THE DESIGN AND SIMULATION OF A WIDE AREA COMMUNICATIONS AND MANAGEMENT SYSTEM FOR CIM CAPABILITY

by

Michael J. Tibodeau

Project and Report submitted to the faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for degree of

MASTER OF SCIENCE

in

Systems Engineering

APPROVED:

Benjamin S Blanchard
B.S. Blanchard, Chairman

F.J. Ricci
F.J. Ricci

Ste M. Stennett
S.M. Stennett

December 3, 1996

Blacksburg, Virginia

Keywords: Systems Engineering, Network, Communications, CIM

C.V

THE DESIGN AND SIMULATION OF A WIDE AREA COMMUNICATIONS AND MANAGEMENT SYSTEM FOR CIM CAPABILITY

by

Michael J. Tibodeau

Committee Chairman: Professor Benjamin S. Blanchard
Systems Engineering

(ABSTRACT)

A hybrid systems engineering methodology has been developed and applied to design and simulate a wide area communications and management system upgrade strategy for SMC Corporation.

A feasibility study, current system description, and desired system description establish the justification for the development effort. From the operations concept, detailed operational and maintenance requirements are defined and presented to form a program management plan. An evaluation of technical alternatives based upon effectiveness factors is completed after functional analyses allocate system level requirements to the subsystem level. Design characteristics and constraints are then specified and a mathematical model is then presented that demonstrates compliance to requirements compliance and provides for design justification. Team organization, work breakdown, subsystem specifications, and test plans are then addressed. A partial training plan follows and finally recommendations for future work are presented along with conclusions.

ACKNOWLEDGEMENTS

I wish to acknowledge the assistance provided by my committee chairman, Dean Blanchard, in completing this project as well as for his support during my graduate tenure. His extensive knowledge and objective questioning have proved to be invaluable both inside and outside of the classroom.

I am grateful to Steve Stennett for giving up his personal time to serve as a committee member. His friendship, technical knowledge, rationale, and sense of humor have helped my decision-making ability, thought processes, and sanity over the past two years. Don't look at me in that tone of voice.

Sincere appreciation goes to Professor Ricci who is indirectly responsible for starting me on this path. I would neither be here nor at the level in my career today had it not been for my first graduate course with him at the helm.

Special recognition also goes to Carolyn DeMain for the support, encouragement, and perspective she has provided me. Thank you for being unselfish and having the courage to tell me to finish that which I had started. If only the Pentium chip were as reliable as you. Just kidding.

TABLE OF CONTENTS

TABLE OF FIGURES

1. Introduction

All dimensions of manufacturing are becoming more complex, diverse, and distributed. Common products such as cars can have thousands of parts, whereas integrated circuits can have millions of individual elements. Both of these products take years to design, requiring the efforts of hundreds of people from across the country.

SMC, a semiconductor corporation, faces the situation of distributed product development. The company believes that continued dispersion of manufacturing operations is likely, where the skills and capabilities of individual units located around the country must be integrated into a virtual fabrication team. Additionally, SMC believes that shorter lead times will become more important than they are today, stressing the importance of rapid technology deployment to meet growing customer demands. Apart from the pressures for shorter lead times, other challenges are created by new product trends and fabrication processes. The company feels that responding to these challenges will take enormous flexibility, with rapid response times at the appropriate levels to new information and changing constraints.

1.1 Operations

Upper management at SMC is concerned about corporate information dissemination channels. SMC's operations involve the cooperation of engineers from many disciplines. A hierarchical, functional organizational structure has resulted in product design where information is passed "over the wall" to other departments. With shorter deadlines and increased amounts of data, people need excellent means to communicate and to exchange information. Information about requirements, costs, materials, schedules, and constraints must flow freely between the different engineers and managers. For example, information about the appropriate length of time for photolithography exposure or temperature for deposition is needed by the shift supervisor, the on-site crew, the designers, and the machine's manufacturer. The information is crucial for situational awareness. High-speed, reliable information dissemination channels are necessary to the link the workers in the fabrication cells to the schedulers, business, and design systems. Thus, a robust dissemination mechanism must be used to provide data to all sites in a timely manner.

1.2 Corporate Vision

To ready the corporation for new process trends, shorter time to market, customer demands, and other challenges, SMC would like to have a Computer Integrated Manufacturing (CIM) environment for integrated circuit fabrication. Often, CIM is viewed as fully automating a factory. This is a paragon that would include integration of performance specifications, conceptual design, detailed design, fabrication, assembly, and test together with the ability to simulate actual use. Such a comprehensive system will not be possible for many years. It is impractical for SMC to jump directly from current procedures to computer-controlled processes; an incremental approach is necessary. An important part of this vision can be achieved by progressing toward a new definition of CIM, defined as Communication and Information in Manufacturing.

