the network probes and network management sys- tem support RMONv2,

or can be upgraded to this new version of RMON, it can be used to obtain a distribution of traffic through the application layer.

Exhibit 45-2 illustrates the use of NEWTGraph to display different TCP/IP statistics by node. In the example shown in Exhibit 45-2, the author dis- played Interface Errors for his node.

EtherVision

When checking the activity associated with an individual network, users can choose from a variety of network monitoring programs. One such pro- gram is EtherVision, also from Triticom, Inc., of Eden Prairie, Minnesota

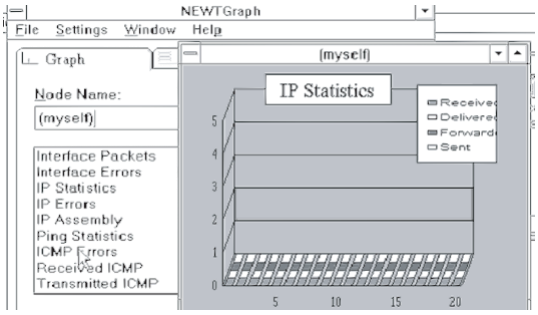

Through the use of the NetManage NEWTGraph program, graphs of different types of TCP/IP statistical information can be displayed.

Figure illustrates the statistics summary display based on the mon- itoring of frames using their source address for constructing a statistical baseline. EtherVision supports monitoring by either Source or Destination address, enabling users to build two baselines. In examining Exhibit 45-3, note that the statistics summary presented indicates the frame count over the monitored period of time, current network utilization in the form of a horizontal bar graph, and a summary of “average,” “now” or current, and “peak” utilization displayed as a percentage, as well as the time peak utili- zation occurred. The latter can be extremely handy as it allows a user to run the program on a workstation connected to an Ethernet LAN and return at the end of the day to determine the peak percentage of network use as well as when the peak occurred.

Although not shown in Exhibit 45-3, an EtherVision user can also set the program to generate a report that will log each period of activity over a cer- tain percentage of network activity. Then, using the logged report, a net- work manager or LAN administrator can easily determine the distribution of network utilization throughout the monitoring period.

In the upper right corner of Exhibit 45-3, note that EtherVision maintains a distribution of frames transmitted on the network based on their size or length, falling into five predefined intervals. By examining the distribution of frames based on their length, users can determine the general type of traffic flowing on a network. This is possible because interactive query- response applications are generally transported in relatively short frames. In comparison, file transfers, such as Web browser pages containing one or more images, commonly fill frames to their full length. In examining the dis- tribution of frame sizes shown in Exhibit 45-3, note that there are relatively few full-sized Ethernet frames in comparison to the total number of frames

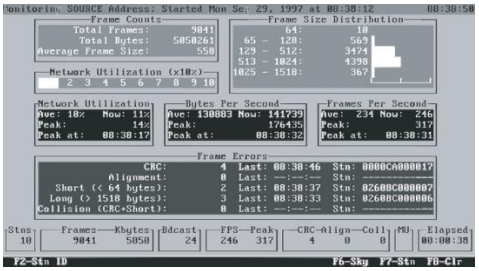

The Triticom EtherVision statistics summary display can be used to obtain information about network utilization and frame distribution.

encountered during the period of monitoring. This indicates a low level of file transfer and Web browser activity occurring on the monitored network.

Although EtherVision provides numeric information concerning net- work utilization, many users prefer to work with charts that note trends at a glance. To accommodate such users, EtherVision includes a number of built-in displays such as the one shown in Exhibit 45-4, which plots net- work utilization over a period of time. By examining a visual display, users can immediately note any potential capacity-related problems. In the example shown in Exhibit 45-4, the maximum level of network utilization is slightly above 46 percent. However, based on the monitored period, net- work traffic rose from 22 to 46 percent numerous times during the monitor- ing period. Because an Ethernet LAN gets congested at utilization levels above 50 percent due to its CSMA/CD access protocol, and the effect of the delay associated with the use of a random exponential back-off algorithm after a collision occurs, Exhibit 45-4 indicates a baseline of network utiliza- tion that justifies careful attention and a scheduled remonitoring effort to ensure traffic on the network does not turn into a bottleneck.

Foundation Manager

Foundation Manager, a product of Network General Corporation, is a sophisticated SNMP Network Management System (NMS) platform that

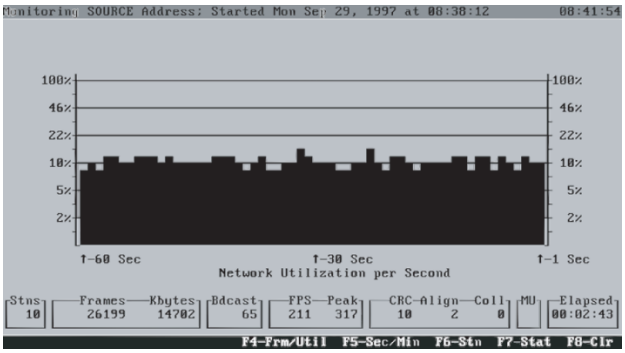

EtherVision supports the display of network utilization over a period of time, which

facilitates observing the changing state of this important baseline metric.

operates on Intel-based computers using different versions of Microsoft's Windows operating system.

Foundation Manager was upgraded to support the emerging RMONv2 standard. When used to gather statistics from an RMON v2-compatible probe, it can provide a summary of statistics through the application layer, allowing it to replace the use of multiple products to obtain equivalent information.

Exhibit 45-5 illustrates the use of Foundation Manager to monitor a local Token Ring network. In the example, two buttons under the Local Token Ring Monitoring bar were pressed to initiate two displays of information from the Token Ring Statistics Group that an RMON probe on the locat net- work accumulates. The first button clicked on is the bar chart icon to the right of the icon with the upraised hand in the form of a stop sign. Clicking on the bar chart icon results in the display of the top row of eight bar charts that indicate the total number of different types of frames and level of network utilization.

For example, the second bar chart located on the left side of the top dis- play indicates that network utilization is at three percent on a 100 percent basis. Other bar charts on the top row indicate the current number of log- ical link control (LLC) bytes and frames, multicast frames, broadcast frames, beaconing frames, purge events, and claim events. The second row

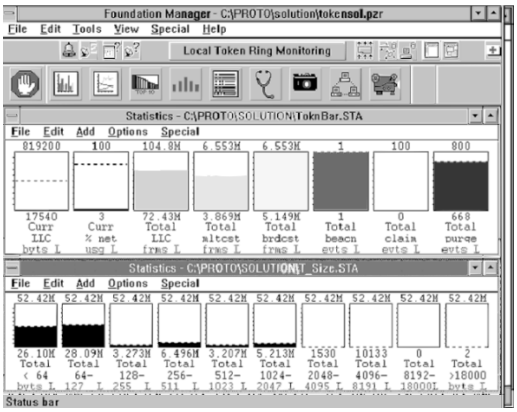

Using Network General's Foundation Manager to monitor the distribution frame by bar charts resulted from clicking on the third icon to the right of the raised hand icon. This sequence of ten bar charts indicates the distribu- tion of Token Ring frames in a manner similar to the method that EtherVi- sion used to summarize Ethernet frame sizes. Foundation Manager follows the RMON standard and provides a more detailed breakdown of the distri- bution of Token Ring frames by their length.

Similarly, when using Foundation Manager to monitor Ethernet net- works, the program retrieves RMON probe-kept frame distribution infor- mation that is more detailed than that kept by EtherVision. However, it is important to note that the retail price of EtherVision is under $500 and it can operate by itself. In comparison, the retail price of Foundation Manager is approximately $5000 and a single probe can cost approximately $1000, requiring an investment of an additional $5500 to obtain an

enhanced level of frame size distributions as well as some additional features.

Two of the more interesting features of Foundation Manager are its QuickStats and discovery, and baselining capabilities. Exhibit 45-6 illus- trates the use of the Foundation Manager Quick Stats feature to display a quick set of statistics for a remotely monitored network. In the example, statistics for an RMON probe connected to a network located in San Diego are displayed.

Foundation Manager is capable of displaying up to eight Quick Stats graphical reports at one time, with each report generated by clicking on an

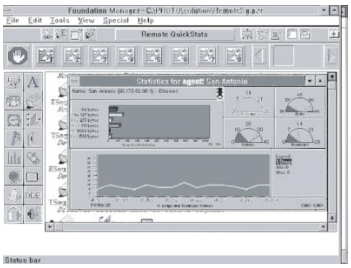

The Foundation Manager QuickStats display provides users with the ability to visually note important network baseline parameters both in real-time and over a period of time.

appropriate icon to the right of the icon with the raised hand in the form of a stop sign.

Each Quick Stats display presents summary information about a moni- tored network in a similar manner. In examining the statistics display for the network located in San Diego, the upper left display presents a distri- bution of frame length for the monitored LAN as a horizontal bar chart. The upper right portion of the display contains four gauges that provide a real- time view of network utilization, bytes transmitted, broadcast traffic, and frame rate. The lower half of the display shows a real-time plot over a period of predefined length for any two of the gauge values. Thus, the use of the Foundation Monitor Quick Stats display provides users with the abil- ity to visually note important network baseline parameters both in real- time and over a period of time.

A second interesting feature built into Foundation Manager is its discov- ery and baselining capability. This capability is available for both local and remotely located networks being monitored, and provides the ability to gather pattern flow information that can be extremely valuable when attempting to determine if the cause of a high level of network utilization results from the activity of one or a few stations on the network.

Figure illustrates the Foundation Manager local discovery and baselining display as a matrix map of network activity. The first portion of the title of the display, “Discovery,” results from the fact that the probe

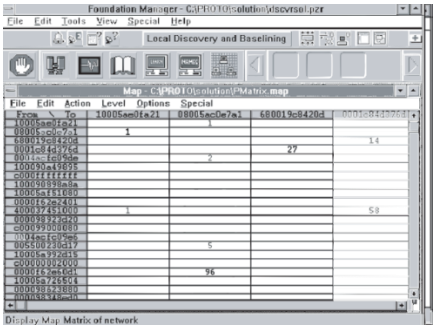

The local discovery and baselining capability of Foundation Manager enables the flow of data between stations to be identified.

examines each frame flowing on the monitored network and discovers its source and destination by examining the source and destination addresses contained in the frame. The second portion of the title of the display, "Base- lining," results from the fact that Foundation Manager extracts information from a matrix table maintained by the probe that denotes the number of frames transmitted from one address to another. Thus, in examining Exhibit 45-7, such numerics as 1, 2, 27, 14, 58, 5, and 96 represent the number of frames transmitted from the address located in the row in the table to the address in the column portion of the table.

When baselining a network, matrix information should be considered as a mechanism to identify the cause of high network utilization. If Quick Stats or a similar display denotes a low level of network utilization, there is no need to use the matrix capability of Foundation Manager or a similar prod- uct to identify the actual flow of data between network stations. This is because even if the user can locate a station using too much bandwidth, a modification of the operation of the station will, at best, have a negligible effect upon improving network performance if the network already has a low level of utilization.

TEXT/ REFERENCE BOOKS

1. Gil Held, "Network Design: Principles and Applications (Best Practices)", Auerbach Publications, 1st edition, 2000.
2. Diane Tiare and Catherine Paquet, "Campus Network Design Fundamentals", Pearson Education, 1st edition, 2006.
3. Larry L. Peterson, Bruce S. Davie, "Computer Networks: A Systems Approach", Morgan Kaufmann Publishers Inc., 5th edition, 2012.
4. William Stallings, "Data and Computer Communications", Pearson Education, 8th edition, 2016.
5. James F. Kurose, Keith W. Ross, "Computer Networking - A Top-Down Approach Featuring the Internet", Pearson Education, 6th edition, 2012.

PART-A - 2 Mark Questions

S. No	Questions (2marks)	CO	Level
1	Illustrate how the multimedia networks are designed?	CO4	2
2	Summarize the concept of simulation used in designing the network.	CO4	2
3	Interpret the concept of bridge	CO4	5
4	Justify the need for router in designing a network	CO4	5
5	Categorize the different tools used in designing the network.	CO4	4
6	Illustrate the need for planning tool used in designing network.	CO4	2
7	Interpret how the various technologies are helpful in designing a good network?	CO4	5
8	In today's scenario we need Intelligent networks. In this regard Appraise the concept of network baseline.	CO4	5
9	Comment on the factors where delays are introduced in the router	CO4	4
10	Justify how the multimedia networks are simulated efficiently with steps.	CO4	2

PART-B - 10 Mark Questions

S. No	Questions (2marks)	CO	Level
1	For establishing a good communication between multimedia networks it is very important to have an effective network. Illustrate the various simulation methods for designing the multimedia network.	CO4	5
2	Appraise on how the remote bridge is used for exchanging the information from source to destination with necessary examples and diagrams	CO4	2
3	As a data analyst you were assigned to develop a software package for detection of router delays. Suggest an appropriate technique which is used to pre-process data so that the effect of delay is removed in detail	CO4	5
4	Without training, the computer Networks cannot converge. As a network engineer expert explain the different procedures used for designing a network.	CO14	5
5	As a design engineer expert you are tasked with the development of a good network. Suggest the suitable the suitable networking tools used for designing the network and also specify how it guarantees a fast convergence?	CO4	5

SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF COMPUTER SCIENCE AND ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE ENGINEERING

Computer Networks and Design – SCSA1502

UNIT 5 -EVOLVING TECHNOLOGIES-SCSA1502

UNIT 5
EVOLVING TECHNOLOGIES

Trends in data communications – Merits of xDSL technology – Preparing for cable modems - Voice and video on the LAN –Internet voice applications – Building IP PBX telephony network – Fax over IP – Videoconferencing over IP networks.

The past few years have been marked by unprecedented technological Innovation in the communications field. Not only have new offerings such as frame relay and SMDS demonstrated their value, but older services such as X.25, T1/T3, and Integrated Services Digital Network are being revital- ized. Making their debut are several new technologies that leverage the car- rier's existing local loop infrastructure in support of high-speed data, including asymmetric digital subscriber line (ADSL) and high-bit-rate digi- tal subscriber line (HDSL). Also being deployed are new wireless technolo- gies, including Cellular Digital Packet Data, that connect mobile computer users to each other as well as to the data center. In addition, the country's entire communications infrastructure is poised for advancement with Synchronous Optical NETwork and Asynchronous Transfer Mode technologies for transmission and switching at multigigabit speeds. Long term, the com- bination of SONET and ATM paves the way for broadband ISDN (BISDN).

Communications at the customer premises is also undergoing rapid advancement. Fiber to the desktop and between LAN hubs is now a reality with the Fiber Distributed Data Interface, which offers transmission speeds of up to 100M bps. Several economical, albeit limited-distance, alternatives are offered over ubiquitous centerpair wiring, including 100Base-T and 100VG-AnyLAN. LAN hubs are also available that offer high-speed ATM switching enterprisewide. Another type of device, the bandwidth control-ler, is available for assembling low-speed 56K or 64K bps channels into eco- nomical high-speed pipes to support high-bandwidth applications over the public network.

Communications already plays an important role in the data center, and recent developments in this area will strengthen its role there. To aid the data center operations manager in keeping up with rapidly changing technologies, this chapter examines the key advances in communications in the context of corporate networking and applications requirements. In particular, this chapter examines developments in

- optical fiber (e.g., line technologies, such as asymmetric and high-bit-rate digital subscriber line)
- cellular digital packet data
- T1 and T3
- X.25 frame relay
- integrated services digital network
- switched multimegabit data services
- asynchronous transfer mode
- synchronous optical network
- fiber-distributed data interface

OPTICAL FIBER

More computer and telephony applications are being developed that require the bandwidth that can only be provided by optical fiber. Carriers have been installing fiber-optic cable when it can be provided at a cost that is comparable to other transport modes, such as copper, or when compe- tition from alternative access providers forces them to

take counter mea- sures. Thus, the carriers are committed to providing a fiber backbone as close to the point of service as is cost effective.

However, if the telephone companies are to offer advanced services in support of such bandwidth-intensive applications as multimedia, docu- ment imaging, and videoconferencing, they must increase as much as pos- sible the transmission capacity of existing twisted-pair wiring in the local loop — and do so without further delay. Fiber to the customer premises would be best, but replacing existing arrangements could cost billions of dollars and will likely take 25 years to complete.

Asymmetric Digital Subscriber Line. One of the most promising local-loop upgrade technologies is asymmetric digital subscriber line (ADSL). This technology allows for the transmission of more than 6M bps over existing twisted-pair copper wiring. ADSL carves up the local loop bandwidth into several independent channels suitable for any combination of services, including LAN to LAN data transfers, ISDN, and plain old telephone service (POTS). The electronics at both ends of the ADSL compensate for line impairments, increasing the reliability of high-speed transmissions.

High-Bit-Rate Digital Subscriber Line. An alternative to ADSL is high bit- rate digital subscriber line technology (HDSL), which uses two full duplex pairs, each operating at 784K bps. This technology is an electronic technol- ogy for conditioning lines for heavy data usage without the use of repeaters

Over distances of up to 12,000 feet from the central office. Because approx- imately 85 percent of the local loops nationwide are within this distance, HDSL promises to have a significant impact on the embedded copper net- work. With the addition of a "doubler" technology, transmission distances can be increased even more, enabling HDSL to be used in all local loops andfurther easing the bottlenecks from fiber facilities.

Although HDSL was invented to solve telephone companies' T1 provi- sioning problems in the local loop, it is being extensively used in private networks as well. Universities, military bases, hospitals, corporate com- plexes, local governments, and other campus environments where multi- ple buildings require connections over relatively short distances, are ideal for HDSL solutions. Previously, campus applications, such as LAN-to-LAN, videoconferencing, CAD/CAM, and PBX networks were often restricted to 56K bps, due to cost and technical limitations.

CELLULAR DIGITAL PACKET DATA

Cellular digital packet data (CDPD) is a data-over-cellular standard for providing a LAN-like service over today's cellular voice networks. The CDPD infrastructure uses existing cellular systems to access a backbone router network, which uses the internet protocol (IP) to transport user data. PDAs, palmtops, and laptops running applications that use IP can con-nect to the CDPD service and gain access to other mobile computer users or to corporate computing resources that rely on wireline connections.

Because CDPD leverages the existing $20 billion investment in the cellu- lar infrastructure, carriers can economically support data applications and avoid the cost of implementing a completely new network, as most com- peting technologies would require. CDPD also offers a transmission rate that is four times faster than most competing wide-area wireless services, many of which are limited to 4.8K bps or less.

In addition to supporting enhanced messaging services, including mul- ticast,

cellular paging, and national short-text messaging, CDPD extends client/server-based applications from the LAN environment into the wire- less environment in support of mobile computer users. This extension can be used for such applications as database updates, schedule management, and field service support. However, the ultimate success of CDPD is closely tied to industry efforts to standardize its implementation. Using a universal standard for cellular packet data would allow users to roam and would sim- plify the introduction of wireless data services.

T1 AND T3

Traditional T1, which operates at 1.544M bps, is also undergoing innova- tion. For example, local exchanges are now being equipped to support switched T1 for on-demand (e.g., dial-up) service. Setup takes less than two seconds, providing a fast, efficient, and economical alternative to ded- icated lines, which entail fixed monthly charges regardless of how little they are used. Some local-exchange carriers offer switched fractional T1, operating at 384K bps.

Some interexchange carriers offer fractional T3, in which a number of T1-equivalent pipes can be selected by customers to meet the bandwidth requirements of bandwidth-intensive applications without having to over- subscribe.

Nx64. A related innovation is Nx64 service, available from all major car- riers. This allows users to build switched data pipes in bandwidth incre- ments of 64K bps to support such applications as videoconferencing, CAD/CAM, or LAN interconnection. Channels of 64K bps can be added or dropped as necessary to support an application.

Inverse Multiplexer. A relatively new type of T1 device is the inverse mul- tiplexer or bandwidth controller. Inverse multiplexing is an economical way to access the switched digital services of interexchange carriers because it provides bandwidth on demand, without having to subscribe to ISDN, which by comparison is more expensive and more complicated to configure.

Users dial up the appropriate increment of bandwidth needed to sup- port a given application and pay for the number of 56K bps local-access channels needed. Channels can even be added or dropped during the transmission if bandwidth requirements change. Once transmission is completed, the channels are taken down. This eliminates the need for pri- vate leased lines to support temporary applications.

Fractional T3

T3 represents the equivalent of 28 T1 lines operating at the DS3 rate of 44.736M bps. With T3, users gain the additional bandwidth needed for a new generation of bandwidth-intensive applications, as well as traditional LAN-to-LAN and host-to-host interconnection. The absence of an optical standard for DS3 restricts the user's ability to mix and match equipment from different manufacturers. Customers must negotiate the types of optical interfaces to be placed in the interexchange carrier's serving office. In con- trast to the more widely available T1 lines, T3 requires special construction for the local channel, from CPE to POP. Thus, T3 local-access lines are pro- vided on an individual basis and usually entail high installation costs.

To broaden the appeal of T3, some carriers are offering fractional T3. With fractional T3, users can order bandwidth in T1 increments up to the full T3 rate of 44.736M bps. This service is designed for users who need more than the 1.544M-bps rate offered

by T1, but less than the full bandwidth offered by T3 to support the interconnection of Token-Ring or Ethernet LANs. This enables corporate locations to share such high-bandwidth appli- cations as document imaging, CAD/CAM, and bulk file transfers between LANs or hosts. In the case of LANs, a bridge or router is used for interconnection through the public network. The public network appears as an extension of the LAN.

X.25

X.25 was developed before digital switching and transmission technology became available. Because networks suffered at that time from interference from noise, X.25 relied on a store-and-forward method of data communica- tion to ensure error-free transmission. When errors arrive at a network node, a request for retransmission is sent to the originating node, which retains a copy of the packets until they are acknowledged. This process is repeated at each network node until the data is delivered to its destination.

Although X.25 is highly reliable, today's digital networks have rendered unnecessary its stringent error-correction and other overhead functions. Network throughput can be greatly increased by leaving out these process- ing-intensive functions and relegating them to the customer premises equipment, as is done with frame relay. Despite the advances made with frame relay technology, the popularity of X.25 remains high.

Because frame relay is often compared to X.25, managers have become more informed about X.25 and its role in supporting value-added and dial- up applications, data entry, short file transfers, and financial and point-of- sale transactions. X.25's error correction is particularly useful for data communications to international locations where high-quality digital facil- ities still are largely unavailable. Although X.25 has been eclipsed by frame relay and switched multimegabit data services (SMDS), there have been advances in X.25.

Most X.25 networks operate at 56K bps, but some vendors have increased the transmission speed of their X.25 offerings to T1 rates and faster, making them more effective for LAN interconnection. Speeds as high as 6M bps performance are supported by some packet switches. Currently available are integrated switches that support both circuit and packet switching — the most appropriate switching method is selected in real time according to applications requirements. Also available are X.25 gate- ways to frame relay, SMDS, and other data services.

FRAME RELAY

Frame relay is a packet technology that offers performance advantages over X.25 while allowing users to more easily interconnect high-speed LANs over the WAN.

The concept behind frame relay is simple protocol sensitivity, unneces- sary overhead functions, and associated processing at each network node — all characteristic of X.25 — are eliminated to obtain higher trans- mission rates. The reliability of digital links enables frame relay service, because error correction and flow control already exist at the network and transport layers of most computer communication protocol stacks. Because these functions have been relegated to the edges of the network rather than placed at every node along a path, as in X.25, bad frames are simply discarded. Upon error detection, customer premises equipment at each

end of the path requests and implements retransmissions.

Frame relay is optimized to transmit traffic in bursts, which is the way applications traverse the LAN. Therefore, when interconnecting geograph- ically separate LANs, organizations should consider frame relay service. It is often cost-justified for sites only 750 miles apart, roughly the distance between New York City and Chicago. Frame relay also allows a variable frame size to make the most efficient use of available bandwidth. Frame relay's variable-length frames also mesh well with the variable-length pack-ets used in TCP/IP, OSI, and DECnet.

Today's frame relay services are based on permanent virtual connec- tions (PVCs) that correspond to the organization's network nodes. Node addresses are stored in each switching point on the network so that frames can be routed accordingly. For each PVC, the customer chooses a commit- ted information rate (CIR) that supports the application, and the carrier bills for it accordingly. In frame relay, a 256K bps virtual connection can handle bursts of up to 1M bps. However, too many users exceeding their CIRs at the same time creates the possibility of the network becoming con- gested and of frames being discarded. Fortunately, carriers overprovision their frame relay networks to guard against this situation.

As companies continue to move from the host-centric data center to the distributed computing environment, the major carriers are starting to offer managed services that address the specific needs of System Network Architecture users. The advantages of frame relay over leased lines are clear-cut, especially for SNA users with many remote locations that must be tied into one or more hosts. Among them are the following

- Frame relay PVCs replace expensive SDLC and BSC multidrop net- works between the host and branch offices.
- Consolidating connections through frame relay eliminates costly serial line interface coupler (SLIC) ports on FEPs while increasing per-formance.
- WAN access extends the useful lives of SDLC/BSC controllers and 3270 terminals
    - The availability of systems network architecture (SNA) connections is increased by allowing controllers to take advantage of WAN connec- tions with multiple host paths.

A managed frame relay service includes the frame relay access devices (FRADs) — leased or purchased — that transport SNA traffic over the PVCs. FRADs are more adept than routers at congestion control and rank- ing traffic according to priority. Some FRADs multiplex multiple SNA/SDLC devices onto a single PVC, instead of requiring a separate PVC for each attached device, resulting in even greater cost savings.

For a legacy SNA shop that does not have the expertise or resources, a carrier-managed frame relay service is a viable option, especially because frame relay networks are much more difficult to configure, administer, and troubleshoot than private lines. All of this activity can be outsourced to the carrier. The frame relay services themselves are priced attractively. On aver-age, the frame relay service costs about 25 percent less than the equivalent private network. In some cases, discounts of up to 40 percent are possible.

INTEGRATED SERVICES DIGITAL NETWORK

Services built around the Integrated Services Digital Network (ISDN) pri- mary rate interface (PRI) rely on switched T1 facilities. Of the 24 64K bps channels, 23 are bearer channels used for voice or data applications, and the twenty-fourth — the D channel — supports call-management functions.

ISDN has enabled early users to eliminate modem and infrequently used dedicated lines and to economically back up dedicated lines. The ISDN automatic number identification feature has also enabled some users to build new applications that integrate the traditionally separate domains of computer databases and voice communications. In this type of applica- tion, customer data is retrieved from a database and displayed at a termi- nal as the call is routed to a customer service representative. This arrange- ment improves customer response and employee productivity.

Another innovative use of ISDN comes in the form of improved call rout- ing. Without actually connecting the call, ISDN's signaling channel first determines whether a PBX or automated call distributor (ACD) can handle it. If not, the call is forwarded to a PBX or ACD at another location that can take the call. This arrangement is useful for businesses spread across dif- ferent time zones in that they can extend normal business hours. It also provides failure protection, so that if one location experiences an outage, another location can take the calls.

After a slow start, ISDN PRI is finally gaining user acceptance for such practical applications as network restoral, performance control, and peak traffic handling. The ISDN basic rate interface (BRI), which offers two bearer channels of 64K bps each and one 16K bps signaling channel, also is undergoing a resurgence, supporting such high-demand applications as computer-telephony integration (CTI), telecommuting, and Internet access. Many communications servers on the LAN now support ISDN, as do communications controllers in the host environment. Because ISDN ser- vice uses digital lines, users benefit from improved reliability, as well as from faster call setup time.

SWITCHED MULTIMEGABIT DATA SERVICES

Switched Multimegabit Data Services (SMDS) is a high-speed, connec- tionless, cell-based service offered by the regional telephone companies. It is used primarily for linking LANs within a metropolitan area. It offers cus- tomers the economic benefits of shared transmission facilities, combined with the equivalent privacy and control of dedicated networks. SMDS is much easier to provision and manage than frame relay and, over short dis- tances, SMDS can be more economical than frame relay.

Despite its advantages, however, SMDS has not fared well against frame relay. One reason for frame relay's popularity is that it became available nationwide at a very early stage, whereas SMDS was promoted as a regional service. Only recently has SMDS become available from long-dis- tance providers. This may alleviate user concerns about being able to link far-flung corporate sites using SMDS. Another sticking point for SMDS has been that access was not available at speeds lower than T1. Now SMDS is routinely offered by the regional telephone companies at access speeds between 56K bps and 34M bps. With these improvements, the demand for SMDS should grow at a much faster clip in the foreseeable future.

ASYNCHRONOUS TRANSFER MODE

Asynchronous Transfer Mode (ATM), also known as cell relay, is a gen- eral-purpose switching method for multimedia (e.g., voice, data, image, and video). Whereas frame relay and SMDS use variable-length frames, the cell size used by ATM is fixed at 53 bytes. This fixed size facilitates the switching of cells by hardware-based routing mechanisms, enabling oper- ation at extremely high speeds. ATM speeds are scalable and can exceed 2.5G bps over optical fiber.

Despite the need to break larger variable-rate frames into fixed-size cells, the latency of ATM is orders-of-magnitude less than frame relay alone. For example, on a five-node network spanning 700 miles, ATM exhib- its 0.3m-second latency vs. 60m-second latency for frame relay at T1 speeds. (At T3, the latency of ATM is only 0.15m seconds.) Thus, ATM makes for fast, reliable switching and eliminates the potential congestion problems of frame relay networks. A nonblocking switching method, ATM virtually eliminates the buildup of congestion that can hamper the performance of campus LANs and inter- campus backbones. ATM hubs also allow networks to grow smoothly. Only switching capacity needs be added to handle increases in traffic; the user interfaces are not changed. ATM hubs are star-wired with direct links to every attached device. This configuration not only minimizes network management overhead but facilitates the collection of statistics for fault isolation, accounting, administration, and network planning.

ATM provides the features necessary for successful multimedia applica- tions. Specifically, it has the ability to define different traffic types, with each traffic type delivering a different quality of service based on the unique properties associated with it. The traffic type that supports multi- media applications is called constant bit rate (CBR) service. CBR supplies a fixed-bandwidth virtual circuit, which addresses the special handling needs of delay-sensitive multimedia applications — those that contain real-time video and voice, for example. The quality of service is negotiated with the network. The applications themselves can do the negotiation through native ATM interface, or such interfaces as LAN emulation, an ATM Forum standard, and classic IP can perform the negotiation for the applica-tions over ATM.

When the quality of service is negotiated with the network, there are performance guarantees that go along with it maximum cell rate, available cell rate, cell transfer delay, and cell loss ratio. The network reserves the full bandwidth requested by the connection. There is no data rate limit for CBR connections, nor is there a limit on how long a connection can trans- mit at the maximum cell rate, otherwise known as the peak cell rate (PCR). The PCR is the maximum data rate that the connection can support with- out risking data loss. Any traffic above the specified rate risks being dropped by the network, whereas traffic below the specified rate will fail to satisfy the needs of the application.

These and the other advantages of ATM — including low latency, high throughput, and scalability — will one day make it the network of choice for supporting new, high-bandwidth multimedia applications, as well as legacy LAN and TCP/IP traffic. Meanwhile, ATM has been slow to take off because of the start-up costs of implementation. On the WAN side, carriers must invest in a new overlay infrastructure — something they have been slow in doing. Likewise, on the LAN side, companies must invest in new hubs, server interfaces, and workstation adapters. In some cases, new cabling may also be required.

SYNCHRONOUS OPTICAL NETWORK

Synchronous Optical Network (SONET) technology allows the full potential of the fiber-optic transmission medium to be realized. SONET standards specify transmission rates that start at 51.84M bps and reach to 2.488G bps and make provisions for transmission rates of 13G bps. Throughout this decade and beyond, SONET will gradually replace the pro-prietary T3 asynchronous networks of today.

With the same fiber cable that supports asynchronous networks, trans- mission capacity can be increased one-thousandfold by using end-to-end SONET equipment. SONET also supports a variety of current and emerging carrier services, including ATM, SMDS, and BISDN, increasing their reliabil- ity through embedded management functions.

Most of the activity in SONET deployment has been in building fiber rings that offer customers fail-safe data communications in major metro- politan areas. SONET equipment can reroute traffic instantly if a section of the network fails or becomes disabled. This level of reliability is increas- ingly becoming a critical requirement for businesses whose networks are becoming more data intensive with each passing year.

Today's self-healing SONET ring networks are capable of operating at 622M bps. Within 50m-seconds of a cable cut, customer traffic running over the primary path is automatically routed along the backup path, with no loss of data. This recovery process is implemented by the SONET Add- Drop Multiplexer (ADM), which duplicates and sends the data in opposite directions over the dual paths. When the signals reach a common point on the network, they are compared, then one set is discarded, while the other is delivered to the destination. If the primary path is disrupted, the data on the backup path is passed on to the destination.

Although the carriers have been installing SONET rings in major metro- politan areas at a pace that has been keeping up with customer demand, they are expected to step up their deployments of SONET. For many telcos, the cost-benefit threshold has been crossed; that is, the financial incen- tives of SONET deployment now exceed the costs of new installations.

FIBER DISTRIBUTED DATA INTERFACE

There is no question that traditional Ethernet and Token Ring LANs are beginning to get bogged down by the added bandwidth requirements of CAD/CAM, document imaging, collaborative computing, and multimedia applications. Fiber Distributed Data Interface (FDDI) backbones offer 100M bps of bandwidth that can alleviate potential bottlenecks. FDDI uses a token-passing scheme similar to Token Ring and a dual-ring fault protec- tion scheme similar to SONET.

Normally considered a private networking solution for office environ- ments, FDDI is now being offered as a wide-area network service by some carriers. Bell Atlantic, for example, provides a service to extend 100M-bps FDDI LANs across WANs. The service, called Full FDDI, is a response to requests from financial firms that wanted to use this kind of service together with a similar one offered by New York Telephone. The financial companies wanted to interconnect offices in New York City with data cen- ters in New Jersey. The only difference between the two services is that Bell Atlantic places FDDI concentrators on customers' premises, while under New York Telephone's service, called Enterprise Service FDDI, con- centrators are located at

the carrier's central office.

Optical fiber equipment and adapters are still too expensive for FDDI deployment throughout the enterprise. An economical alternative to FDDI in the office environment is the twisted-pair distributed data interface (TPDDI). TPDDI offers the same speed as FDDI over ordinary twisted-pair wiring at distances of up to 100 meters (328 feet) from station to hub, which is enough to accommodate the wiring schemes of most office environ- ments. TPDDI is designed to help users make an easy transition to 100M- bps transmission at the workstation level, in support of bandwidth-inten- sive data applications.

For companies with large investments in Ethernet, other 100M-bps tech- nologies worth considering are 100Base-T, also known as Fast Ethernet, and 100VG-AnyLAN, a non-standard Ethernet extension technology. The major difference between the two is that 100Base-T preserves the contention access scheme of pure Ethernet, while 100VG-AnyLAN dispenses with it.

The Emerging Advantage of xDSL Technology

SOME OF THE FLAVORS OF XDSL TECHNOLOGY have been around for awhile but sold under a different label. However, for the most part, end users just do not care what it is called, so long as it provides more bandwidth at less cost than before. Currently, Digital Subscriber Line (DSL) technology comes in several popular flavors asymmetrical DSL (ADSL), high bit-rate DSL (HDSL), symmetric DSL (SDSL), or very high speed DSL (VDSL). There is also rate adaptive ADSL (RADSL). All of these are collectively referred to as xDSL, where x is the designator for the service.

WHAT HAS CHANGED

The original Digital Subscriber Line (DSL) service was introduced as ISDN in the 1980s. This technology compressed 160Kbps into an 80KHz bandwidth of the local loop. ISDN utilized a four-level PAM modulation (2 Binary, 1 Quaternary) "2b1Q" to reach the range of 18,000 feet.

High Bit-Rate Digital Subscriber Line (HDSL) came along in the early 1990s and used the same 2 Binary, 1 Quaternary (2b1Q) line coding to sup- port T1 services. HDSL made it possible to provision loops of up to 12,000 feet long using 24 AWG. Some vendors are offering equipment that can extend this reach to 18,000 feet.

HDSL is more robust than the old T1 service, which required repeaters every few hundred yards. More advanced HDSL equipment, on the other hand, has eliminated many of the problems associated with provisioning T1 service, which resulted in much lower rates for local T1 access.

SOME OTHER FLAVORS OF DSL

Asymmetrical Digital Subscriber Line (ADSL) service came about in the 1992/1993 timeframe as a vehicle for offering video services to the home. Another DSL technology, Rate Adaptive Digital Subscriber Line (ADSL), came along as a means of allowing a transceiver to automatically adjust line speed to attain the highest level of speed over a given loop.

ADSL and RADSL promise to deliver rates of about 7Mb downstream with upstream links of about 1Mb; and while ADSL and RADSL are supposed to run up to 18,000 feet, to get the promised 7Mb downstream, the user would have to be very close to the serving central office (CO). While all of this technology sounds great, one needs to

focus on the real-life application of technology for thousands of end users. Today, the majority of domestic use is focused on the World Wide Web; however, there is a growing number of Small Office Home Office (SOHO) users that require multi-network access as a means of directly or indirectly earning a living. Therefore, it is ADSL/RADSL service that ultimately offers a solution to the needs of the SOHO user.

Today, business applications have grown to stress higher speed access to public and private network infrastructures. The real-life issue is provid- ing access to the corporate or public network without spending large amounts of money for local loop access.

Initially, HDSL technology reduced the cost for provisioning T1 services because it eliminated the need for extensive engineering, expensive repeat- ers, and the huge labor costs associated with deploying traditional copper wire T1 services. The adoption of HDSL technology allowed the LECs to make use of their copper infrastructure to offer more competitive T1 local loops without the burden of high provisioning costs.

Symmetrical Digital Subscriber Line (SDSL) is similar to HDSL in that vendor equipment supports the same line encoding scheme 2 Binary, 1 Quaternary (2b1Q), which avoids any conflict when installed on the LECs' local copper. In addition, depending on vendor equipment and desired line speed, SDSL differs from HDSL in that the loop reach for an SDSL line has the potential of being somewhat greater than HDSL.

Like its predecessor, SDSL technology has become an enabler for high- speed services at a much more affordable price. This is due to that fact that the application of SDSL technology, like HDSL, causes no change in ser- vices and uses the same embedded copper infrastructure as HDSL. Poten- tially, SDSL technology, where applicable, provides lower install and monthly recurring costs for the installation of a circuit capable of support- ing up to 1.5Mb of service.

Perhaps the largest demand for service has come from the Internet com- munity of users and their access to the World Wide Web for the conduct of their business. Today, access to the Web has become the medium of choice for the dissemination of information to business associates and customers alike.