One of the first steps in migrating from a traditional fabrication environment to CIM is to create a foundation information dissemination and communications facility for better and more reliable access to data. The idea behind a solid, yet flexible dissemination channel is to have a seamlessly integrated system so that information needed from one part of the enterprise by another part will be transported with minimal difficulty. For example, plans for new services, processes, and products will be directly transferred electronically from development into production, significantly reducing the time between design and realization.

The establishment of communications structures helps to reduce the time and cost of coordination to help improve productivity. Through communications, each fabrication cell, design center, and office will have access to information pertinent to the operations of the business.

1.3 Background Information

SMC has six campus facilities located in the western half of the country that operate mainframe computers, personal computers, and department workstations in semiconductor fabrication cleanrooms, product engineering and design centers, and regular business, sales, and support offices.

Due to the nature of the integrated circuit fabrication industry, information is time critical. Fabrication is performed on shifts and requires communications and information to be available twenty-four hours per day, seven days per week.

The six sites are connected via X.25 wide-area network links of 56 kbps bandwidth. Response time from the network is slow due to saturated communications lines. In recent months, transmission delays have increased from two minutes to over eight hours. Data are being dropped from the network, reducing node reliability to less than seventy percent. Moreover, network management is practically non-existent. There is no ability to forecast that a potential problem is forming and no capability for centralized network control.

Due to these problems, costs have increased and productivity rates have dropped. Support and maintenance costs have increased twenty percent each year for the past three years and are expected to continue at this rate. Expansion costs have increased forty percent. With slow data transfer due to the unreliable network, employees cannot guarantee on-time delivery of products and services. Subsequently, morale and work in progress time have deteriorated along with the network.

Network utilization has reached an all time high and continues to increase. With the introduction of CIM systems, the increase in network traffic, from cleanroom to cleanroom and from design center to cleanroom, will be enormous. Projected average traffic levels for all six sites will increase into the megabit per second range, with minimum loads exceeding current system capacity.

The communications network is slowly crippling SMC. An ad-hoc network does not provide the utility or growth capacity for SMC to remain competitive.

1.4 Project Objective

SMC needs a facility to ensure fast, reliable data transfer. The CIO has been chartered to develop, install, operate, and maintain "SeMCom," an upgrade for the existing wide area communications network. She is aware of the difficulties in completing this task. Since there are no corporate standards, she has requested a systems engineering approach to develop the design. There is a budget of six million dollars to implement an upgraded WAN infrastructure for a four year period of time, capable of expanding in year two.

The objective of this project is to apply the systems engineering process to develop an upgrade design of the existing SMC network infrastructure to support information dissemination and communications channels for future CIM capability. There are five associated goals to be achieved.

1. To improve the uniformity and reliability of data communications
2. To improve network management
3. To increase network speed
4. To reduce life-cycle costs
5. To improve expansion capabilities

There are two major drivers behind this framework. The first is to make SeMCom available as soon as possible and the second is to evolve this goal framework as technology and requirements evolve. When completed, SeMCom shall provide a solid foundation for information dissemination and wide area communications by linking the distributed business units.

2. The Systems Engineering Process

The systems engineering process forms the basis for this project. A hybrid version of the process was tailored to meet the project's needs.

2.1 Introduction

The systems engineering process utilizes analytical techniques to establish and evaluate the design, operation, and maintenance of a system. It provides a logical, thorough methodology for system development. The process is useful because it coordinates the design, engineering, and business activities related to a system's life-cycle. These activities include the conceptual, preliminary, and detail designs, along with the installation, operation, maintenance, and retirement aspects of the life-cycle.

2.2 Conceptual Design

The conceptual design allows the engineer to evaluate project feasibility, to define system requirements, and to begin advanced planning. Research is an extensive part of conceptual design. A conceptual design requires a needs analysis with a descriptive and a normative scenario. The descriptive scenario describes the current condition to be corrected or improved by the system starting to be designed, while the normative scenario describes the situation as it will be when the project is fully operative. After the needs analysis is complete, indices of performance are defined to provide measurable goals and criteria. Additionally, cost data is collected and risk and sensitivity analyses are performed. Concurrently, operational concepts are created that later lead to the collection, definition, and refinement of requirements. Finally, advanced system planning, which includes organizational, cost, and work breakdown structures, along with a project master schedule, is completed in preparation for the preliminary and detail designs. A flowchart summarizes the elements of the conceptual design in Figure 1.