SDSL ENABLES ISP SERVICES

Historically, the gateway to any network — either private or public — was through a local copper loop that connected the end user to the network through the central office. The service providers or networks contracted with the serving utility to provide connectivity at the edge of their network. The end user or customer paid a high price for connectivity to the network, with the monthly price being determined by the available bandwidth of the line. The arrival of SDSL equipment, together with enabling legislation that supports competition at the local level, has changed the service arrange- ments at the local level. Because SDSL equipment requires a physical con- nection to the local loop, an ISP must locate a Digital Subscriber Line Access Multiplexer (DSLAM) within 12,000 to 18,000 feet of a subscriber. The ISP can do this by locating a DSLAM adjacent to the central office. Unconditioned "dry copper" is ordered from the LEC to connect the sub- scriber's location through the central office to the ISP's DSLAM adjacent to the central office.[1] T1, multiple T1, or T3 facilities connect the DSLAM to the ISP central hub. This arrangement gives the end user a direct link to the Internet at speeds of from 64Kbps through to 1.5Mb depending on the dis- tance from the DSLAM.

Another variation is the location of a DSLAM on the premises of an industrial park or campus. Located at a central location (usually the facil- ity's demarcation point [DEMARC]), DSL lines are extended to users around the industrial park or campus area. The DSLAM serves to concentrate all of the DSL lines and provide high-speed access to the Internet via T1/T3 access to a central ISP hub. In a multitenant indus- trial park, it is estimated that at least 200 or more users would be required to make this business model a success, although fewer tenants would work where higher-speed access is required. On the other hand, enterprise net- work operators have found that the application of dry copper can be used to support the connection of local offices or remote buildings that are served by the same central offices. They have also used dry copper and a local DSLAM as a more cost-effective method for concentrating a number of LANs or desktop terminals into the Internet.

Until the advent of more advanced xDSL technology, local loop lengths were limited to end users located within a 2.5-mile radius of the central office. However, the newer SDSL equipment utilizing proprietary technol- ogy can potentially lengthen the reach of a dry copper circuit's effective- ness to 30,000 feet at a lower bandwidth of 64Kbps. However, with incre- mental speed increases to 1Mb, the circuit length conversely would get shorter. For example, Bellcore, the former research arm of the Regional Bell Operating Companies, has conducted SDSL tests with single pair ser- vice, which was extended out to 24,700 feet at 192Kbps. These extensions are possible through continual advances in SDSL technology. For example, TUT Systems of Pleasant Hill, California, deploys a patented process called FastCopper™ technology that removes ambient electronic noise and other distortions in the environment where copper pairs are used. The focus of

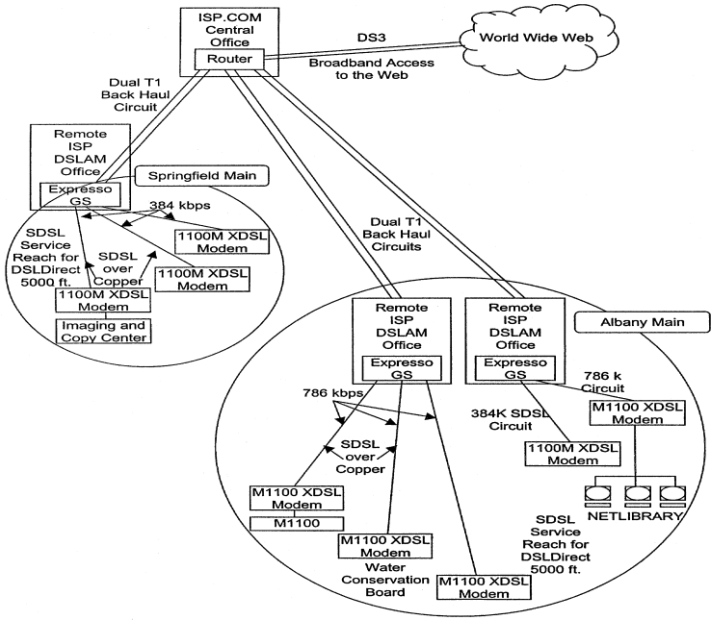

Figure ISP.COM application of TUT Systems Expresso SDLAM.

this technology is aimed at noise reduction circuits, analog and digital sig- nal processing circuits, and digital modulation. This technology makes possible the deployment of dry circuits to within a radius of five to six miles, rather than the more limiting factor of one to two miles. An extension in the central office service range provides support for increased band- width for private metropolitan networks as well as access to the Internet at high speeds for power users.

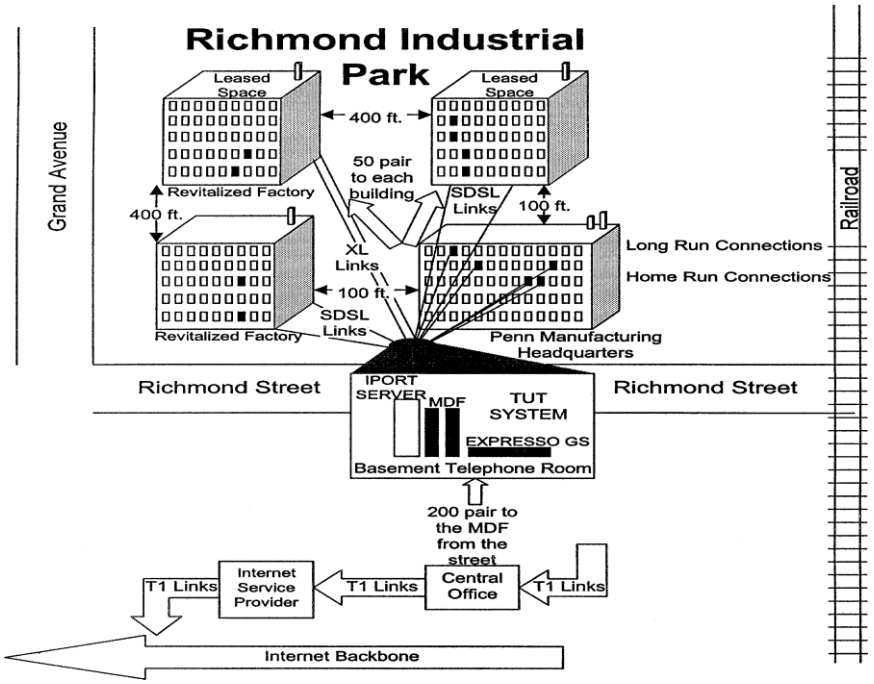

Figure Richmond Industrial Park redeploying copper lines using DSL technology.

DISP.COM

ISP.COM is an Eastern Internet Service Provider, serving eastern busi- nesses with high-speed Internet access. ISP.COM provides direct support to a number of downtown office buildings in several eastern cities. At present, these internal connections are brought to an Ethernet switch at the DEMARC and then in turn brought back to a TUT Systems Expresso DSLAM that is connected to the ISP.COM DS3 central hub via a fiber link. (See Exhibit 57-2.)

DSLDirect

ISP has recently introduced DSLDirect for direct high-speed Internet ser- vice for small and medium-sized Albany and Springfield businesses. The DSLDirect service utilizes TUT Systems Expresso DSLAM units to provide direct access via dry copper, which is a more cost-effective solution than previous provisioning methods. Using these low-cost copper circuits, the ISP is able to provide 384Kbps and 768Kbps Internet connection speeds in direct competition with other larger service providers. The ISP service is not a bridged connection, as is the local RBOC service offering. All of the ISP DSLDirect end-user packets are shipped directly to the Internet, with- out being routed over several switches to the designated Internet switch.

With ISP DSLDirect service, the customer is linked directly to ISP Expresso DSLAM, which, in turn, is linked via multiple T1 links back to the ISP hub and their DS3 connection to the Internet. This in effect provides customers with their own dedicated circuit into the Internet.

The Expresso DSLAM, when first put on a dry copper pair, will adjust its speed depending on the line condition. When it finds its level, the speed remains constant. This has been a very helpful feature because customers can order a 384Kbps line but get a slightly better level of service. Because the level of service on the Expresso DSLAM is

software adjustable from 64Kbps through to 1.2Mb in 64Kbps increments, ISP has a great deal of flex-ibility in serving the needs of its customers using SDSL technology. To use this service, the customer must purchase or lease a TUT Systems 1100M DSL modem/router for about $495, or $30 per month. There is also an ISP charge of about $175 (384Kbps), $275 (786Kbps), or $375 (1.2Mb per month) for the desired level of Internet service. These charges can vary from one region of the country to another.

Presently, DSLDirect is available to businesses located in the immediate vicinity of the Albany Main and Springfield Main (central office) locations. Additional central office sites will be added to the DSLDirect service offer- ing in the future.

ANOTHER SERVICE MARKET

Just about every service provider and equipment manufacturer is expanding into the enhanced multiple dwelling units (MDU) Internet ser- vice provisioning market. This market is comprised of over 40,000 office buildings, 3.5 million hotels, and as many up-market apartment complexes whose tenants have two or more PCs. Companies like Copper Mountain, Paradyne, and TUT Systems have organized special marketing efforts to support the multiple-dwelling market.

For example, TUT Systems has put together an integrated support ser- vice through a network of value-added resellers (VARs) that will provide the hotel or a multiple-dwelling property owner like a real esatate trust (REIT), with a turnkey operation. This turnkey solution utilizes the TUT Systems HomeRun™ 2 technology together with an IPORT™ Internet

System premium Content Billing Platform System to provide an integrated solution called the Connected Community. This is a complete service package sold by local VARs that provides a full package of services.

The advantage of the Connected Community package to the property owner is the fact that there is nothing for the property owner to do. The VAR handles all the arrangements for Internet access with the local CLEC/ISP service provider. All arrangements for equipment configuration and installation are arranged by the VAR. In addition, the equipment can be configured and priced based on the property owner's immediate require- ments. Therefore, costs for equipment can be based on present need. As the property owner's end-user requirements expand, additional equipment can be purchased as needed. Another cost-saving feature is the ability of the TUT Systems HomeRun™ technology to use existing standard tele- phone lines to deliver service without affecting the existing voice service. This solution eliminates much of the cost associated with rewiring a hotel or MDU for data transmission and Internet access.

MORE THAN JUST SPEED

The xDSL market focus until recently has been on speed and low-cost access. While DSL is a very compelling technology, it is worth mentioning that there are other competitors that offer other products and services. For example, a growing number of companies offer cable modems and Internet appliances that will support delivery of competing services over cable.

Recognizing the requirement for something more than bandwidth, ser- vice providers like GTE, U.S. West, Bell Atlantic, and others have announced an ADSL product. With over 700 million phone lines worldwide, service providers can take

advantage of this technology to offer an exten- sive number of voice and data services.

ADSL offers great promise as an alternative to cable because it supports voice and data over the same twisted pair. Further, advanced xDSL line cod- ing algorithms allow for the effective division of the frequency spectrum on copper telephone wire to support voice and data. (See Exhibit 57-4.) With the voice spectrum confined to the 4KHz baseband, the upstream and downstream channels can be dedicated to data. This allows the service provider to offer multimegabit data service while leaving voice service intact. Thus, ADSL connects two different entities — voice and data — all over the same physical wire pair. Recently, the International Telecom Union (ITU) has determined G-lite (G.992.2) to be the standard for ADSL-based service products. Many of the DSLAM manufacturers and service provid- ers have settled on this standard so as to support interoperability. This will greatly simplify the acquisition of CPE (customer premise equipment) and the cross-integration of services across the country and internationally.

Preparing for Cable Modems

Thousands of corporations, universities, government agencies, and individuals creating home pages on servers, while tens of millions of users surfed the World Wide Web. As corporations began to recognize the value of the Internet for building software applications, promoting products and services, and locating as well as disseminating information, the addition of graphics to World Wide Web home pages literally slowed Web surfing operations to a crawl, adversely affecting user productivity. Whereas the replacement of 14.4K bps modem by state-of-the-art 28.8K bps devices has assisted many users in speeding up their Internet search operations, even at that operating rate the display of a typical Web page containing one or two graphic images can result in a delay of 10 to 15 seconds as the picture is “painted” on a monitor.

Recognizing the operating limitations associated with transmissions via the public switched-telephone network, as well as looking for an additional source of revenue, several cable television (CATV) companies initiated broadband access trials to the Internet during 1995. Each of these trials involved the use of cable modems, which enable a personal computer (PC) to access the Internet via a common CATV coaxial cable at operating rates up to tens of millions of bits per second. Although cable modems are in their infancy, both independent market research organizations and many cable operators predict that within a few years, the installed base of this new type of communications device will rapidly grow to over 10 million modems.

Due to the advantages associated with obtaining high-speed Internet access, as well as the potential economics associated with the use of cable modems to obtain such access, data center managers should consider pre- paring their facility for the infrastructure required to use cable modems.

This chapter discusses the nature of cable modems and describes their operation. The scope of the discussion also includes the cabling infrastructure being developed to provide a megabit transmission facility to residences and businesses. The chapter outlines the cabling require- ments for installation within buildings, requirements that are necessary to access this new high-speed information highway via the use of cable modems. The data center manager should have a background of knowl- edge concerning a rapidly evolving new technology and be able to support its use when corporate policy begins to include Internet issues.

MODEM FUNDAMENTALS

The ability to appreciate why cable modems are able to provide a trans- mission capability that is an order of magnitude or more than conventional modems used for transmission on the switched telephone network, requires knowledge of certain transmission concepts, including the Nyquist theorem. This section concentrates on the operation of conven- tional analog modems that are used on the switched telephone network. This can provide the data center manager with an understanding of why analog modems' operating rate is limited and how they may be able to overcome that operating rate limitation.

A conventional analog modem commonly used to transmit information over the switched telephone network is limited to a maximum operating rate of between 28.8K bps and 33.6K bps, with the rate achievable depen- dent upon the quality of the connection and according to the modulation technique employed. In theory, the maximum operating rate of an analog modem that has been designed for use on the switched telephone network is limited by the 4K Hz bandwidth provided by the communications carrier for a switched telephone channel.

In 1924, Nyquist proved, in what is now referred to as the Nyquist theo- rem, that the maximum signaling rate of a device is limited to twice the available bandwidth; beyond that rate, inter-symbol interference occurs and adversely affects the transmission. As an example, for the 4K Hz tele- phone channel, this means the maximum signaling rate of a modem used to transmit on that medium is limited to 8000 baud. Baud is a term used to indicate signal changes per second.

The Quadrature Amplitude Modulation Technique

The most commonly used modem modulation technique, quadrature amplitude modulation (QAM), uses a combination of phase and amplitude to convey the settings of a group of bits in one signal change, enabling four bits to be represented by one baud change. This in turn enables an 8000 baud signaling rate to transport data at a rate of 32K bps when QAM is used for modulation.

Due to the 4K Hz telephone channel limitation, however, data transmis- sion rates are limited to approximately 32K bps, with a slightly higher rate of 33.6K bps recently achieved by a few modem vendors using a modified QAM technique. Although the incorporation of data compression into modems provides a potential doubling to quadrupling of modem through- put, to between 67.2K bps and 134.4K bps, the ability of a modem to com- press data depends upon the susceptibility of data to the compression algorithm being used. Because that susceptibility varies considerably as a modem user performs different operations, the end result is a variable compression rate; even though it is not noticeable during file transfer oper- ations, that variable rate becomes extremely noticeable during interactive operations. In addition, even with the ability to compress data at a high rate, the resulting information transfer rate of 134.4K bps pales by compar- ison to the operating rate obtainable through the use of cable modems. It is clear, however, that advances in modem and cabling technology are lim- ited with respect to increasing the performance of modems used to com- municate via the switched telephone network.

CABLE MODEMS

The key difference between an analog modem designed for use on the pub- lic

switched telephone network and a cable modem is in the bandwidth of the channels they are designed to use. Cable TV uses RG-11 cable for the main CATV trunk and RG-59 cable from trunk distribution points into and through residences and offices. Both types of coaxial cable have 75 ohms impedance and support broadband transmission, which means that two or more chan- nels separated by frequency can be simultaneously transported on the cable.

From Unidirectional to Bidirectional Systems

A cable TV broadcasting infrastructure uses 6M Hz channels within the bandwidth of RG-11 and RG-59 cable to transmit a TV channel. Most CATV systems are currently unidirectional, which means that TV signals are broadcast from the CATV system operator without any provision for receiving a return signal. This transmission limitation is gradually being overcome as CATV operators begin to add bidirectional amplifiers to their networks that, when they are installed, will support transmission from sub- scribers in the reverse direction to conventional TV signal broadcasts. This will enable CATV systems to support the standardized transmit fre- quency range of 5M Hz to 42M Hz, and receive a frequency range of 54M Hzto 550M Hz, with 6M Hz cable TV channels.

By using one or more 6M Hz cable TV channels, a cable modem obtains the use of a bandwidth that is 1500 times greater (6M Hz/4K Hz) than that provided by a voice channel on the switched telephone network. This means that the modem can support a signaling rate of twice the bandwidth, or 12M baud, on one TV channel, based upon the Nyquist theorem, before the occurrence of inter-symbol interference.

The primary difference between cable modems currently being used in field trials is in their use of one or more 6M Hz TV channels within the band of channels carried by a coaxial cable, and their methods of attachment to the CATV network. One cable modem manufactured by Zenith Network Systems, a subsidiary of Zenith Electronics of Glenview, Illinois, operates on 6M Hz channels at 4M bps to the subscriber, using a special filtering technique to prevent data channels from interfering with adjacent informa- tion, which can be in the form of either data or video, that would coexist with the data transmission provided by the cable modem. The uplink or return data rate occurs at 500K bps. Modem modulation is biphase shift key (BPSK), which means that two bits (bi) are encoded in each phase change, and the modem's phase changes are shifted in phase from one to another. This modem is also frequency-agile, which means it can be set to operate on any standardized channel on a broadband CATV system.

The Zenith cable modem is actually one portion of a series of compo- nents required for a PC to use the modem. A complete transmission system requires the use of a Zenith cable modem, Ethernet 10Base-T adapter card with a 15-conductor pin connector, and a 15-conductor shielded cable to connect the cable modem to the adapter. Exhibit 58-1 illustrates the cabling required to connect a PC to a CATV network via the use of a Zenith Network Systems cable modem.

When the adapter card is installed in the PC it, in effect, turns the com- puter into a client workstation. Because the adapter is an Ethernet 10Base- T card, this means that the channel being used by the cable modem oper- ates as one long CSMA/CD Local Area Network, with each PC user competing

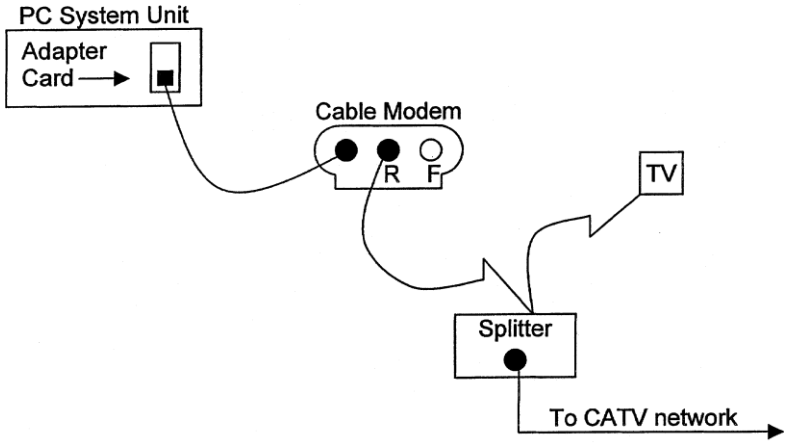

Cabling for a Zenith Network Cable Modem System.