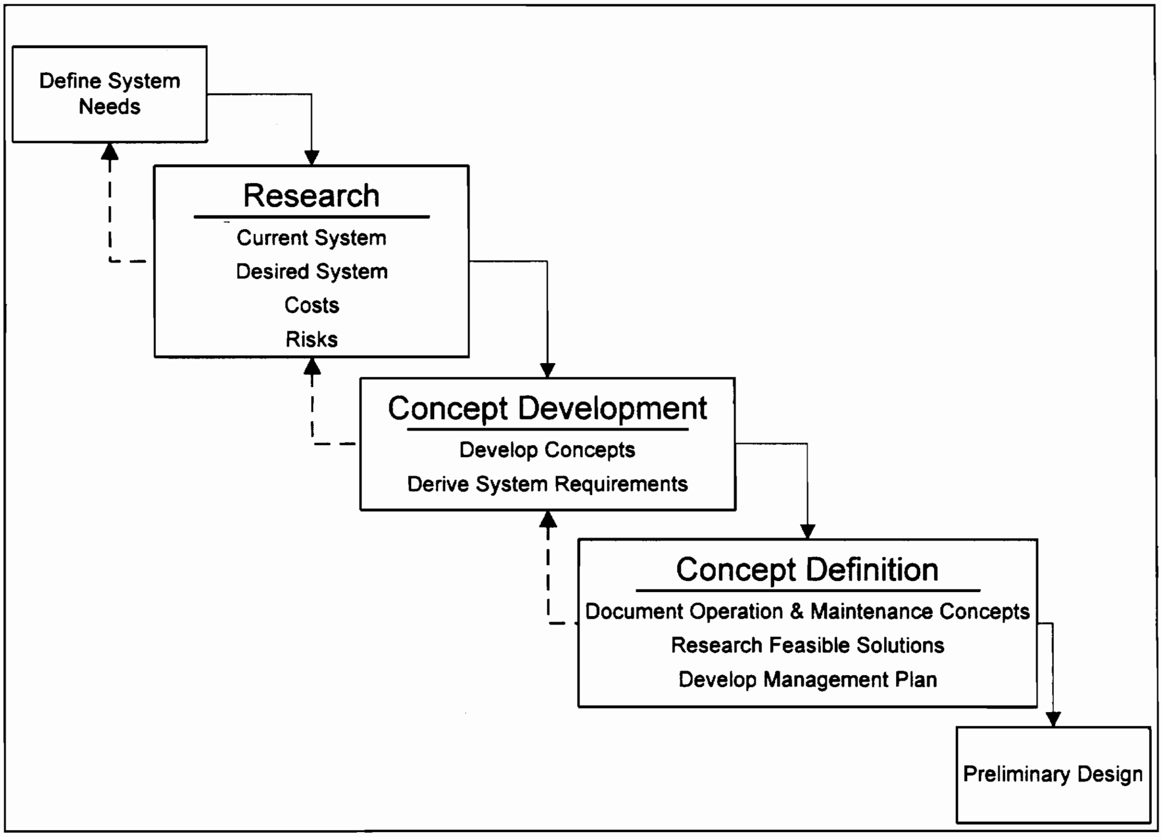

Figure 1: The Conceptual Design

2.3 Preliminary Design

Preliminary design begins with the technical baseline obtained from the conceptual design. Functional analysis and requirements allocation help to refine initial operational concepts. Refinement of functional requirements is completed via the iterative systems engineering process, incorporating continuous evaluation and feedback at all levels. This is illustrated in Figure 2. Requirements allocation provides a baseline for the development of alternative candidate solutions. Lastly, comparing and evaluating specific functional requirements, operational requirements, and system constraints refines the design to improve the system.

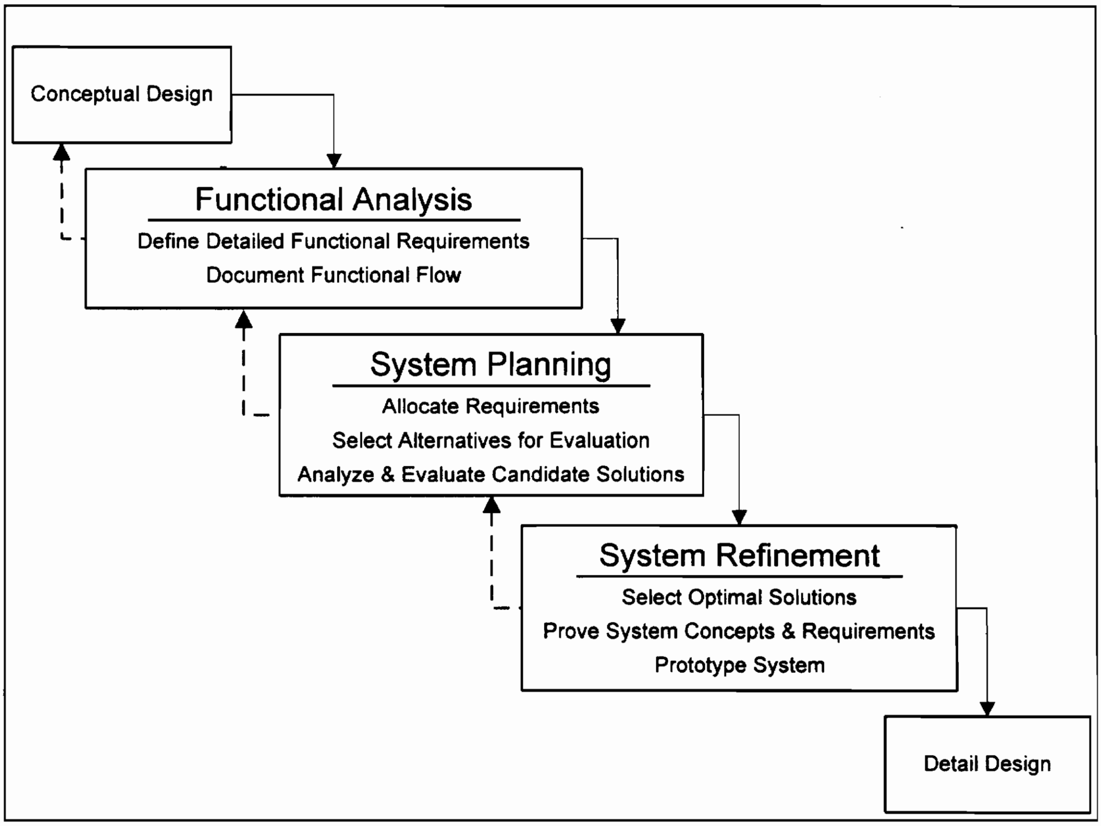

Figure 2: The Preliminary Design

2.4 Detail Design

Detail design begins with the configuration established in the preliminary design and continues to further define system and subsystem specifications. Here, the engineer should provide detailed design documentation to establish the development of a final simulation or model prototype. Again, the iterative approach is necessary to provide continuous evaluation and feedback mechanisms, as shown in Figure 3.

2.5 Summary

This project is a hybridized application of the systems engineering process that incorporates methods from Blanchard, Fabrycky, Gibson, and Sage. The overall process includes a feasibility study establishing justification for a system upgrade, detailed operational and maintenance requirements, a program management plan, and a draft training plan. Functional analyses provide the medium to allocate requirements to the subsystems. Based on budget constraints and performance measures, alternative candidate solutions are formulated and evaluated. Then, system specifications, simulation, testing, implementation, and operation are addressed. Finally, conclusions and recommendations for future work are presented.