With other PC users for access to the channel. Because of this, the CATV operator should segment its cable distribution system to limit the number of cable modems attached to any segment, similar to the manner in which conventional LANs are limited with respect to the maximum number of workstations that can be connected to the LAN.

The connector labeled "R" on the rear of the cable modem is a reverse cable connector designed for networks that use a single coaxial cable. The second connector, labeled "F," represents a forward cable connector that would be used if the modem were connected to a cable system that uses two cables. In such a system, one cable is dedicated to conventional CATV broad- casting through one-way amplifiers, which precludes reverse transmission on the same cable. This type of system also requires the use of a second cable to obtain a transmission capability in the reverse direction.

A High-Speed Cable Modem Architecture

In addition to the previously described cable modem based upon the exclusive use of RF technology and biphase shift key modulation, Zenith Electronics Corporation announced a high-speed cable modem architec- ture. This architecture is based on the use of 16-VSB (vestigial sideband), a technique developed by Zenith as part of the organization's high-defini- tion research, as well as the 256 quadrature amplitude modulation technol- ogy. Through the use of more complex modulation techniques for which more data bits can be represented by one signal change, the Zenith modem architecture can support data rates up to 40M bps on a 6M Hz cable channel.

Recognizing the fact that many cable TV systems will be limited to one- way transmission in the foreseeable future because of the time and cost associated with upgrading the CATV infrastructure, Zenith plans to sup- port a range of options and speeds for upstream data transmission. According to Zenith, both telephone (analog modulation) and RF return path transmission capabilities will be incorporated into different versions of this new family of cable modems. For many cable modem applications, such as Internet operations, the use of the switched network for a return path should provide an acceptable level of performance. The rationale for this is best noted by examining the communications interaction between a potential cable modem

user and the cable network as a user searches out and accesses various points on the World Wide Web.

On the Web

When users access a Web page, they transmit a universal resource loca- tor (URL) address that represents the document they wish to view. This address is transported using the HTTP within a packet. The HTTP consists of an address that totals fewer than 100 characters, which are used to frame the address to which the message is being transported, as well as the address of the originator of the request. The destination Web server uses the document address to locate the requested page on the server, retrieves it from disk, and forms a packet using the source address from the incom- ing packet as the destination address for the outgoing packets. If the requested document contains a full screen of text, the packet contain close to 2000 characters, because a full screen of text consists of 80 columns by 24 rows of data (i.e., 1920 characters). However, because a typical Web page contains one or more graphics, the total amount of data transmitted from the server to the user will be, in actuality, substantially more than 2000 characters. For example, it is assumed that the Web page in question includes a 3 in. x 3 in. photograph, drawing, or schematic diagram that has been scanned using a resolution of 300 dots per inch. Regardless of the color of the image, each square inch of the image requires 11,250 bytes of storage. If the image was scanned using a 256-color resolution, each pixel requires a byte to represent its color, resulting in 90,000 bytes of storage per square inch. Thus, a 3 in. x 3 in. color image requires 270,000 bytes of storage.

Because HTTP breaks large files into small packets for transmission, the image might be carried by a sequence of approximately 100 packets, each roughly 2700 bytes in length, to include packet overhead. Thus, the short, 100-character transmission from a user can result in a response of 280,000 bytes. Because a user connected to the Web typically clicks on hotlinks that represent document addresses to view other documents, most Web operations represent asymmetrical transmission, that is, more transmis- sions return to the user than the user actually originates. Thus, a high- speed cable channel with a low-speed reverse path occurring over the switched telephone network may actually be sufficient for most data trans-mission applications.

The previously described asymmetrical transmission operation of users was also recognized by Intel Corporation, which took it into consideration when designing its CablePort cable modem system. That cable modem is designed to provide an average downstream data rate of 27M bps and a 96K bps upstream rate. One interesting difference between Zenith and Intel concerning their cable modem systems is in the type of adapter card required to be used in the PC. Because Intel provides a higher downstream operating rate than what is usable by a 10Base-T adapter card, the user must install a Fast Ethernet(100M bps) adapter card in the PC to be able to use the Intel cable modem. Although no commercial costs were provided by Zenith or Intel for field trial operations, it is worth noting that a Fast Ethernet adapter has a retail cost of approximately $250, whereas a 10Base-T adapter can be obtained for less than $50.

A second difference between the Zenith and Intel modems concerns their upstream capability. Although Zenith's new architecture permits support of the switched telephone network for locations where CATV operators can- not provide reverse direction transmission, the Intel system did not offer this capability when this chapter was researched.

Voice and Video on the LAN

VOICE AND DATA CONVERGENCE IN THE LAN has become a hot topic in the industry, thanks to advances in switching and processors, as well as the H.323 standard. This chapter first looks at the business reasons for consid- ering the deployment of voice and video over the LAN and then discusses the technical issues and requirements. Topics include the value of voice and video on the LAN, infrastructure efficiencies, LAN technologies for integrated voice and video, and standards for LAN-based voice and video applications

Most desktops in enterprises today are equipped with two network con- nections a LAN connection to the PC or workstation for data communica- tions and a phone connection to the PBX for voice communications. The LAN and the PBX exist as two separate networks with little or no connec- tivity between them. Each has evolved to meet the very specific and differ- ing needs of data and voice communications, respectively.

Despite much talk in the industry about the convergence of computers and communications, LANs and PBXs have not really moved any closer together during the last decade. In the mid-1980s, some PBX vendors sought to bring data services to the desktop via ISDN technology, but the advent of PCs requiring far more than 64K-bps communications bandwidth favored the emerging LAN standards of Ethernet and Token Ring. So far, most LAN vendors have not attempted to support voice communications on the LAN. But all this is about to change.

There are three key factors at work today that suggest why voice and data convergence in the LAN has become a hot topic in the industry

1. Widespread acceptance of advanced LAN switching technologies, including ATM, which makes it possible for the first time to deliver reliable, high-quality, low-delay voice transmissions over the LAN
2. emergence of the first standard for LAN-based videoconferenc- ing and voice telephony, H.323, which removes objections about the use of proprietary protocols for voice and video over the LAN
3. the deployment of the latest generation of Intel processors, featur- ing MMX technology, which makes high-quality software-based, real- time voice and video processing feasible for the first time, and the new PC hardware architectures with Universal Serial Bus that per- mit voice and video peripherals to be attached without additional hardware inside the PC

This chapter first looks at the business reasons for considering the deployment of voice and video over the LAN and then discusses the tech- nical issues and requirements.

THE VALUE OF VOICE AND VIDEO ON THE LAN

There are essentially two main kinds of motivation for considering voice and video on the LAN the need to support new types of applications that involve real-time communications and the desire to improve the overall cost effectiveness of the local communications infrastructure.

New Types of Applications

Desktop videoconferencing, real-time multimedia collaboration, and video-based training are all examples of new kinds of applications that can benefit from the delivery of voice and video over the LAN.

The uptake of desktop videoconferencing has been held back by a com- bination of high costs and the difficulty of delivering appropriate network services to the desktop. Standards-based H.320 desktop videoconferenc- ing systems require costly video compression and ISDN interface hard- ware, as well as the provision of new ISDN connections at the desktop alongside the LAN and the phone system. New systems based on the H.323 standard and designed to run over the LAN will leverage the processing power of the latest PCs and the existing switched LAN infrastructure, to lower cost and simplify deployment dramatically.

Desktop videoconferencing may be used either to support internal meetings and discussions between groups located at remote sites or to support direct interaction with customers and clients. For example, some enterprises in the mortgage lending business use videoconferencing to conduct mortgage approval interviews with potential borrowers, so as to greatly reduce the overall time to complete a mortgage sale.

Real-time collaboration applications, involving any mix of video and voice with data conferencing to support application sharing and interac- tive whiteboarding, provide a new way for individuals and small groups to collaborate and work together remotely in real time. This emerging class of applications, typified by Microsoft NetMeeting, is being evaluated by many enterprises, particularly for help desk applications.

By contrast, video-based training is already widely used in enterprise LANs. By delivering self-paced video learning materials to the desktop, training needs can be met in a more timely and less disruptive fashion than traditional classroom methods.

The growing popularity of these kinds of applications should be noted by network planners and designers. A preplanned strategy for local LAN upgrades to support voice and video will reduce the lead time for the deployment of these applications and enable the enterprise to move swiftly when the application need has been identified, to obtain the busi- ness benefits with the least possible delay.

Infrastructure Efficiencies

A single local communications infrastructure based on a LAN that han- dles data, voice, and video has the potential to cost less to own and oper- ate than separate PBX and data-only LAN infrastructures.

The average capital cost of a fully featured PBX for large enterprises is between $700 and $750 per user, according to TEQConsult Group, a leading U.S. telecommunications consultancy. Furthermore, this is expected to rise slightly over the next few years as users demand more sophisticated fea- tures from their phone systems. It is not difficult to see how a switched LAN that has been enhanced to handle voice could provide a solution for tele- phony at a fraction of this cost.

Most large PBX installations are equipped with additional facilities such as voice mail and Interactive Voice Response systems for auto-attendant operation. These systems are typically connected directly to the PBX via proprietary interfaces, and they, too, represent major capital investments. With voice on the LAN, such voice-processing applications could be based on open server platforms and leverage the low-cost processing power and disk storage that is a feature of today's PC server market, thereby lowering the system's capital cost still further.

Separate PBX and LAN infrastructures each incur their own manage- ment and operational costs. For example, moves, adds, and changes require separate actions to

patch physical LAN and voice connections and to update LAN log-on and voice directories. With telephony provided over a voice-enabled LAN supporting combined directory services, the manage- ment effort required to administer moves and changes would be substan- tially reduced.

These cost of ownership benefits come with a raft of usability improve- ments for telephony. The PC (with phone handset attached) becomes the communications terminal for making and receiving phone calls, and the processing power and graphical user interface of the PC can be leveraged to provide point-and-click call launch and manipulation. Features of PBXs such as call transfer, divert, and hold, which are hard to invoke from a phone keypad, become very easy to use from a Windows interface.

Incoming callers can be identified on the PC display by matching Calling Line Identifier with directory entries. And with voice mail and e-mail sup- ported on a unified messaging platform such as Microsoft Exchange or Lotus Notes, all messages are accessible and manageable via a single user interface.

These usability benefits for voice telephony over the LAN also extend to videoconferencing — a single consistent user interface may be applied to both video and voice-only calls.

LAN TECHNOLOGIES FOR INTEGRATED VOICE AND VIDEO

The LAN technologies in widespread use today — Ethernet, Fast Ether- net, FDDI, and Token Ring — were not designed with the needs of real-time voice and video in mind. These LAN technologies provide "best effort" delivery of data packets, but offer no guarantees about how long delivery will take. Interactive real-time voice and video communications over the LAN require the delivery of a steady stream of packets with very low end- to-end delay, and this cannot generally be achieved with the current LAN technologies as they stand.

Asynchronous Transfer Mode (ATM)

At one time, there was a belief that ATM networking to the desktop would be embraced by LAN users to solve this problem. ATM is a network- ing technology that was designed specifically to handle a combination of the low-delay steady-stream characteristics of voice and video and the bursty, intermittent characteristics of data communications.

The ATM Forum, the industry body responsible for publishing ATM specifications, has developed a number of standards that enable desktops connected directly to ATM networks to support existing LAN data applica- tions as well as voice telephony and videoconferencing. The ATM Forum standards for the support of voice and video over ATM to the desktop typ- ically avoid the use of traditional LAN protocols such as IP, and instead place the voice or video streams directly over the ATM protocols.

While it is clear that ATM to the desktop provides an elegant and effec- tive solution for combining voice, video, and data over the LAN, this approach does imply a "forklift" to the LAN infrastructure and the end sta- tion connection. The cost and disruptive impact of such an upgrade tend to limit its appeal, and as a result desktop ATM is not expected to be widely adopted.

However, the ability of ATM to provide "quality of service" — that is, to deliver

real-time voice or video streams with a guaranteed upper bound on delay — makes it an excellent choice for the LAN backbone where voice and video over the LAN is needed.

Shared and Switched LANs

It is generally accepted that shared LANs are unsuitable for handling real-time voice and video because of the widely varying delays that are seen when multiple stations are contending for access to the transmission medium. The CSMA/CD access method used in shared Ethernet is particu- larly poor in this respect. Token Ring, on the other hand, is based on a token-passing access method with multiple levels of priority. Stations wait- ing to send data packets can be preempted by other stations on the ring with higher priority voice or video packets to send. As a result, Token Ring has excellent potential to handle real-time voice and video traffic, though this potential has yet to be realized in currently available networking products.

LAN switching does much to overcome the limitations of shared LANs, although today's products are still a long way from providing an answer for voice and video over the LAN. It is now cost-effective to provide users with dedicated 10M-bps Ethernet connections to the desktop and 100M-bps Fast Ethernet uplinks from the wiring closet to the backbone.

However, despite the vast increase in bandwidth provision per user that this represents over and above a shared LAN scenario, there is still conten- tion in the network leading to unacceptable delay characteristics. For example, multiple users connected to the switch may demand file transfers from several servers connected via 100M-bps Fast Ethernet to the back- bone. Each server may send a burst of packets that temporarily over- whelms the Fast Ethernet uplink to the wiring closet. A queue will form in the backbone switch that is driving this link, and any voice or video pack- ets being sent to the same wiring closet will have to wait their turn behind the data packets in this queue. The resultant delays will compromise the perceived quality of the voice or video transmission.

The only way to overcome this problem is to find a way of treating real- time voice and video packets differently from data packets in the network and to give them preferential treatment when transient data overloads cause queues to form on busy network links. In practice, this means that LAN packets must be tagged with some kind of priority information that enables switches to identify which packets need to jump the queue. The IEEE 802, which oversees standards for LAN technologies, has initi- ated a project identified as 802.1p, which is concerned with "Traffic Class Expediting" in LAN switches.

The principal problem faced by 802.1p is that there is no spare informa- tion field in the standard Ethernet packet format that could carry the required priority tag. As a result, it has been necessary to propose a new Ethernet packet format with an additional 4 bytes of information in the packet header that can contain a 3-bit priority tag field (offering eight lev- els of priority), together with some other information concerned with vir- tual LANs.

With the new Ethernet packet format containing a priority tag, end sta- tion applications can identify real-time voice or video packets by assigning them a high-priority value in the tag. LAN switches that have been enhanced to process the priority tags can separate high- and low-priority traffic in the switching fabric and place them in separate queues at outgo- ing switch ports. The LAN switches need to implement a

queue scheduling algorithm that gives preference to the higher priority queues on outgoing ports, and by this means it is hoped that real-time voice and video can be carried over the LAN without incurring unacceptable delays during peri- ods of heavy data traffic.

Hybrid ATM Networks

The discussion of ATM described how it offers guaranteed quality of ser-vice for real-time voice and video streams. Today, ATM is increasingly used as a LAN backbone for pure data applications, because it offers greater scalability and fault tolerance than other LAN technologies. Ethernet and Token Ring LANs are connected to ATM via "edge switches" equipped with ATM uplinks, typically supporting the ATM Forum standard for carrying LAN traffic over ATM, know as LAN emulation.

It is possible to enhance ATM edge switches to enable desktops con- nected via Ethernet or Token Ring to enjoy the benefits of ATM quality of service across the LAN backbone. Two techniques have been proposed to achieve this.

The first technique, known as "Cell-in-Frame," extends the native ATM signaling protocols over dedicated Ethernet connections from the edge switch to the end station. The voice or video application in the end station places the voice or video stream in ATM cells using the ATM Forum stan- dards for native ATM transport, and then encapsulates the ATM cells in Ethernet packets for transport to the edge switch for onward transmission onto the ATM network. Effectively, this is ATM to the desktop, but using physical Ethernet with standard Ethernet adapter cards as a kind of phys- ical transport layer for ATM traffic. The second technique makes use of an emerging standard protocol for end stations to request quality of service for IP-based voice or video appli- cations, known as the Resource Reservation Protocol, or RSVP. The enhanced edge switch intercepts RSVP requests originated by end stations and converts them into ATM signaling to request the setup of connections across the ATM backbone with the appropriate quality of service. The edge switch then distinguishes between IP packets containing data and those containing voice or video, using the information provided by RSVP, and steers voice and video packets onto ATM connections that have quality of service.

At the time of writing, the technique described here for RSVP-to-ATM mapping enjoys somewhat broader industry support than cell-in-frame, perhaps because of its relationship with Internet technology.

Until LAN switches supporting 802.1p priority tagging have proved themselves capable of meeting the very stringent end-to-end delay require- ments for real-time voice and video communications, hybrid approaches based on ATM in the backbone and switched Ethernet or Token Ring to the desktop are likely to find acceptance as the solution of choice for voice and video over the LAN.

Standards for LAN-Based Voice and Video Applications

Standards for voice and video over the LAN fall into two categories those designed for native ATM protocols and those intended for general- purpose LAN protocols, particularly IP.

Standards for native ATM protocols, such as the ATM Forum's Voice Telephony over ATM (VTOA), are appropriate only for ATM-connected desktops or desktops

running Cell-in-Frame over Ethernet.

Standards for applications that run over IP are applicable both to ATM- connected desktops as well as desktops in general Ethernet or Token Ring environments. The most important standard in this space is H.323, which was developed by the International Telecommunications Union. While
H.323 is designed to be independent of the underlying networking proto- col, it will most often be deployed running over IP.

H.323 references other existing standards for the digital encoding and compression of voice and video signals and describes how audio and video streams are carried in the payload of IP packets with the aid of the Real Time Protocol (RTP), which provides timing and synchronization informa- tion. H.323 also covers the handling of data streams for application shar- ing, shared whiteboarding, and real-time file transfer (referencing the T.120 standard) and includes signaling based on ISDN messaging protocols for call setup and teardown.

The H.323 standard is flexible and accommodates any combination of real-time voice, video, and data as part of a single point-to-point or multi- point conference call. It may be used with a voice stream alone as the basis of a LAN telephony solution. H.323 enjoys the broadest support in the industry as a proposed standard for Internet telephony.

Additional Components Gateways and Gatekeepers

Creating a LAN infrastructure that can consistently deliver voice and video streams with sufficiently low delay is an absolute prerequisite for integrating voice and video on the LAN, but it is by no means the complete answer to the problem. There are two other key components of a complete voice and video solution, which in H.323 parlance are known as the gate- way and the gatekeeper.

An H.323 gateway provides interconnection between voice and video services on the LAN and external voice and video services typically pro- vided over circuit-switched networks such as ISDN and the public tele- phone network. The gateway terminates the IP and RTP protocols carry- ing the voice and video streams and converts them to appropriate formats for external networks. For videoconferencing, the conversion is most likely to be to H.320, another ITU standard that specifies how voice and video are carried over ISDN connections. For voice-only connections, the conversion will be to the G.711 standard for digital telephony. This allows voice interworking with any phone on a public network or connected to a PBX.

An H.323 gatekeeper is a pure software function that provides central call control services. While it is possible to run H.323 voice and video com- munications over the LAN without a gatekeeper, in practice this function is extremely useful. At the most basic level, the gatekeeper provides direc- tory services and policy-based controls applied to the use of voice and video communications. For example, the gatekeeper can bar stations from accessing certain types of external phone numbers at certain times of day. The gatekeeper can be thought of as the “server” in a client/server model of LAN-based telephony and videoconferencing.

At a more sophisticated level, the gatekeeper may be able to support supplementary services, including call transfer, hold and divert, hunt groups, pickup groups, attendant operation and so on — features that are typically found in high-end

PBXs for controlling and managing voice calls. While the H.323 standard does not explicitly describe how supplementary call control features may be supported, the standard does provide a frame-work for the addition of these advanced capabilities.