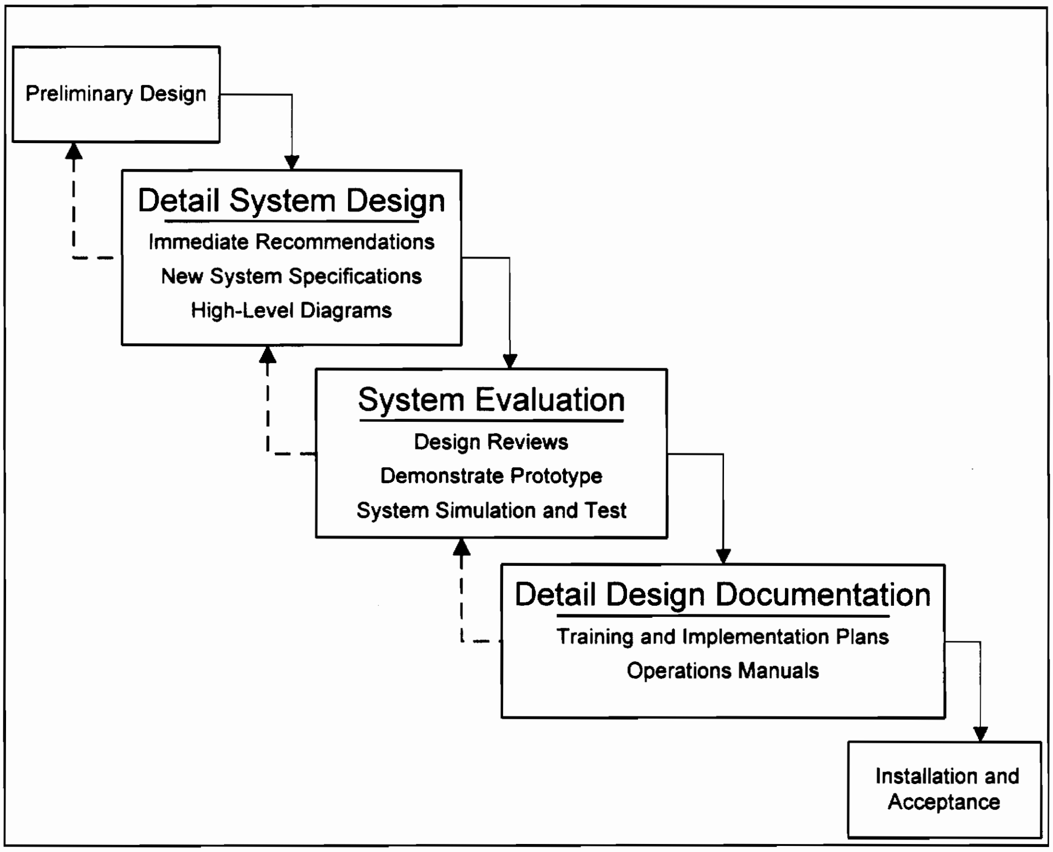
Preliminary Design
Detail System Design
Immediate Recommendations
New System Specifications
High-Level Diagrams
System Evaluation
Design Reviews
Demonstrate Prototype
System Simulation and Test
Detail Design Documentation
Training and Implementation Plans
Operations Manuals
Installation and Acceptance

Figure 3: The Detail Design

The system engineering activities outlined above establish high-level guidance. This is not meant to be a cookbook approach to engineering a communications network. Rather, it is an initial design description of SeMCom using the foundation of the systems engineering methodology. The methodology requires thorough analysis and continuous evaluation and feedback. Such an iterative approach forces the review of and correction of the system and its supporting processes, with the objective being to advance the communications network upgrade through its conceptual, preliminary, and detail designs.

3. Conceptual Design

The conceptual design requires a significant amount of research and information gathering. Fact-finding missions help the engineers to make sound decisions on ideas and proposals. Research provides the data necessary to evaluate the feasibility of a concept, develop requirements, and begin advanced system planning. A needs analysis documents the current system and its inefficiencies and maps them to future system functional requirements to provide justification for the project. Information gathering includes identifying alternatives, examining current trends and developments, and incorporating life-cycle costs and associated risks. Based upon a feasible alternative, operational concepts and system objectives are developed. Advanced system planning, including project organization, is completed in preparation for the later design phases.

3.1 Feasibility Study

The feasibility study provides justification for SeMCom, the SMC communication system upgrade. A needs statement describes the current condition to be corrected or improved by the proposed system and the situation as it will be when the project is fully operative. An extensive investigation for alternative candidate solutions identifies products, services, and methods capable of fulfilling the current system's deficiencies. Life-cycle cost evaluations, risk assessments, and further analyses are then conducted once the decision to proceed with the upgrade design is made.

There are several feasible alternatives that meet the minimum requirements for connectivity. Various technical approaches were evaluated for communications and information dissemination systems including mainframe computers, satellite relays, and various private ground networks. The current state of mainframe computing technologies and their probability of inexpensive and quick advancement at this time is inadequate in meeting project performance and timeline requirements. Although a satellite system could meet time deadlines, it would more than likely exceed budget limitations fourfold. Only a few public and private ground networks have demonstrated the potential for meeting the performance and timeline requirements. It is recommended that the project proceed with a baseline design concept being a scaleable, standards-based internetwork with both public and private components.