A BRIEF HISTORY OF VOICE COMMUNICATIONS

Communications via telegraph, radio, telephone, and cellular technol- ogy each struggled through an early period of disbelief, limited accep-tance, and technical hurdles. Each communication medium also had the potential to provide substantive solutions to real needs. Commercial acceptance and widespread deployment came only with the creative and effective application of these technologies.

For the moment, available Internet telephony packages have captured the curiosity and excitement of consumers in the same manner citizens band (CB) radio entranced the American public in the 1970s. CB was a con- venience technology that lured users by the millions with the promise of free and easy communication with friends, family, and business associates. Demand grew so dramatically that the Federal Communications Commis- sion (FCC) was forced to open the citizens band spectrum from 23 to 40 channels.

While CB is still active as a commercial and emergency communications tool, the vast majority of radios built now lie dormant in closets, garages, radio repair shop parts bins, and landfills. What happened?

Unpredictable Service Can Doom a Good Idea

Congestion happened, among other things. For many once-impassioned CB cowboys, cowgirls, and rangers, the thrill of the fad wore off quickly with the realities of use. The excitement of anonymous pranksterism was dulled when users became equally susceptible to the same.

Security was nonexistent. Solar noise often relegated useful medium- and long-range communications to late-night hours. Hackers with illegal 100-watt linear amplifiers would dominate the channels with the best prop-agation and flame anyone who dared to talk. The inconveniences of terrain, keeping antennae properly tuned, and coordinating important calls were also discouraging to many once enthusiastic CD users.

Worse, there simply was too much demand for limited bandwidth. The sheer number of simultaneous users on a given channel produced back- ground noise levels that were difficult or impossible to communicate over, even using the best equipment at short range.

The unregulated load placed on the spectrum resulted, effectively, in lost or error information that made the medium nearly unusable. With- out quality of service guarantees, commercial implementation was limited to the very few whose needs were well-suited by this unpredictable service.

STATE OF INTERNET TELEPHONY TODAY

The Evolution

Internet telephony is evolving through a period of CB radio-like applica- tion. Wander into one of the Internet Relay Chat (IRC)-based voice call serv-ers and you will

find hundreds of users chatting about a wide variety of interests.

Some users coordinate offline to meet in private conversations in lieu of a standard telephone call. Many congregate in multiuser chat sessions on numerous topics of common interest — a natural evolution of the IRC relay chat channels. A few, cutting their teeth on limited demo versions of soft- ware, place calls randomly into any open channel in search of a modern "radio check."

Current Audio Products and Applications

But these applications represent old technology now. In fact, updated, improved versions of Internet audio and telephony software surface almost daily.

The energy in this industry manifests itself with an intense competitive urgency. New players and products are introduced and fade, accepted, rejected, or absorbed. The dominant vendors seek to outpace each other and grab market share. Alliances are struck, technologies acquired, and Web page press releases trumpet feature lists that would make a PBX salesperson envious. This industry is vibrant, alive, and here to stay — but in what form?

There are three basic audio product types approaching maturity on the Internet

1. Audio broadcast. Products in this category provide real-time, one- way transmission of press conferences, announcements, or enter- tainment such as music and talk radio.
2. Group conferencing. These products enable multiuser voice confer- encing.
3. Telephony. Such products enable person-to-person telephony via a personal computer or workstation, conference calling, and voice mail. Some products provide whiteboarding capabilities, permit multiple simultaneous calls, and support collaborative computing.

THE WORKINGS OF INTERNET VOICE TECHNOLOGIES

Internet audio and telephony applications employ efficient software- driven codecs on a personal computer or workstation to digitize and pack- etize voice information for transport via internetwork protocols such as SLIP, PPP, TOP, or UDP. These applications use the computer's multimedia hardware for input/output devices (e.g., microphones and speakers) and analog-digital conversion.

Delay. Delay-handling mechanisms are implemented in most commer- cially available packages, since voice communications are particularly intolerant of delay and delay variation. While not as disconcerting as prop- agation delay over satellite, significant but tolerable delay is noticeable in these applications at most connection speeds. Delay is only significant in two-way communications.

Delay is accounted for by buffering a certain amount of voice informa- tion in order to compensate for variations in network transit time between callers. While this actually adds slightly to the overall delay perceived by the user, it is a necessary acknowledgment of the random, unpredictable delay present in today's Internet.

One-Way Audio Broadcasts. Broadcast applications enable entities such as radio stations, news services, or corporations to transmit one-way audio to a potentially unlimited number of users. A centralized server dig- itizes the audio and either transmits it in real time to Internet-attached users, or stores the compressed audio on the server.

If users are unable to attend the broadcast, they may later access a com- pressed audio file on the server. Users are not required to download the file to their local workstation in order to listen; the file may be played out in real time by the server, avoiding a long, unproductive wait.

Several Internet-only "radio stations" have already sprung into exist- ence, offering an eclectic variety of music, talk, and commentary program- ming. One intriguing advantage of Internet-based audio programming is the ability to archive shows for access when the user finds it convenient to lis- ten. If a favorite show or episode is missed at broadcast time, it may be accessed remotely at a later time or even downloaded to the user's local disk drive for repeated listening.

Two-Way Conferences. Two-way conferencing applications are very pop- ular among noncommercial users. A modified chat-server implementation is used to connect multiple users. Half-duplex communication is the norm, but some available applications allow limited whiteboard and collabora- tive computing capabilities.

Connection Speeds and Sampling Rates. Voice quality is directly affected by connection speed and the quality of the codec design. While some of the available codecs are capable of sampling rates above 35 KHz, only a minor- ity of users currently have Internet access speeds that top 14,400 bps. All available software specifies a minimum connection rate of 14,400 bps, which provides usable but grainy voice quality.

With the better codec implementations, voice quality approaches then surpasses that of toll-quality conventional public telephony because higher connection speeds allow greater sampling rates. This is an exciting promise when contrasted with the fixed 300 Hz to 3,000 Hz bandpass and 8 KHz sampling rate used in conventional telephony.

Software Enhancements. Because Internet telephony applications are software-based, creative enhancements to the virtual telephone set are possible. These enhancements come in the form of integration of familiar tools, such as autodialing, address lists, directory services, caller ID, note- pads, and voice mail. More sophisticated users may opt for concurrent support of whiteboarding and other image transfer and collaborative com- puting. The potential for integration with other software applications also warrants consideration.

USAGE ISSUES AND IMPEDIMENTS

Several success factors have been achieved as Internet voice applica- tions have matured. Codec design, voice quality potential, and usability have been well-received if shelf sales and trial downloads are any indica- tion of market interest.

The response of power users and techno-junkies does not, however, necessarily indicate a broad, long-term acceptance of these technologies. For Internet voice products to continue to thrive, several obstacles have yet to be overcome.

Support for Full-Duplex Operation. Only a small portion of existing sound cards support full-duplex operation, limiting even full-duplex-enabled tele- phony software to half-duplex capability. Given the rapid growth in home computer sales, and the dearth of installed multimedia hardware in exist- ing business computers, this should not be a long-term concern.

Likewise, 28.8K bps modems are rapidly becoming the norm, and high- speed

cable access technology trials promise greater availability of usable multimedia bandwidth. (What user wants to abandon a multiuser role-play- ing Net game or telecommuting session to take an Internet phone call?)

Interoperability Standards. Of greater concern is the general lack of standardization in codec and transport implementations. Few Internet voice implementations interoperate, requiring callers to use the same software to communicate.

Ease of use leaves much to be desired. Currently, both users must have an active Internet connection in order to place a call, which for many requires that a dial-up modem connection be made to the Internet before attempting to locate and connect a voice call to another user. One notable exception at this time is an Internet service provider (ISP) that proposes to allow a user to connect across the Internet through a server to a conven- tional telephone switch, from which users of traditional telephones may be dialed.

IP Addresses. Divining a user's Internet "telephone number" can also be challenging. All of the approaches currently available require significantly more coordination between users than conventional telephony.

Some products currently require users of a given application to use a private IRC-based server network. Some products allow "dialing" of a spe- cific IP address, which gives more universal access to other users, should their client software be compatible with yours and they have a fixed address. Newer releases allow scanning a wide range of IP addresses for user IDs, which makes it possible to locate a specific called party, should they actually be online with their ISP at the time you wish to call.

An alternative approach to locating users with dynamically assigned public IP addresses employs an e-mail page to notify the called party that you wish to place a call to them.

Unpredictable Nature of the Internet. Perhaps the greatest limiting factor is the unpredictability of the Internet itself.

Traffic management is almost nonexistent, providing little or no quality of service guarantees to users. The vast majority of Internet voice applica- tions are designed to manage delay and traffic loss only within strictly defined limits.

Internet provisioning practices do not account sufficiently for traffic loss, congestion, and delay. Current practices almost universally involve throwing more bandwidth and routers at the problem, which only provides temporary relief and does not ensure fairness among users. Retransmis- sion of lost broadcast audio packets is somewhat acceptable within strict limits, but multisecond delays and large-scale discards render two-way voice communications unusable.

Multiplatform Support. Support for non-PC platforms is currently very limited, though it is only a matter of time before broader support of Macin- tosh and UNIX-based platforms is common.

Internet access capabilities are also being delivered in PBX platforms and PC-based telephony servers. While the PBX approaches are primarily aimed at 56K bps integrated services digital network (ISDN) Internet data access, it is only a small jump to providing Internet voice services on the same platform.

FUTURE DIRECTIONS

Technology Outlook

Rapid advances in hardware and software technology offer much hope for Internet telephony. Faster, more capable platforms will enable more sophisticated codec implementations, and improvements in peripherals such as sound cards, microphones, and modem speed will enhance voice quality. Multitasking and collaborative applications will benefit from increased platform capacity and connection speed as well.

Given sufficient processor power and connection bandwidth, sophisti- cated multimedia capabilities may be integrated with Internet telephony, bringing the capabilities of the corporate conferencing center to the desktop.

In addition to conversation, applications that employ images, video, whiteboarding, and console-sharing may be performed in real-time. Newer high-end home and business PCs approach these capacities today. Internet access speeds are increasing steadily and will achieve widespread avail- ability in megabit increments within the next 10 years.

The growing acceptance of telecommuting will have a complementary impact upon the acceptance and deployment of Internet voice applications.

Success Factors

If Internet voice is to gain broad and permanent acceptance, it must offer value equivalent or superior to the existing public switched telephone network. Like the common telephone, interoperability must be universal and without question. Directory services, numbering, and billing must be effectively dealt with. Dialing must be effortless and intuitive, and feature sets must be standardized. While competitive pressures are understand- ably high, vendors and service providers must come to agreement on core standards and interoperability issues.

ALL-IMPORTANT INFRASTRUCTURE ISSUES

Quality of Service Is Everything

Commercial radio, the public-switched telephone network, and cellular telephony have survived because commercial providers have taken the steps necessary to ensure that a minimum quality of service (QOS) can be guaranteed of their respective offerings.

Quality of service encompasses many factors that vary by the service offered. For voice communications, the most important factors are

- Reliability. The service must be functioning at least 99.8 percent of the time. Users must have an extremely high degree of confidence that the network will function and provide the desired service every time it is used.
- Performance. The service must provide a minimum guaranteed level of performance that the user finds desirable and of value.
- Predictability. The service must provide consistent, predictable per- formance at or above a specified minimum level of quality.
- Fairness. The service must provide fair, equal access to network resources for all users.
- Accessibility. The service must be readily and easily available to a number of users

sufficient to make the service of value to subscribers.

Each of these quality of service factors is affected by a number of infra- structure design issues that require analysis, planning, and monitoring by the service provider in order to ensure the long-term viability of the service.

Becoming as Reliable as the Telephone. During the years of federal regula-tion of the public telephone network, the telephone came to be considered an essential item of daily life for more than conversation. To be more than another household or business accessory, Internet telephony must at least come close to meeting the reliability of the conventional telephone.

Telephone networks have been designed and built for decades to have sufficient redundancy to function in all but the most catastrophic circum- stances. Floods, earthquakes, tornadoes, hurricanes, and other natural disasters are examples of conditions that challenge even minimal network operation yet place a terrible urgency upon the provision of at least a mar- ginal level of service. Telephone switch central offices operate on self-sus- taining, battery-backed power systems for this reason. Redundancy in power systems, switching hardware, cable capacity, and cable paths are all considered to be essential and critical baseline design criteria for any public telephone system. While federal deregulation has somewhat eased the metrics by which public telephone networks are built, a deeply ingrained cultural design ethic exists within the telephony com- munity that preserves such conservative design practices.

Likewise, the public customer base has for generations grown up with an unquestioning dependence upon the reliability of the network. If the power goes out, one generally expects to be able to pick up the telephone and notify the power company. Day after day, customers depend on the telephone and think nothing of it, unless it fails to function.

Responsibilities of Service Providers. Despite federal deregulation, the tele- phone service industry remains heavily regulated. Service providers are required to contribute significantly to infrastructure development, to main-tain minimum service levels, and to guarantee universal access to all users, including accommodation of special needs such as TDD terminals for the hearing-impaired or low-cost basic service to the infirm.

As of this writing, the rapid rise in Internet telephony has caused many public telephone service providers to protest to regulatory bodies. The perception is that if Internet telephony is to be given a free ride without contribution to infrastructure pools such as the Universal Service Fund, conventional service providers will be at a competitive disadvantage.

Notably absent from the protest are telephone service providers who were adventurous enough to also have well-established Internet service offerings. These regulatory issues are sure to be addressed as the Internet matures.

It seems unlikely that Internet telephony will completely supplant the existing telephone network for quite some time. Internet voice quality and service levels require dramatic improvement. To achieve this level of quality and service, massive investments made over decades by the public tele- phone service providers would need to be matched. Internet access devices and circuits must be simplified to an appliance level and match the reliability and survivability of the existing network.

Moreover, the core infrastructure of the Internet itself requires switch- ing equipment and design practices that make efficient use of wide-area bandwidth while providing redundancy in switching capacity, subsecond rerouting, and intelligent, dynamic bandwidth allocation. All of these core network factors depend on the inclination and ability of the service provid-ers to develop the Internet infrastructure to such a level.

Performance Metrics — Quantifying "Perception." Service performance includes many elements, such as the time required to connect a call, peak instantaneous call capacity, voice quality, noise levels, and the ability to support enhanced services. The capacity of the overall network to support calls at a guaranteed minimum level of quality under normal and disaster conditions is also significant.

Traditional metrics used in the design and benchmarking of data net- work performance focus on objective factors such as error rate, through- put, and latency. Voice networking exposes an entirely different range of subjective, human perception factors that are difficult to quantify yet are critical to the practicality of a commercial voice offering.

Bell Labs long ago developed a standard, the Overall Reference Equiva- lent (ORE), to attempt to bring the human factor to bear in the design of the public telephone network. The ORE metric quantifies the user's tolerance for circumstances, such as time to connect to another user, blocked calls due to network overload, echo, volume, and noise. These metrics are a fun- damental component of all current voice network design and testing processes.

A successful mass deployment of Internet telephony will hinge on the same factors. The Internet infrastructure and user devices will require low delay to inhibit echo and double-talk. Low delay variation and low data loss will be necessary to avoid distortion and noise.

High-performance switching and efficient utilization of switching, trunk- ing, and access capacity in the Internet backbone will be essential to achieving these goals. Backbone switching hardware and software must be scalable to accommodate high-capacity, long-term growth of the network that will be necessary to effectively meet user demand.

Predictability — Key to a Universal Service. While performance and reli- ability are paramount, consistent performance of the network is also of extreme importance.

Users will expect a universal service to perform in essentially the same manner every time it is called upon. If a significant variance in performance exists from call to call, the lower range of performance will be perceived as unacceptable. This will be cause for contention between service providers and clients and can only be addressed through effective service network design and provisioning practices.

As Internet access and traffic grow, we are already witnessing dramatic variances in a given server's performance as a consequence of demand on the host processor and contention for transport bandwidth. Much like today's commercial telephone switches, sufficient resources must be pro- visioned at the host site to ensure that the desired level of service is pro- vided to users.

In addition, wide-area transport and switching capacity in the Internet backbone require significant improvement in order to avoid the unpredict- able network overloads and outages seen today. Delay, throughput, and error rate must remain relatively

consistent not only throughout a call, but from call to call.

Fairness — Meeting Guaranteed QOS Levels. What is fairness? Users should be able to depend on a level of service proportional to their invest- ment in the service, regardless of the state of the network at any given moment.

For example, a user with an expensive high-speed connection should expect more total throughput during network congestion or impairment than a user with a less expensive low-speed connection. Each user should receive a degree of service that is directly proportional at all times to their contracted quality of service.

In a similar fashion, active user connections should not have their throughput, delay, or error rate reduced below minimum guaranteed levels at any time as a result of new connection requests or variance in demandby other established connections.

Fairness also includes the assurance that once a call is placed and in progress, quality of service cannot be disrupted or usurped by new calls. Fairness of this sort equates to an "all circuits are busy" message from the public switched telephone network. The implicit message is that other users were there first, and courtesy (quality of service) dictates that an in-progress call may not be disconnected or impaired in preference for a new call.

The current frame-switched Internet backbone architecture generally allows unfair, unpredictable degradation of service to established connec- tions due to its minimal support and enforcement of quality of service metrics.

Accessibility to a Broad User Base. Virtually universal access exists in today's public switched telephone networks. Practically anyone in the world desiring connection to telephone service may obtain it.

Likewise, while the worldwide telephone network is composed of multi-tudes of independent local, regional, national, and international service providers, interoperability is truly universal. Any caller on any telephone service may connect to anyone in the world with a telephone, including cellular and commercial mobile-radio telephones. Internet telephony requires the same degree of interoperability.

Of course, public telephone networks were of value long before such wide access was provisioned. Internet access is today sufficiently broad to support an Internet telephony user base. The promise of interoperability With the billions of installed conventional telephones is extremely encour- aging. In fact, Internet-style telephony seems to have quite a strong chance of becoming the dominant voice communications medium in the next cen- tury. Instrument costs (e.g., PC sound cards and software) will certainly be driven continually downward, and stand-alone consumer Internet tele-phone instruments are absolutely feasible.

Building an IP PBX Telephony Network

INTERNET TELEPHONY HAS BEEN INCREASINGLY EXPLORED and implemented as a viable communication tool in large corporations. A main component of enterprise IP voice is the IP PBX, which functions the way a traditional PBX does. It allows calls to be transferred throughout the organization, it allows easy intra-enterprise calls, and it operates automatically.

An IP PBX is different in almost every other respect. Not only is it easier and less costly to operate and maintain, it operates with different technol- ogy. The IP PBX has

paid off for the corporations using it through reduced manpower and by eliminating an entire (telephone) network. This chapter provides other payoff ideas and an explanation of the technology behind the IP PBX.

THE PBX

Yesterday's PBX fulfilled a simple need it allowed users to talk together, and also allowed users to talk out to the PSTN (public switched telephone network). PBX (premise branch exchange) manufacturers fulfilled this need by installing a mainframe computer into the enterprise and connect- ing a proprietary line card interface to either analog phones or proprietary digital phones. The connection out to the PSTN was established through a trunk interface card.

Today's PC-based PBX similarly fulfills a need. Phones on the enterprise side and the PSTN on the outside can be connected together. The approach with a PC-based PBX is fundamentally the same as the mainframe PBX architecture. The big difference is the use of relatively inexpensive PCs instead of hefty mainframe computers.