3.2 Descriptive Scenario

Dissemination channels are currently supported by a combination of mainframe, network, and direct access communications services. Much of the information technology at SMC can be characterized as "islands" that are unable to communicate with one another due to incompatibilities in their representation of similar information. Network management is almost non-existent. The staff can neither track current problems nor forecast that a potential problem is forming. Additionally, there is no capability for centralized network control. Communications are insufficient to keep fabrication processes, stations, and engineers on schedule and in synchronization.

SMC has six fabrication facilities that are connected via X.25 network links of 56 kbps bandwidth. Workstations are unable to exchange data with mainframes and even other workstations because response time from the network is slow and data packets are being dropped due to communications lines being saturated with data. In recent months, SMC has experienced significant degradation and failures on its communications network. Transmission delays have increased from two minutes to over eight hours. Overall node reliability has been reduced to less than seventy percent. Costs have increased and productivity rates have plummeted. Support and maintenance costs have increased twenty percent each year for the past three years. Expansion costs have increased forty percent due to the aging of the system. Due to the nature of the semiconductor industry, information is time critical because semiconductor components must traverse a number of re-entrant processes, with inspections, re-work, and approvals at major points in their production.

Because of the delays and gaps in communications, production times on the shop floor have increased. It can take hours or even days for engineers to receive manufacturing reports and design information and subsequently take appropriate actions, if they receive the information at all. Up to seventy percent of a semiconductor wafer's cycle time is spent in queue, awaiting information for inspection and approval.

The lack of effective communications is slowly crippling SMC. An ad-hoc, non-standardized computing infrastructure no longer provides the utility, response time, or capacity for SMC to remain competitive because it limits flexibility and productivity. The state of the major services of the existing system are described below.

3.2.1 Directory Services

The directory system is central to all network applications and services. It provides the means for applications to locate and share information with one another. Therefore, it must be available at all times, resistant to failure, and scale well to provide excellent performance independent of the number of users or objects. Additionally, the directory must be able to accommodate change with minimal disruption to operations.

A number of directory service products are currently used to provide directory interfaces and structures. These include Novell Directory Services (NDS), Windows NT Directory Services (NTDS), Domain Name Server (DNS), Network Information Service (NIS), and Systems Network Architecture. The directory services enable synchronization of the local and wide area network directories. However, there are incompatibilities between the directory services that lead to more overhead and increased load.

3.2.2 Network Management Services

The framework used to describe SMC's network management can be characterized as ad-hoc. Networks grew unexpectedly fast and the need for management appeared. The CIO's staff also reports that approximately 500 trouble calls to network operations personnel are made daily, but that they cannot track trouble tickets or forecast that a potential network problem is forming. Currently, there is little management capability and no capability for centralized network control.

3.2.3 Campus Connectivity

The campus networks consist of groups of Ethernet LANs constructed according to the Institute of Electrical and Electronics Engineers (IEEE) 802.3 specification, now superseded by the International Standards Organization/International Telecommunication Union (ISO/ITU) 802.3 specification. As the number of hosts required to be connected grew, repeaters were inserted to extend the network and the number of hosts capable of being supported. Since the above specifications limit the number of repeaters that may be used, bridges were used to extend the network further. This created congestion and reduced reliability. To increase reliability, redundant bridges were installed that caused problems with broadcast management and protocol isolation.

Besides speed and reliability, the foremost problem facing SMC's network is the lack of containment of broadcast packets. Protocols such as Address Resolution Protocol (ARP), Route Information Protocol (RIP), IPX Service Access Protocol (SAP), and IPX RIP rely on broadcast packets. For these layer three broadcasts to reach all hosts on their segment, they are transmitted with the broadcast Media Access Control (MAC) address. Because bridges act on layer two information, all broadcasts must be forwarded. With a large network, broadcasts increase and consume an appreciable percentage of bandwidth. Broadcasts must be processed by all workstations, increasing the amount of workstation central processing unit (CPU) time devoted to the processing of these protocols.

The rollout of Microsoft Windows NT workstations has contributed to a steady increase in network bandwidth and CPU utilization. An analysis of Sniffer traces revealed an excessive number of NetBIOS name translation broadcast packets traversing the network. The nodes are configured with helper addresses and user datagram protocol (UDP) forwarding to facilitate the xdmcp broadcast first method of connecting to the servers and balancing load. Windows NT uses the NetBIOS name resolution mechanism to translate IP addresses to network resource names in an attempt to establish connections. Broadcast packets associated with UDP ports 137 and 138 are used to facilitate the NetBIOS name resolution. The segments where the servers associated with the helper addresses reside are flooded with broadcast storms. Here, network utilization has increased twenty-five percent.