The third generation, tomorrow's PBX, is the IP (Internet Protocol)- based PBX. Again, it fulfills a by-now well-known need, but with a lot of other benefits. Instead of using a line interface card and circuit-switched card , it uses the TCP/IP network switching voice packets through an Ether- net, ATM, frame relay, ISDN, or whatever satisfactorily carries TCP/IP.

THE IP-PBX

Full PBX capabilities over IP LAN/WAN networks promise to substitute and replace traditional enterprise PBXs, and are an important step toward full voice and data convergence. In the IP PBX, voice traffic is digitized and compressed, placed into data packets, and transmitted across the packet network directly between the stations or WAN interfaces. End stations communicate with a call control server only when a call processing func- tion, such as transferring a call, creating a conference call, or sending a call to voice mail, is required or requested.

Standards and the IP PBX

An IP PBX operates within the ITU (International Telecommunications Union) Standards (H.323 and T.120) that define how data equipment worksin a data environment and define the signaling, call control, and audio com- pression for packet delivery of voice and video communications on IP net- works. Without these standards in place and strictly followed, interopera- bility would not be possible.

Components

An IP PBX requires three components the desktop telephone, call man- ager software, and a WAN/IP gateway. These three components are attached to existing LAN/WAN infrastructure.

The Desktop Telephone. Users have two desktop phone choices

1. an IP Ethernet phone that plugs directly into an Ethernet jack
2. handsets or headsets that plug into their PC

The IP Ethernet telephone resembles a normal digital PBX set, but instead of connecting to a proprietary PBX port, it plugs into a standard Ethernet LAN jack. An IP telephone delivers audio quality comparable to that of a PBX telephone and is easy to use with single-button access to line appearances and features. The IP telephone can operate as a standard IP device with its own IP address. A fully H.323-compatible IP phone can talk to any other H.323 device. The following are key characteristics of the IP telephone.

- connects directly to any 10 Base-T Ethernet (RJ45) network
- programmable buttons for features, speed dialing, or line appearances
- IP address and signaling (TCP/IP) to call manager
- H.323 standards
    - built-in compression G.711; G.723 (ITU standards), on a call and feature basis
    - IP address assignment and configuration with DHCP keypad or BootP
    - administration and button configuration through a Web browser
    - built-in encryption for privacy protection during voice conversation
    - 3rd-pair or phantom powered to permit power backup in the event of building power failure
    - one-button collaboration (T.120) with PC and NetMeeting for features such as application sharing, video, chat, and whiteboarding
    - built-in repeater port for cascading Ethernet devices

The Call Manager. The call manager provides the network intelligence to enable simple-to-use and feature-rich IP communications. Call manager software is designed to work seamlessly with existing telephony systems (PBX or Centrex) or can provide full PBX functionality on its own. It can be deployed as a single IP PBX in a single office, or as a single IP PBX with mul- tiple geographically dispersed users. With total switch and network inde- pendence, administrators can create a truly virtual campus environment utilizing a common Web browser.

By installing the call manager software on a Windows NT server in the IP network, features such as call, hold, call transfer, call forward, call park, caller identification, and multiple line appearances are provided to the IP phone. The SMDI interface on the call manager provides connectivity to various voice mail and IVR systems along with CDR reporting for call accounting and billing.

The call manager provides the call processing functionality for the IP PBX. It manages the resources of the IP PBX by signaling and coordinating call control activities. The call manager sets up a call by instructing the calling party to set up an RTP audio stream to the other device, either tele- phone or gateway. Once an audio stream is set up between two devices, the call manager is idle until a new request (such as transfer or disconnect) is made. In the event the call manager fails during a call, the two parties stay connected and can complete their call. Various signaling protocols, such as Q.931 for ISDN WAN control and H.225/H.245 for IP packet control, are man- aged and controlled by the call manager.

The call manager also manages calling zones to ensure efficient band- width performance at the maximum audio quality. When a call is routed over a low-bandwidth IP pipe, the call manager will instruct the IP phone to use a lower bit rate audio compression, such as G.723. For calls toward the PSTN, the call manager will have the

phones use G.711, which is the com- pression required for PSTN calling.

The call manager offers a standard directory service that allows other applications on the network to access the call directory. It can be overseen Via a Web browser and provides remote management for diagnostics and maintenance from anywhere in the world. The browser provides an intui- tive interface for administrators and users. Upon administrator approval, users can access and configure their own phones. Call records are kept ina standard CDR database for billing and tracking activity.

The Gateway. IP-based telephony systems today need to connect to the PSTN and the existing PBX. Gateways are specifically designed to convert voice from the packet domain to the circuit-switched domain.

The gateway converts packetized voice to a format that can be accepted by the PSTN. Since the digitized format for voice on the packet network is often different than on the PSTN, the gateway will provide a type of conver- sion called transcoding. Gateways also pass signaling information.

Based on the various PSTN interfaces, there is a need for both a digital and analog trunk version. Gateways must all support supplementary ser- vices, such as call transfer and hold across subnets in the IP network and should be easily configured using the Web browser. Support for supple- mental services is in the H.323 Standard and allows for the RTP audio stream to be redirected to different IP ports.

Configurations and Applications

The IP PBX is not defined by physical hardware limitations, as is a tradi- tional PBX or even the newer "un-PBX" systems. Traditional PBXs or un- PBXs have constraints that limit scaling the system. For example, the cir- cuit switch matrix that defines how many connections can be made at one time is based on the specific model of the PBX that has been installed. Once the limit has been reached, the entire PBX usually must be replaced.

Another limitation is the hardware line cards required for every tele- phone device or trunk interface. These cards fit into cabinets and, when the growth of the system requires more cards than cabinet space the entire system again must be replaced.

IP PBXs are very different in their architecture. Instead of a circuit switch matrix to make connections, the IP PBX uses LAN bandwidth to make voice connections. For telephone calls, the voice traffic does not pass through a central server or call manager. The call manager only per- forms signaling to set up and manage call states. Therefore, it can handle a large number of calls with fewer restrictions or limitations.

In addition, because of the scalability of LAN architectures, the IP PBX can scale linearly from one port to thousands of ports. When more ports are needed, additional hubs and switches can be added to grow the system without replacing the current investment.

IP Telephony off an Existing PBX. This configuration extends the existing PBX within the campus using the IP network as transport. The IP PBX con- nects to the PBX using either an analog or digital gateway, depending on the expectations of voice traffic and the number of users. The call manager software runs on an NT server in the data center.

This application allows a business, enterprise, university, or other large organization to extend normal telephony services using the existing IP LAN. The call manager provides feature functionality to the IP telephones, with features such as transfer, secretarial call coverage, and parallel dial plan used by the PBX. With the gateway interface to the PBX, users can call users with PBX telephones or call to the PSTN with the same privileges andrestraints set by the enterprise administrator.

Remote Offices over an IP Network. This application is simply an exten- sion of the previous configuration with the inclusion of IP WAN connectiv- ity to remote sites. The same basic rules apply for the IP PBX, just as they would for a single-site deployment. The call manager can remain on the central site, or a secondary call manager can be deployed at the remote location.

This configuration is a common initial application for the IP PBX product line. Companies with multiple sites can now easily install full telephony systems while leveraging the IP data network already in place. This saves costs for long-distance calling, as well as eliminates the cost to install a sec- ond network at each remote location. This option also enhances flexibility for growing or shrinking locations based on business conditions and mak- ing changes.

Using the analog access gateway at the remote site, the remote workers have local calling. Long-distance calling can be muted over the IP WAN linkand consolidated from the central site to maximize long-distance calling costs and administration. With the IP PBX capability to configure audio compression based on call routing, calls destined to the main location would use a lower bit rate compression to conserve bandwidth.

Network Deployment

The configuration of an IP PBX as a network-based service (such as an ISP) has characteristics similar to the previous configurations, except the call manager and the gateway are located in the WAN. On premise would be IP phones and possibly a smaller analog gateway for local calling and backup, in case the IP link to the network is unavailable.

In addition to local and long-distance calling, the network provider can also provide traditional services like voice mail and call center services with the applications residing either at the remote location or in the network.

IP provider can also provide billing and management services for the cus-tomer a range of telecommunications services in addition to long-distance routing and Internet access. The configuration options are based on the flex- ibility and power of IP networking.

PRACTICAL ADVANTAGES OF THE IP PBX

The IP PBX is expected to offer significant advantages in large-scale tele-phony. The earliest advantages pertain to cost. The benefits multiply, how- ever, and include

- Cost. Using the existing datacom network for voice transport, there is no need for the circuit-switched card or line interface card, and those expenses are avoided.
- Total cost of ownership. When one moves a phone on a circuit-switched PBX, one must call a PBX administrator, who makes an entry in a data- base that moves the phone from one physical port to another. It is logistical agony! IP phones are simpler

and less costly in every way.

- Maintenance. One can plug in an IP phone directly out of the box. It automatically configures with a call management server, and it gets a directory number. Maintenance and configuration are simpler and easier.
- Support. There is no need for external support from field technicians from a proprietary PBX manufacturer. Additionally, there is a vast hir- ing pool of people who know Windows NT, TAPI, and TCP/IP — much greater than the number of people who know a particular vendor's cir-cuit-switched PBX.
- Extensible. On a distributed campus with a unified dial plan and uni- fied feature management, one can browse into the call processing server and manage the database from any point on the network.
- Availability. It is not necessary to pay for the extra availability the PBX vendors design into the system. One can pay a lot of money for very good PBX design work. But with an IP PBX, one does not pay for the extra capacity if it is not needed.
- Capacity. Using a dual Pentium Pro 300MHz server, one can run 500 to 600 phones. With the advent of inter server signaling, it will be theoret-ically possible to scale the system up to 100,000 lines, or larger.

Payoffs

There are several ways an IP PBX will save a company money.

- Long-distance charge savings. In many international markets, espe- cially highly regulated ones, communications carriers have artificially high tariffs, as compared to carriers in deregulated markets. Addition- ally, these carriers have lower tariffs for data connections. There is

- short-term opportunity to exploit these differences — until carriers close the gap between voice and data costs. Longer-term cost savings will come from consolidation and management of all WAN connec- tions, the Internet, local calling, and long distance through a single gateway/router device.
    - Data and voice convergence. Data and voice conversion will facilitate new business practices, enabling people to work more effectively. This technology will release customers from barriers imposed by pro- prietary solutions, allowing organizations to develop.
    - Cutting acquisition and operating costs. In 1997, the capital cost of build- ing a LAN PBX system was slightly higher than the cost of building a traditional PBX. The changing marketplace has changed this model, however. The cost of swiftly evolving LAN equipment has fallen below the also-declining cost of traditional PBX equipment.
    - Administration costs. This is the single largest opportunity to reduce costs. One will manage a single network instead of two parallel net- works. Today's PBX requires a full-time staffer to manage the PBX database. In the traditional PBX, it costs $60 to $80 to move a phone. With the IP PBX, this cost is eliminated. It is also easier and cheaper to add a phone extension. General management of the IP PBX is identical to that of the IP network, which means that the same people with the same knowledge can be used in both arenas.

Fax over IP

FAXING IS ONE OF THE MOST COMMON FORMS OF COMMUNICATION IN THE BUSINESS WORLD. Faxing does what no other form of communication can it provides hardcopy information in realtime (almost) anywhere there is a telephone. Studies have shown that most people select a communications method based primarily on the urgency of the message, and faxing is still the preferred method for important documents. With 55 to 65 million fax machines worldwide generating a phone bill of about $30 billion per year, faxing represents a crucial part of the telecommunications industry.

Growth Industry

Recent studies show that faxing accounts for about 40 percent of a typ- ical Fortune 500 company's yearly phone bill, with the average fax machine shared by 21 to 23 people. In the United States, the annual spending on fax long-distance has been estimated at about $12 to $15 billion (Murata Busi- ness Systems, 1996). One industry touchstone, the annual Gallup/Pitney Bowes fax study, estimates that U.S. long-distance fax minutes grew to 140 billion in 1999. With the advent of e-mail, express mail, overnight delivery services, courier services, voice mail, and even videoconferencing, one might expect the use of faxing to decline, but it has not. Fax usage is grow- ing, not shrinking. In fact, according to a major polling organization of users who fax on a daily basis, 60 percent were faxing more than in the pre-vious year.

Preferred Form

Studies show that faxing is preferred for immediate hardcopy transmis- sion of urgent documents and documents under review, where handwrit- ten comments must be passed along. Companies also widely prefer fax for sending documents internationally. Large companies use fax extensively within their own organizations, sending almost half of their fax traffic from one company location to another. Finally, faxes get more attention and quicker response than any other medium except the live phone call.Many desktop fax software packages are available today, but using them requires a dedicated phone line and fax modem (to connect the desktop to the Public Switched Telephone Network). Studies show that only 20 percent of corporate PCs have a dedicated phone line (Gallup/Pitney Bowes, 1996).

Thus, it is apparent that bringing together the power of the desktop PC with the unique advantages of the departmental fax machine is highly desirable, but requires a somewhat different approach from the traditional modem-and-dedicated-phone line method.

Corporate Fax

According to Gallup/Pitney Bowes, about 37 to 40 percent of the typical Fortune 500 company's telecom bill results from fax. This is often a difficult expenditure to track, due to the difficulty of distinguishing voice from fax charges. In addition, corporations face many fax-related infrastructure expenses, including high-cost LAN fax servers, dedicated phone lines for fax machines (with an average of over 300 fax machines per F500 com- pany), and fax modems and phone lines for PCs equipped to do traditional PSTN-based desktop faxing.

The average fax machine in the average medium-to-large business sup- ports 23 people. It is often considered to be a less-than-efficient device, requiring that "trip down the hall" to use, sometimes a wait for access, often manual phone number entry, and then busy-signal retries. On the receiving end, there are several well-known drawbacks as well. It is no secret that faxes are not private; in fact, it is impossible for co-workers not to "peek at" faxes while sorting through jobs looking for their own (another time-consuming task). It is also not possible to retrieve faxes when travel- ing away from the office with the ease that one collects e-mail. From the standpoint of the corporate IS manager, every analog phone line that comes into the organization is a potential "open window" for intruders to hack their way into the corporate infrastructure.

Broadcast Fax

Certain groups in the organization send out faxes to groups of recipi- ents. Investor relations, human resources, PR, sales, legal, and marketing departments must often contract outside fax service bureaus at significant expense to handle broadcast faxing. Technically, a broadcast fax is any fax sent to more than one recipient, although broadcasts typically reach doz- ens, hundreds, or even thousands of destinations.

Most broadcast faxers are business organizations that need to send the same message to a large constituency, or to a smaller but very important one. Broadcast faxing is commonly used for financial results, news releases, or important internal information for the organization. Normally, one would like to have complete control over such important and urgent communications, but the time-consuming, serial nature of traditional fax- ing makes that impractical.

Even preprogramming a fax machine to send out many jobs overnight one-at-a-time falls short of the bandwidth needed to complete a medium- sized broadcast in time. Fax bureaus harness many fax devices in parallel in order to speed up the process, and provide assurances of delivery within a certain timeframe.

However, outsourcing broadcasts to fax service bureaus takes away control of important (often mission-critical) projects from in-house depart- ments. Control and feedback are necessarily reduced. And it is also expen- sive, with typical charges running to 30 cents per page or more.

Small Office Fax

Small businesses, branch offices, and consulting/vendor organizations rely on faxing in their own crucial way. Faxing may be more locally focused, but that is not always the case. Many small organizations do broadcast faxes and therefore make high expenditures on fax service bureaus. The desire to get the most out of every capital equipment dollar is very high, and install- ing dedicated phone lines and fax modems is an expensive proposition.

Small businesses often pay higher international and long-distance phone rates than large corporations, so they are even more motivated to gain control over fax charges, while lowering expenditures on equipment and service bureaus.

ADVANTAGES OF INTERNET FAX

Using the Internet to transmit fax documents from one computer or fax machine to another computer or fax machine should be as natural as send- ing e-mail, a form of

communication that is already used by millions around the world. In fact, e-mail has become an easy way to share documents that originate in electronic form, such as word processing or electronic spread- sheets. Desktop scanners are available that allow users to scan in hard- copy documents containing written material such as signatures, comments, and corrections. Scanned documents in graphical format can be sent from desktop to desktop as e-mail attachments.

But what if hardcopy, in the form of fax output, is desired or required at the receiving end? Desktop fax applications that require a fax modem and dedicated phone line have drawbacks the difficulty and expense of provid- ing the phone line, the need to disconnect from other communications in order to use the fax phone line, etc.

Internet fax connects the PC (or a standard fax machine) to a fax server in the Internet. The fax server routes jobs through other servers in the Internet to a final server located physically near the destination. That server then makes a phone line connection to the destination fax machine and delivers the job at the lowest possible cost. Alternatively, the deliver- ing fax server may send the job to a fax application in a desktop system, or as a graphic file e-mail attachment to a destination e-mail account.

There are three basic fax "on ramps" to the Internet

1. LAN-connected computer to Internet
2. fax modem dial-up connection to Internet via PSTN
3. fax machine to Internet via PSTN

At the receiving end, there are also three types of "off ramps"

1. Internet to LAN-connected computer
2. Internet via PSTN to computer with modem
3. Internet via PSTN to fax machine

The key to Internet faxing is the network of specialized fax servers inside the Internet. These servers accept fax jobs from PCs or specially equipped fax machines, then either deliver the jobs directly to the destination device, or route the jobs over the Internet (using least-cost routing) to a server near the destination for delivery.

Internet fax servers are connected both to the IP network and to the ana- log PSTN simultaneously. Each server can accept incoming fax jobs from the Internet or over analog phone lines. Likewise, they can send fax jobs over the Internet through the IP connection or deliver them via PSTN to destination fax devices.

When fax jobs are handled by computer, rather than telephony equip- ment only, many benefits become available. Users can view the status of all faxes, send faxes to large recipient lists, print faxes out on high-quality printers, forward faxes electronically, and distribute documents to a mix of fax devices and e-mail systems. By receiving faxes through an Internet server, users can ensure that they capture all their faxes, and can access them electronically at all times, even while away from the office.

Individual User Advantage

Individual users gain by faxing over the Internet. By adding fax send and receive capability to the PC, users come one step closer to a universal mes- saging center (adding voice to the PC will one day complete the solution). Low-cost desktop scanners that take up very little desktop "real estate" (some are built directly into a keyboard) can give the PC full "fax machine" capability. Because the average worker sends up to 25 fax pages

per week, personal productivity goes up by eliminating trips to the stand-alone fax machine. Receiving faxes via an electronic inbox gives another boost to productivity and security, as the user previews faxes privately online and prints hardcopy from a high-quality laser printer.

Power User Advantage

Groups such as investor relations, PR, sales, and marketing send out faxes to large lists of recipients (broadcast faxing). Internet-based faxing lets users manage the entire process right from an in-house PC, including setting up recipient lists, sending the broadcast, and monitoring status of each broadcast fax job. And because the Internet "does the work" of trans- mitting the fax document to each recipient, the in-house PC is freed up immediately for other work, as soon as one copy of the document, and the recipient list, is sent to the Internet.PC

Organization Advantage

Organizations view Internet fax as an opportunity to outsource a signif- icant piece of their communications infrastructure. Many IS and telecom managers welcome the chance to shift capital equipment and technology investments to the hands of specialized service providers. Faxing is a par- ticularly good candidate for outsourcing. The combination of empowering individuals in the organization to do more while at the same time increas- ing control over usage and saving money on capital equipment and telco charges is a powerful incentive.

Many fax machines connected to dedicated analog lines actually bypass the corporate PBX, and may in fact cost the organization more than the offi- cial corporate rate. So, Internet faxing can produce dramatic savings while capitalizing on the investment made to put all those PCs on all those desks.

From a security standpoint, fewer analog phone lines into the organiza- tion means fewer "open windows" for intruders to hack their way into the corporate infrastructure. Getting visibility into the usage of 40 percent of the organization's phone bill is also a good thing. The key ingredients that corporate managers look for are reliability and quality. Here, the NSP is uniquely positioned to leverage its track record as an infrastructure ser- vice provider.

Smaller organizations can win big as well. They often pay higher interna- tional and long-distance phone rates than large corporations. Savings gained through lower rates, lower equipment costs, and elimination of ser- vice bureau charges hit the bottom line very quickly.

Top Ten Advantages

1. The PC is easier to use, and more convenient than a fax machine.
2. It is faster no trip to the fax machine, redials, etc.
3. It is less disruptive to send faxes directly from within applications such as word processors, or from paper using a handy desktop scanner.
4. Inbox users never miss a fax, retrieving their faxes like e-mail or as e-mail attachments. Inbox faxes are private and secure, and easily previewed on the PC.
5. No dedicated phone line, fax modem, or LAN fax server is required.
6. Performing broadcast faxing easily from the desktop eliminates ser- vice bureau

charges.

7. It automatically captures and reports online status of faxes.
8. Saves money on transmission costs, especially internationally.
9. Faxes can be sent to fax machines or converted to e-mail for delivery to PCs.
10. NSP-based service produces itemized billing for easier usage tracking.

BUILDING AN INTERNET FAX SERVICE

Take a step-by-step look at how an Internet fax service offering can be created. Why offer Internet fax? Every network service provider (NSP) is asking itself where to put its future emphasis. There are many competing demands for resources and priorities. But a few facts stand out

1. The most valuable resource the NSP has is its customer base.
2. Creating new revenue streams from that existing customer base is not easy. Today, connectivity is a commodity. Where is a new, prof- itable offering going to come from?
3. Establishing an infrastructure for metered, "pay-as-you-go" services is essential to break out of the current "flat monthly fee" pricing environment.

Every major service industry, from telecommunications to overnight deliv- eries to airlines, relies on sophisticated rate structures and pricing pack- ages to optimize charges against usage. Usage-based pricing funds improvements in the infrastructure and enables the creation of additional offerings. Network service providers need the ability to authorize and authenticate users, allocate resources, meter transactions, and charge for measured usage of their service offerings.

Internet fax is a great "first value-added-service" opportunity. Fax trans- mission lends itself to the current technology of the Internet. Internet fax- ing is an "easy reach" from current desktop capabilities, and, as pointed out, fits right in with the current suite of communications activities. Faxing is an extremely large business, international in scope, that offers great rev- enue potential. Internet-based faxing is the first of many billable value- added services that can be offered by an NSP. Future offerings could include streaming video, videoconferencing, voice applications, data access services, Internet-based metered gaming, information "push" ser- vices, and more.

KEY REQUIREMENTS

Here are the key requirements for a successful Internet fax service offer- ing Offer a compelling value proposition to the end user. Customers per- ceive value from improvements in their daily lives and business opera- tions. Many of these benefits were listed in the previous section, including personal productivity, privacy, and security, easier and more manageable IS/telecom, better usage of computer resources, and greater control over important/urgent hardcopy communications.

Savings on capital equipment costs are important, reducing dependency on LAN fax servers, dedicated analog phone lines, fax modems, etc., and provide savings over telco fax transmission rates. By "letting the Internet do the work" of long-haul fax deliveries, most users can experience a price advantage from Internet fax versus a traditional telco service. Rather than positioning Internet fax as a premium service, it can be promoted as way to reduce the overall cost of fax transmissions.

Provide Both Outbound (Sending) and Inbound (Receiving) Services

To completely outclass the traditional fax environment, NSPs should deploy fax services that beat the fax machine "coming and going." Fax receiving services offer personal fax numbers, privacy, security, 24-hour-a- day coverage, and unmatched accessibility. The power of a complete send- and-receive solution is both dramatic and easily marketed.

DEFINING THE IP FAX SERVICE

Every service business begins with the customer. The implementation of Internet fax will depend on how one estimates two specific customer characteristics

1. What will the adoption rate be over time for the service? The rate of acquisition of new customers will depend on many factors, includ- ing how many business vs. consumer users are in the customer base, how aggressively one plans to promote the service, and what kind of introductory offers one plans to implement, if any.
2. What is the overall faxing profile of the customer base? According to Gallup/Pitney Bowes, corporate faxing patterns are roughly 10 to 12 percent international, 50 to 60 percent long distance, 25 to 40 per- cent local/regional.

The average fax user sends five fax pages per business day. These figures are a good starting point for estimating traffic loads. Over time, one will develop an accurate model of the particular customer base.

Look at three different adoption-curve scenarios based on a 20-day "business month." For now, exclude broadcast users and assume that everyone sends an average of five fax pages per business day. Note that the following rates of adoption are estimates. They are presented here for pur- poses of comparison only, and are not based on data gathered through experience.

Deploying Internet Fax Servers

At first, it makes sense to do a limited rollout of servers to high-traffic POPs, because servers deployed in high-traffic localities provide immedi- ately profitable local fax delivery. Meanwhile, a centralized pool of servers delivers all other fax traffic via long-distance lines. However, as in the above service analysis, inbox accounts can increase revenue when config- ured into the network of deployed servers. So, a sensible strategy is to deploy local fax servers to provide inbox accounts and local fax deliveries, as well as a pool of "long-distance" servers. (The deployed servers provide local area code numbers for inbox customers.)

Over time, additional servers can be deployed to expand the local cov- erage of the system and to fine-tune it for greatest efficiency and profitabil- ity. This "division of labor" between distributed "local only" servers and pooled "long distance only" servers makes the provisioning of each server straightforward. In this approach, each server is set up with either local T1 and ISDN ports or with long-distance T1 ports.

A sophisticated least-cost-routing algorithm behind the scenes ensures that servers are delivering packetized fax jobs in the most reliable, time- efficient manner, adapting flexibly to changes in network traffic loads and other operating variables. The more fax servers an NSP deploys, the greater geographical coverage it can achieve. By deploying a fax server to every POP in the NSP network, the vast majority of local user faxes will be delivered locally, and more long-distance faxes will also be delivered for

the cost of a local call. The ultimate goal of the service is to establish an essentially "distance-independent" alternative to the "distance-sensitive" model offered by traditional telco networks.

BASIC INTERNET PROTOCOL CONVENTIONS

Protocols that together manage packets on the Internet build on several widely accepted conventions/foundations. The principal components of the network are data connecting equipment (DCE), such as routers, and data terminating equipment (DTE), such as desktop computers (also called "hosts").

Without any special adjustments for the unique requirements of different media types, the network layers work in concert to transmit packets of user information. In its simplest implementation, the flow of the data from and to end points over an IP network is monitored or verified by a simple layer 3 communications protocol (e.g., Transmission Control Protocol, or TCP).

The advantage is that there is very little communications overhead associated with components in IP networks "talking" to one another. TCP running over IP in effect takes care of this. Each packet allocates the maxi- mum number of bits to the user's information.

Together, protocols ratified by the Institute of Electrical and Electronics Engineers (IEEE) and the Internet Engineering Task Force (IETF) ensure that packets of data are reliably transmitted under any condition, as quickly as possible (e.g., bandwidth or network load on the segments of the networks tying together two or more points).

Several protocols have been developed to manage the unique require- ments of real-time streaming data. To understand the importance of these developments, this chapter begins with a high-level discussion of video- conferencing and visual collaboration using networked multimedia desk- top computers.

PRINCIPLES OF VIDEOCONFERENCING AND MULTIMEDIA COMMUNICATIONS

When a video camera and microphone pick up real-life events, the imag- ery and sound can be turned into digital formats for communications between Properly enabled end points over local or wide area networks. For the user to perceive the moving images and intelligible sounds, the digital information moves from transmitter to receiver (transmitters are simulta- neously receiving in the case of two-way videoconferencing) in a highly consistent fashion. Compressed in real time, the data streams over a net- work in such a way that frames of video can be reconstructed and synchro-nized with audio with the least end-to-end delay.

Quantitative Quality of Service Parameters

End-to-end quality of service (QoS) in videoconferencing and visual col- laboration is defined as the level of satisfaction a user has with a given ses- sion. It is a function of many independent and interdependent factors (e.g., window size, processor speed, network bandwidth), which together influ- ence frame rate, bit depth, image clarity and resolution, audio clarity, lip synchronization, and latency.In contrast to

conventional data applications in which data transmis- sion is bursty, digital video and audio applications require continuous data transmission. In IP environments, precise bit rates during the transmission vary. In the following list, the factors influencing bit rates (bandwidth) dur- ing videoconferencing and collaborative computing are presented as a function of the media (i.e., video, audio, and data)

Video factors include

- Bit depth (number of colors)
- Resolution (size of the image being captured, compressed, transmit- ted, and decompressed) — Resolution is contrasted with window size, which is the size of the image that is displayed for viewing. If the win- dow size is different from resolution, then interpolation is used to gen- erate a new image that fits the window.
- Q factor (sharpness of edges in any given frame)
- Smoothing (this is a result of — and dependent on — motion estima- tion algorithms and content changes from one frame to the next)
- Frame rate (frames per second) — For example, NTSC video, the U.S. standard, is 30 fps; PAL, the European standard, is 25 fps.

- Audio factors include

  - Sampling rate (the number of audio samples captured, compressed, and transmitted, expressed in KHz cycles per second) — For example, telephony is 6.3 KHz, FM radio is 36 KHz, music CDs 44.1 KHz (See Exhibit 63-1).
  - Bit rate (the number of bits the system has in order to accurately rep- resent different tones; for example, 8-bit or 16-bit)
  - Mono and stereo sound

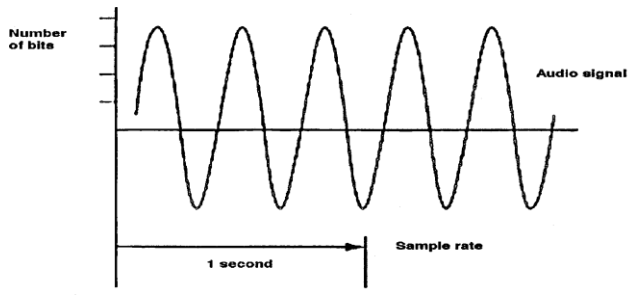

Note:
Frequency response is how high and low the audio signal can go. Clipping occurs if the actual audio signal becomes higher or lower than the signal the analog-to-digital or other chips can handle. Sample rate refers to how many audio samples are taken in a 1-second interval. Numbers of bits (usually 8 or 16) represents the accuracy of each audio sample.

Interdependence of quality of service parameters in time- dependent streaming media.

- Data factors that includeds

- Quantity of data
- Frequency of transmissions
- Latency of transmission (how long it takes to send at a given bit rate)

The central processor, any additional compression/decompression (codec)

circuits, network infrastructure, and the user's network connec- tion directly affect these quantitative factors.

Qualitative QoS Experience

Depending on the quantitative parameters, the user's experience with multimedia (which is a combination of audio, video, and data) may be qual- itatively different. Exhibit 63-2 is intended to help readers understand the interdependence of various quality-of-service parameters in time-depen- dent streaming media.

The user's experience with bidirectional live video and visual collabora- tion over any network can be expressed quantitatively and qualitatively; in general, though, the objective is to reproduce a live meeting or conversation.For the interaction to be as close to natural as possible, it is especially impor- tant that both (or all) users in a videoconference experience a uniformly low latency (minimum delay). Any variation in the frame rate between points is perceived as jitter. Poor synchronization between lips and audio is also dis- tracting.

Thus, the most important factors in the user's qualitative experience of videoconferencing system are

- synchronization of audio, video, and data
- end-to-end latency
- window size
- jitter
- richness and clarity of audio
- image clarity (bit depth, sharpness, smoothing, and resolution)

All these quality-of-service concepts are critical to the reader's overall understanding of the pros and cons of selecting IP networks for video- conferencing.

PROS AND CONS OF VIDEOCONFERENCING OVER IP NETWORKS

Founders of the Internet were academics driven by four guiding principles

- reliability (guaranteed delivery of packets), not efficiency
- end systems' interoperability, and information or packet loss recovery
- variable quality of service, not guaranteed bandwidth, so that any network bandwidth can be accommodated by a single protocol
- no support for charging mechanisms, since commercial traffic was not envisioned

Internet protocols have withstood the test of multiple applications at the user interface, new operating systems in end points, and ever-changing transport media in the physical layers (e.g., SONET, ATM). For videoconfer- encing and visual collaboration, there are fundamental principles that determine bandwidth use in any session.

Bandwidth Allocation

Exhibit 63-3 illustrates the following types of bandwidth allocation

- Total bandwidth. This is the maximum bandwidth available for user data and network management overhead.
- User-available bandwidth. This is the total bandwidth minus the net- work

management overhead.

- Application-specific bandwidth. This is the bandwidth requested by the application.

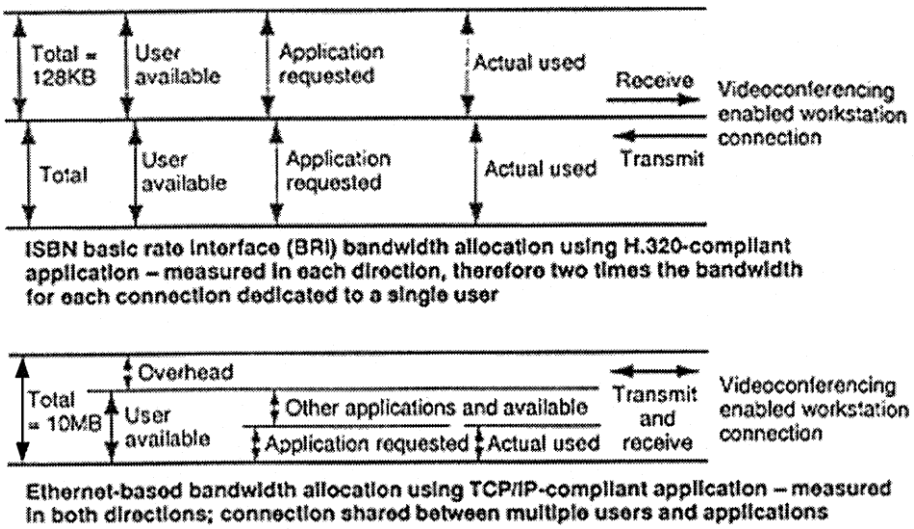

Examples of bandwidth allocation.

- Actual used bandwidth. This is the minimum bandwidth of either what the application requested or what is available over the end-to-end net- work during a particular session (whichever is lower).
- Allocated bandwidth. This is the portion of the total user-available bandwidth that can be reserved; it only applies when a reservation mechanism is in place, requested, and granted.

IP Network Advantages

Compared with other wide area network communications protocols, IP has four principal advantages for videoconferencing

- Low management architecture overhead and high carrying capacity make IP cost effective.
- No guaranteed quality of service and low overhead make IP band- width scalable (64K bps to 100M bps and beyond).
- IP architecture provides for many simultaneous virtual circuits, thereby enabling multiple services and multiple connections (through well-recognized sockets) at the same time.
- Packet structure and network design make it possible for both broadcasting and multicasting to occur in the same network, with- out requiring packet redundancy.

Using extensions of IP to modify the payload format and reduce over- head associated with packet acknowledgment between end points, the communications between end points dedicated to user information is high. In some cases as much as 95 percent of the bandwidth can be assigned to user traffic. The actual bandwidth a data stream uses is a function of the bit rate the end points have agreed to send upon call establishment.

When two endpoints have only very low bit rate capacity, owing to the local bus architecture or modem technologies over a POTS (plain old tele- phone service) line, the IP-based videoconference session operates within this constraint. On the other hand, when the same software (e.g., user application) is used over a high-bandwidth connection, a videoconference can take advantage of the increased capacity without any change in the application or communications protocol stack.

Networks without support for Internet Protocol Independent Multicast- ing (PIM) must recreate a data stream for each of the desktops to which the user wishes to communicate simultaneously. This puts extra pressure on the source (i.e., the host has to create, manage, and transmit the same packets more than once) and on the shared resources connecting all users, whether they are multimedia communications enabled or using traditional transactional functions between clients and servers. Multicasting on IP enables multiple participants (so-called multipoints) to experience the same real-time conference at the same time. The importance of this feature cannot be underestimated.

IP Network Disadvantages

Many practitioners believe that connectionless communications and best-effort delivery of packets is inappropriate for isochronous network traffic (e.g., where data needs to be sent sequentially and arrive at its des- tination at specific intervals). Such as that generated by digital video and audio. The principal disadvantages of IP network for videoconferencing are basically

- Lack of guaranteed bandwidth, unless all the network components are controlled by a central manager that has chosen to implement net- work-wide bandwidth reservation schemes.
- Lack of international standard that would support interoperability of products from different vendors for intra- and intercompany communications.

The most common wide area network solution for videoconferencing has been integrated services digital network (ISDN). With a dedicated con- nection (i.e., a communications circuit) between two endpoints, band- width is guaranteed and, consequently, the quality of service remains con- sistent throughout a session.

The next most common wide area network for point-to-point desktop videoconferencing has been POTS, which despite its low bandwidth, ensures, like ISDN, a dedicated circuit for a conference between two end points.

Users find that they prefer running business applications over networks in which the bandwidth allocated, requested, and used are all the same (i.e., the circuit-switched environment). And there are many competitive, standards-compliant offerings for desktop videoconferencing over basic rate ISDN.

The natural consequence of a large — and especially a standards-com- pliant — installed base is that there are more people with which a system can interoperate without modification of any software or hardware. Until more vendors deliver standards-compliant, interoperable solutions for vid- eoconferencing, the installed base will be confined to pockets of propri- etary products that interoperate among themselves, but do not permit endpoint application independence.

Specifications for session setup, management, and compression were under development (e.g., IETF Audio/Visual Transport Working Group and the International

Telecommunications Union's [IUT] H.323), and became commercial products during 1997. However, the standards have been revised three times since then, which makes it difficult for vendors to fully comply with the changing standard.

Compared to other protocols on the local area network for managing desktop videoconferencing (e.g., IsoEthernet or the specialized multime- dia operating system, such as MOS, by First Virtual Corp., for 25M-bps ATM at the desktop), IP communications over 10Base-T Ethernets are prone to suffering from network contention and congestion. The relatively low bandwidth available for each user on a 10Base-T network has, to date, been unsuitable for business-quality videoconferencing using inexpensive soft- ware codecs.

GETTING STARTED WITH IP-BASED VIDEOCONFERENCING

Disadvantages and drawbacks aside, it is clear that many organizations will choose IP-based videoconferencing for its

- low cost of entry (e.g., integration into existing infrastructure)
- low cost of ownership (e.g., low maintenance and no telecommunica-tions charges)
- ease of use (e.g., accessibility of the global IP network compared toISDN provisioning)

For many years, real-time packetized audio and video over IP networks were tested and used on an isolated portion of the Internet the multicast backbone (Mbone). The Mbone is operating, expanding, and improving. Some of the protocols used on the Mbone (developed by the Audio/Visual

Working Group of IETF) are being ratified by the IETF and have migrated from this relatively exclusive academic and industrial environment into commercial routers for Internet and intranet deployment. Subsequently, IETF protocols for managing video and audio packets would be widely incorporated in enterprises and on the Internet in general.