3.2.4 Wide Area Communication

The wide area communication network is based on the packet-switching standard X.25. This ITU specification defines an interface protocol that allows a Data Terminal Equipment (DTE) to interface to a Data Circuit-Terminating Equipment (DCE). There are three levels of technology: the physical, data link, and packet levels. The packet level specifies a virtual circuit service. Virtual circuit service is offered in two types: a virtual call using a switched virtual circuit (SVC) or a permanent virtual circuit (PVC). SVCs are dynamically established, consisting of setup, transfer, and clear phases, whereas PVCs are permanently established and do not require setup and clear phases.

The capabilities of X.25 networks are being challenged by heavy-bandwidth applications like Computer Aided Design (CAD), CIM, and multimedia. CAD, CIM, and multimedia applications are becoming increasingly important to enhance productivity and reduce cost. Thus, these and other heavy-bandwidth applications will continue their unprecedented growth in the semiconductor industry. As a result, X.25 is becoming outdated. A high-level diagram of the network is shown in Figure 4.

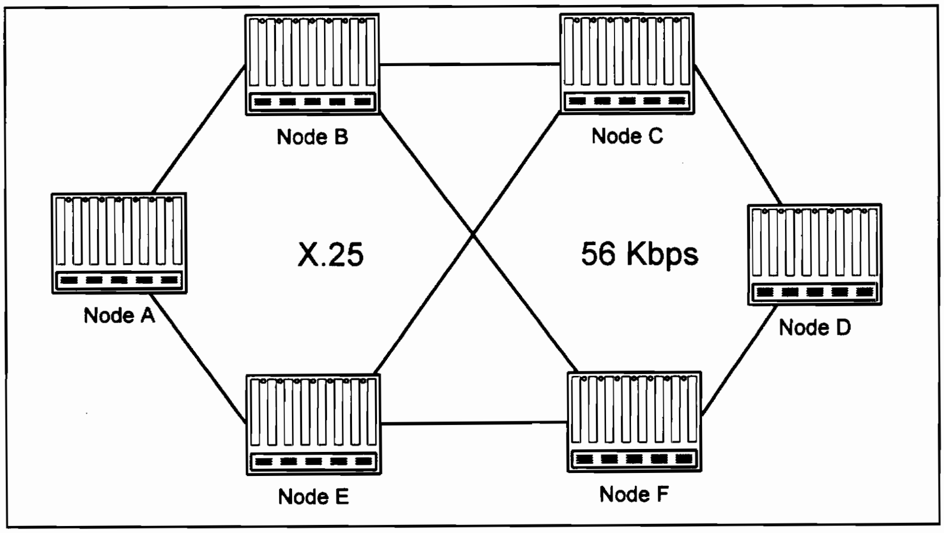

Figure 4: SMC's Current WAN Infrastructure

The network is very limited in its capacity. Figure 5 shows today's demand for WAN throughput and what the current network is capable of handling. Notice that the average demand of sixty-three kilobits per second exceeds the network's capacity.