This section examines the components users need to add to their desk- tops for videoconferencing on IP. Local (LAN), metropolitan (MAN), virtual (VAN), wide area (WAN), and global network managers need to modify and prepare their networks to support the types of traffic generated by video- enabled desktops. The scope of these networking component changes and alternatives is discussed in more detail later in the chapter.

Desktop-enabling Technologies

To experience desktop videoconferencing on the Internet (or intranet) firsthand, the user needs only a camera for video input, a microphone for audio input, speakers (presuming the user wants to hear what others say), software to give the user access to connection initiation and answering, session management, compression algorithms, a video display board, an IP network interface, and a premium CPU.

CU-SeeMe and Other Software Supporting Multicasting. One of the more unique and proprietary of these desktop components is the user application and interface software. The first, and consequently one of the most widely deployed, application designed for videoconferencing on the Internet CU- SeeMe, originated at Cornell University. Distributed as freeware/shareware for the first several years, the application

satisfied the needs of many Mac- intosh users in academic and nonprofit institutions for distance learning and research applications.

In 1995, Cornell University issued an exclusive license for commercial distribution of CU-SeeMe to White Pine Software. Since then, White Pine Software has ported the application to other platforms and greatly enhanced the functionality (e.g., adding color video, password security, and whiteboard capabilities).

Cu-SeeMe, like three or more competing user applications currently offered on the Internet — for example, VDOnet's VDOphone, CineCom's CineVideo/Direct, Apple Computer's QuickTime Conferencing, and Intelli- gence at Large's Being There — provides a directory management system, call initiation, answering and management software, and some utilities for controlling video and audio quality during a session.

The Desktop's Connected to the Mbone. Percept Software had developed multicast audio/video server and viewer products for Windows 3.1.1, Win- dows 95, and Windows NT to help enable the PC/Windows world to join the Mbone community. The viewer, called FlashWare Client, can receive Mbone sessions transmitted with Livermore Berkeley Laboratory's vat 4.0 in real- time transport protocol (RTP) mode (selected via the –r option) using PCM, DVI, or GSM audio-encoding algorithms and vic 2.7 using its default
H.261 video codec.

On the Precept Web site is a program guide that lists Mbone sessions using these protocols; users can launch the client automatically from there. The client is built as a media control interface (MCI) device driver so it can be invoked through Microsoft's Media Player, a Netscape plug-in, or other applications using the MCI A-PI. Playback of audio and video is syn- chronized using the time-stamping mechanisms in RTP and real-time trans-port control protocol (RTCP).

The IETF Audio/Visual Transport Working Group's RTP and RTCP proto- cols have been developed to facilitate the smooth transmission, arrival, and display of streaming data types. When end point applications support RTP, packets leave the sender's desktop with a time stamp and content identification label. Using this information, and through congestion-moni- toring facilities at either end, the proper sequences of frames can be more reliably recreated at the receiving station, using a specified delay buffer (generally less than 100 milliseconds). Netscape's Cooltalk is another example of an architecture for streaming video and audio with RTP-ready endpoints.

Compressing Audio and Video Streams. In all but the most exceptional conditions (e.g., broadcast-quality production requirements), digital video and audio need to be compressed for superior management. Subsequently, the information must be decompressed (decoded) upon arrival so that it can be displayed on its destination screen.

A comprehensive discussion of compression technology and the ensu- ing debates over the virtues of different algorithms are not within the scope of this chapter; however, it must be noted that digital video com- pression has a marked impact on the quality of the experience users can expect when videoconferencing over an IP network.

All freeware applications for IP-based videoconferencing bundle a soft- ware codec for encoding and decoding the audio and video streams at the appropriate

bandwidth for the station. Software codecs deliver lower qual- ity audio and video than hardware in which there are optimized digital sig- nal processors (DSPs) for these functions. Currently, there are no standard compression algorithms for use on the IP-based networks, so users receive the codec specified by the desktop application.

In the case of freeware developed by Livermore Berkeley Laboratory, as well as Apple's QuickTime Conferencing, the architecture can accommodate Any number of compression algorithms, including H.261, which is the basis of all H.320 systems. The products comply with a new specification — H.323
— for videoconferencing over IP networks and use H.261 as a codec; how- ever, a new and more efficient version (H.263) is less bandwidth consump- tive and will quickly replace H.261 on IP networks.

Hardware Implementations for Business-Quality Videoconferencing. In gen- eral, videoconferencing on IP networks is like any other commodity the customer gets what he or she pays for. The software packages mentioned so far are considered suitable for academic, nonprofit, and perhaps "per- sonal" applications. For now, customers who seek "business-quality" video and audio will need to evaluate and select desktop videoconferencing sys- tems that have been implemented in hardware.

Video and audio compression hardware for IP-based conferencing is available for Industry Standard Architecture (ISA) as well as Peripheral Component Interconnect (PCI), Sbus interfaces, and all major operating systems. Examples include those offered in Intellect Visual Communica- tions Corp.'s TeamVision family of products (using very large scale integra- tion [VLSI] chips and Mosaic's own design) and Netscape's Communique Line of products (using an Osprey Technology board and Lucent Technol- ogies' AVP chips for DSP-assisted compression and decompression of audio and video).

Network Interface Hardware. All vidioconferencing systems require a network interface adapter for LAN or WAN access. Most institutions with intranets or T1 access to the Internet provide an Ethernet adapter at each desktop. The bandwidth and suitability for video depend on whether this interface adapter is 10Base-T, IsoEthernet, or 100Base-T and an assortment of network design issues. IP-based videoconferencing can also run locally over Token Ring networks with routers providing connectivity to and from the wide area networks.

For those who do not have a dedicated connection to an IP network, dial-up access to the Internet is accomplished with a point-to-point proto- col (PPP) or serial line IP (SLIP) connection via a modem or an ISDN termi- nal adapter through an Internet services provider (ISP). In general, dial-up IP network interfaces accommodate consumer applications adequately, but are not suitable for business-quality video and audio supported with specialized hardware.

VIDEO-READY NETWORKS

In the previous connectivity scenarios, an Internet communications protocol stack in the host operating system negotiates and monitors con- nections. This section, however, focuses on the steps needed to address bandwidth, as well as the IP facilities and the internetworking software commonly used and currently being proposed for desktop videoconfer- encing in IP environments.

Network Upgrade and Management Issues

Preparing a network for any new application, including multimedia, requires careful analysis of existing components and user requirements. As far as network upgrades for videoconferencing are concerned, an intranet is quite different from the public Internet.

In the private network (e.g., intranets over LANs, MANs, VANs, and WANs), technologies can be more consistently deployed, more effectively maintained by a central IT group, and are often more economical to pur- chase when large site licenses are negotiated. This said, jurisdictional (i.e., workgroup) management of LANs is increasingly popular.

In contrast to the situation with LANs and private WANS, new protocols and architectures take much longer to deploy in the public/commercial IP environment. There is an inherent lack of control in this progress, espe- cially if new management challenges are associated with upgrades. Video- enabling upgrades clearly fall in this category.

One way for LAN administrators and managers to approach the design of a video-ready network is by working from the endpoints toward the com- mon infrastructure (e.g., the Internet).

Endpoint Performance. Initially, users and planners should evaluate the endpoint CPU performance. If the CPU is involved in any general data appli- cation management and compression or decompression tasks (which is almost always the case in the desktop videoconferencing applications dis- tributed as freeware, and less the case when add-on compression hard- ware is necessary), then low performance at the desktop will translate to poor quality of service and less efficient bandwidth usage patterns. When endpoints are enhanced and capable of compressing video frames, band- width will be more efficiently utilized between desktops.

Network bandwidth requirements for desktop videoconferencing vary with applications as well. Some applications — especially precision medi- cal or surgical applications or high-quality entertainment and advertise- ment production — require many megabits per second to transmit lossless (i.e., compressed without any loss of information) or nearly lossless video between points.

In most business scenarios, however, the combination of efficient com- pression algorithms, network management software, and user tolerance of less than TV quality video keeps the bandwidth requirement (per bidirec- tional session) between 28.8K and 768K bps.

Because, in general, users' lowest level of tolerance is the highest per- formance they have had the privilege of using, it is safe to assume that IP networks in place today need modification to deliver acceptable business- quality video in real time. On a shared network, such as an Ethernet, Token Ring, or fiber distributed data interface (FDDI) network, or the commercial Internet today, all stations have equal opportunity to send and receivedata. This is known as "time division multiplexing."

Several options exist for changing network designs to accommodate the demands of streaming data types. One of these is to supplant or augment bewst-effort protocols in order to prioritize video and audio streams in such a way that end points receive consistently low latency. This approach is discussed in greater depth a little later in this

chapter.

If, prior to changing the data management, an enterprise decides to deploy high-speed LAN technologies (e.g., 100Base-T, 100VG-AnyLAN, ATM), there also needs to be upgrades to WAN infrastructures. Options and issues in this arena are the focus of many books and current articles and outside the scope of this chapter.

Evolving Bandwidth Management Protocols. From an integration per- spective, however, one of the most important advantages of planning the netwok using IP is that IP is well adapted to LAN as WAN environments.

DVMP. At the netwok achitecture level, piorritization schemes in the IP specifications issued by IETF working groups hold the greatest promise for improved management and distribution of video over IP networks. In the late 1980s, the IETF ratified the distance vector multicast routing protocol (DVMRP) to transport live video feeds in IP multicast mode over the Mbone. DVMRP works by essentially "flooding" all available routes with a broadcast message, something which could be tolerated more easily before the Intenet grew in popularity.

Ipv5 or Streaming II Protocol. About the same time DVMRP was intro- duced, the Steaming II (ST-II) protocol was proposed by Bolt Beranek and Newman (BBN). A connection-oriented routing protocol, ST11, is used in endpoint and router software and offers a call setup facility that lets the originator control bandwidth in a video and audio session by allocating bandwidth through the router upon request. Virtual links are established for the duration of the session and resources are allocated along the virtual links.

ST-II, also known as Ipv5 and sometimes called ST-II+, is evolving to address connection setup delays and options for allowing both receivers and senders to open sockets without a conference administrator's approval. Today's version of ST-II+ is not backward-compatible with ST-II.

Protocol Independent Multicasting (PIM). To address the inherently "unscalable nature of DVMRP," the PIM system was proposed. This proto- col designates so-called rendezvous points for registration of both senders and receivers of multimedia multicasts.

Because the protocol (implemented in routers such as those shipping from Cisco Systems) is not restrictive, it also works with any unicast rout- ing protocol (as in the case in a private videoconference over an IP WAN). Dense mode PIM, which applies where the volume of multicast traffic is high and senders and receivers are in close geographic proximity to one another, uses reverse path forwarding and operates much like DVMRP.

RSVP. The bandwidth management protocol with the most enthusiastic following to date is known as the reservation protocol (RSVP). Imple- mented in endpoint and router software on the Mbone and currently under review for IETF ratification, RSVP guarantees bandwidth allocation in con- nectionless networks according to a receiver-driven model.

RSVP is fixed-bit-rate allocation, with routines to handle available bit rate in the future. It is also technology independent and can run on ISDN and private network connections such as Ethernet-based intranets.

Prototype support for RSVP has been demonstrated by several different router vendors and became available in many products in 1997. With these products, RSVP was quickly deployed throughout intranets.

Billing and Related Issues. In addition to the inpediments cited so far — namely, complex management challenges associated with video — current Internet pricing models do not reflect guaranteed bandwidth allocation. As a result, most commercial Internet service providers will be reluctant to implement RSVP in their routers because, in using this protocol, a few users could potentially monopolize router resources without appropri- ately compensating the service provider. Research at BBN and in the IETF's Internet Services Working Group has addressed the problem with specific billing protocols built into endpoints and routers.

Researchers at the University of Illinois-Champaign are exploring a solu- tion to circumvent the successive layers of management code over IP. The video datagram protocol (VDP) eliminates TCP and works at the IP level to move video, audio, and data simultaneously. The protocol itself addresses the delivery timing issues by dynamically using a best-effort adaptive flow control methodology.

HOW DESKTOP-TO-DESKTOP VIDEO AND AUDIO CONNECTIONS OPERATE IN IP NETWORKS

Given the large number of freeware and shareware solutions for video- conferencing on IP networks and proposed standards, and the rapidity of new developments on the IP landscape, it is difficult to make generaliza- tions about the manner in which desktop-to-desktop connections are nego- tiated, maintained, and torn down in all user applications. This section explores some of the approaches different products use to enable video- conferencing and collaboration over IP networks.

Negotiating/Establishing Desktop Connections

Some applications have relied primarily on Internet (and later intranet) servers for negotiating desktop connections over LANs and WANs. White Pine Software's Reflector or software, for example, supports unicast, broadcast, and multicast sessions.

Unicasting and multicasting are achieved by specifying a reflector as the destination and sharing (publishing) the appropriate IP address of the reflector in question with other conference participants. To control unwanted participants, the reflector lets network managers issue pass- words for different conferences. In addition, a roster of conference partici- pants (one for a unicast, more for multicasts) is published dynamically to all participating desktops.

Several freeware and shareware programs are available to initiate and administer online conference over IP networks. Confman is one such tool for conference initiation and monitoring that employs certain Mbone tools vat (for audio data), vic and nv (for video transmission), and wb (for white- boarding) on the Internet.

Confman does not handle multimedia data, but helps the user to plan, setup, and control a conference by letting the network manager choose meeting participants (by IP address), the start time, and the tools (and codecs) the session members need to run on their endpoints. Conferences can be held in two different modes

1. Closed mode. Using this mode with more than two participants requires a server

process to route the multimedia data. This process might be regarded as a conference room. All connections are unicast connections; no multicast features are required.

2. Multicast mode. Multimedia data is sent via multicast. To restrict access, data has to be encrypted. A conference management tool distributes the encryption key to the selected participants in advance.

Microsoft's and Intel's Internet telephony products (and subsequent IP- based videoconferencing offerings) use the standard User Locator Service (ULS) to negotiate/establish calls between desktop videoconferencing users on IP networks. In June 1996, ULS was submitted to the IETF for con- sideration to become a standard and incorporated with LDAP (lightweight directory addressing protocol).

Ideally, all intranet and Internet service providers will have standards- compliant directories. For users not on a corporate network, the ISP's direc- tory will automatically associate the e-mail address with a person. Every time users connect to the Internet, the network service can then pick up the IP address and initiate a local call to the end point nearest the recipient.

In a corporate environment, where there is a Novell network with ULS and computer telephony integration, simply having a desktop computer turned on and connected to the network should suffice to identify the end point for any incoming video and audio (or audio-only) sessions.

Maintaining/Modulating a Videoconference Session

As long as the sockets between endpoints, or between the endpoint and server, are open, the IP session is maintained. Another way of expressing this is that within a conference, the virtual circuits are always present.

Some applications use a single circuit for all audio, video, and data; other applications use a separate circuit for each media type. Having a sep- arate socket for each permits higher error recovery and, therefore, fewer chances for problems to occur during a live conference.

Multicasting

The focus of this chapter is bidirectional real-time video and audio between two or a few desktop systems, otherwise known as desktop vid- eoconferencing or, in Internet terms, strictly a unicast transaction.

For most who participate in unicast or simple point-to-point sessions, there comes a time when applications, especially meetings, require more than two participants. To execute on a network limited to unicast sessions, applications must generate a unique copy of each packet and send those packets to each participant's desktop (by specifying the endpoints' IP address). This is inherently inefficient if there is an alternative. With Inter- net multicast protocols, an application generates a single copy of each packet and sends it to a group address. Endpoints (e.g., clients on the cli- ent/server network) can selectively choose to listen to the multicast address. Multicasting minimizes network traffic and gives all users on a network grreater flexibility.

For an application to take advantage of multicasting, the IP stack in the network software on the host must support multicast and broadcast addressing. Multicast is implemented at both the data link layer (layer 2) and the network layer (layer 3). For multicast confined to a single LAN, the data link layer implementation is sufficient.

When a multicast application extends into different network media, such as frame relay, FDDI, or Ethernet, network layer implementations are recommended. Therefore, multipoint applications with both LAN and WAN participants must implement in both layers.

For all vendors of endpoint applications and network components to interoperate in multicast IP networks, several parameters must be defined. The IETF has standards specifying the addressing (i.e., mechanism for clients to inform the network that they wish to be a member of a specific multicast group), and multicast routing (e.g., DVMRP and PIM).

Monitoring and Managing a Videoconference Session

In general, reflectors, as the servers are sometimes called, provide net- work bandwidth control, video "pruning," audio prioritization, and a range of conference control software. Using network management utilities, the reflector/server can adjust transmission rates of specific individual users on-the-fly, if packet loss is running too high because of heavy network traffic.

If contention is too heavy for a reliable conference, the transmission remains at the lowest setting and only moves up when the network is less congested. This is an important tool for network managers and ISPs who have to be concerned with balancing the needs of their other nonvideocon-ferencing network users during peak load.

Network monitoring utilities allow network managers to control the maximum bandwidth allocated per videoconference and the maximum number of simultaneous videoconferences. In this way, sufficient band- width is reserved for other users who have conventional network applications.

One of the freeware network monitoring tools for videoconferencing is Rtpmon.Rtpmon is a tool for viewing RTCP feedback packets from a ses- sion using the real-time transport protocol. It presents loss rate and jitter information from RTCP receiver report (RR) packets in a tabular format. The table can be sorted by various parameters and the recent history of reports from a particular receiver can be viewed in a stripchart.

TEXT/ REFERENCE BOOKS

1. Gil Held, "Network Design: Principles and Applications (Best Practices)", Auerbach Publications, 1st edition, 2000.
2. Diane Tiare and Catherine Paquet, "Campus Network Design Fundamentals", Pearson Education, 1st edition, 2006.
3. Larry L. Peterson, Bruce S. Davie, "Computer Networks: A Systems Approach", Morgan Kaufmann Publishers Inc., 5th edition, 2012.
4. William Stallings, "Data and Computer Communications", Pearson Education, 8th edition, 2016.
5. James F. Kurose, Keith W. Ross, "Computer Networking - A Top-Down Approach Featuring the Internet", Pearson Education, 6th edition, 2012.

PART-A - 2 Mark Questions

S. No	Questions (2marks)	CO	Level
1	Illustrate how the multimedia networks are designed?	CO4	2
2	Summarize the concept of simulation used in designing the network.	CO4	2
3	Interpret the concept of bridge	CO4	5
4	Justify the need for router in designing a network	CO4	5
5	Categorize the different tools used in designing the network.	CO4	4
6	Illustrate the need for planning tool used in designing network.	CO4	2
7	Interpret how the various technologies are helpful in designing a good network?	CO4	5
8	In today's scenario we need Intelligent networks. In this regard Appraise the concept of network baseline.	CO4	5
9	Comment on the factors where delays are introduced in the router	CO4	4
10	Justify how the multimedia networks are simulated efficiently with steps.	CO4	2

PART-B - 10 Mark Questions

S. No	Questions (2marks)	CO	Level
1	For establishing a good communication between multimedia networks it is very important to have an effective network. Illustrate the various simulation methods for designing the multimedia network.	CO4	5
2	Appraise on how the remote bridge is used for exchanging the information from source to destination with necessary examples and diagrams	CO4	2
3	As a data analyst you were assigned to develop a software package for detection of router delays. Suggest an appropriate technique which is used to pre-process data so that the effect of delay is removed in detail	CO4	5
4	Without training, the computer Networks cannot converge. As a network engineer expert explain the different procedures used for designing a network.	CO14	5
5	As a design engineer expert you are tasked with the development of a good network. Suggest the suitable the suitable networking tools used for designing the network and also specify how it guarantees a fast convergence?	CO4	5