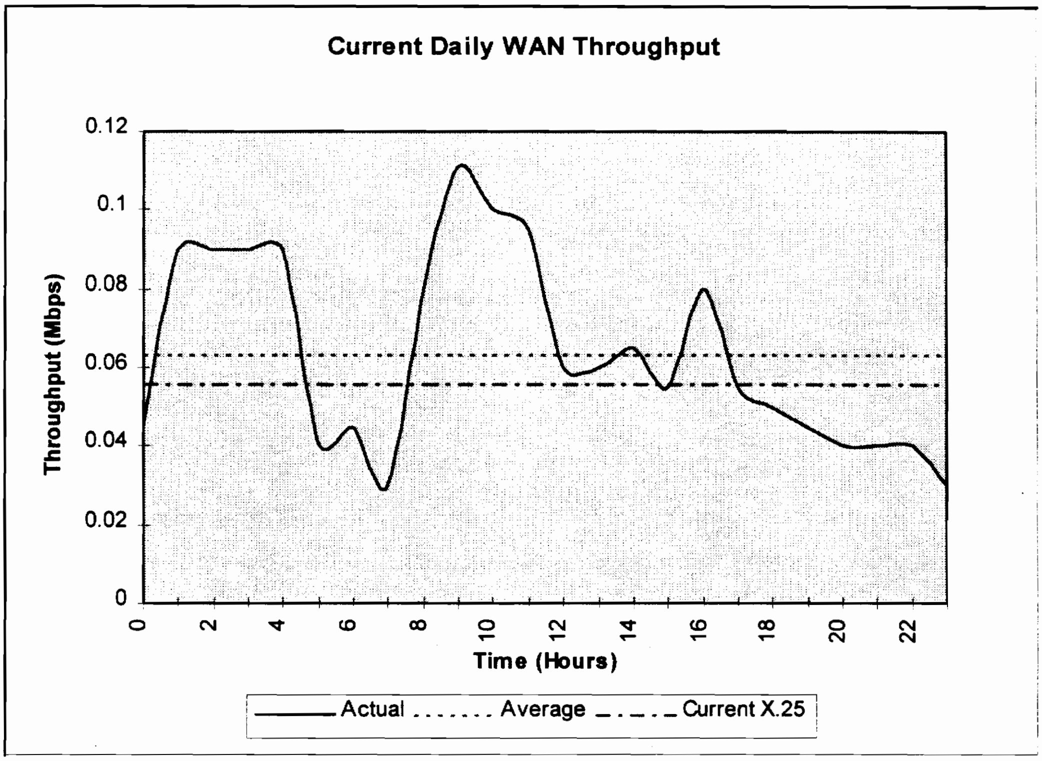

Figure 5: Current Daily WAN Demand

3.3 Problem Description

The network lacks the flexibility and scaleability to incorporate advancements in technology and to accommodate the ever-increasing system load. If a problem or break occurs, there is no workaround solution. The only alternative is to shut down temporarily and start the activity again, resulting in a significant loss of time and revenue.

3.4 Statement of Need

SMC needs a means for a fast and reliable information dissemination and communications because the current system provides limited capability to send and receive important manufacturing data. The system must be able to adapt to and manage a dynamic environment. In a time of continuous business change in the semiconductor industry, the communications needs to be as responsive as the company and its employees to add competitive advantage. To support production and business services, the system must be able to respond quickly to specific needs and eliminate unnecessary processes, queue times, and discarded data to allow the required production and design activities to proceed as quickly as possible.

SMC requires a physically distributed communications architecture. This will replace the separate and autonomous communications systems in use for the Unisys computing system, personal computer clusters, and UNIX workstations. The communications infrastructure is the first step in the evolution of the Computer Integrated Manufacturing system. Once the foundation network is in place, the production support capabilities and derivative product creation applications will be migrated to the open-system architecture and integrated by an upgrade to the communications infrastructure.

3.4.1 Description of Need

Unisys mainframe computers and the communications network are nearing their life cycle end of utility. They have surpassed full operating and throughput capacity for semiconductor fabrication processes. Requirements for support to the design, engineering, and fabrication environments necessitate that SMC improve its capability to rapidly distribute data. To avoid prohibitive support and maintenance costs, the SeMCom system will be acquired to provide a virtual factory environment, ensuring continued production and distribution support capabilities. SeMCom will improve capabilities of the current systems to provide more efficient production and dissemination of design and fabrication information for the virtual factory.

3.5 Normative Scenario

In concept, a virtual factory is one in which business units are connected electronically. Although manufacturing firms have been linked by telephone, fax, and autonomous computer systems for years, a virtual factory demonstrates significantly reduced transmission times by providing a common ground, eliminating many of the factors contributing to delays in the design and fabrication processes. Units distributed around the country could more easily coordinate their operations and schedules. An open architecture provides an efficient environment, giving SMC the ability to communicate more effectively with diverse systems, ensuring that all components share the same high-level communications facilities. It enhances the information infrastructure supporting the manufacturing and design departments, including both the internal infrastructure used within the corporation and the external infrastructure as well that links the corporation with its suppliers and customers.

For the virtual factory, SMC needs a communications facility to ensure fast, reliable data transfer. There are a number of goals to reach. First, SMC wants to be able to transfer all of its internally generated information electronically. There must be support for multiple protocols and multiple speeds for time-critical delivery of CAD drawings, database queries, CIM routines, still images, documents, and other files. Second, the infrastructure should be uniform and modular to allow for easy modification. Services must be continuously available, made possible through fault-tolerant techniques. Third, in the future, SMC expects to implement video conferencing. Therefore, the system must be upgradeable to support heavy multimedia bandwidth requirements and respond to the projected annual ten percent rate of growth. Fourth, a management system is needed immediately to manage resources, reduce life-cycle operations and maintenance costs, and promote increased system security and trustworthiness.

The use of a virtual factory by all kinds of personnel and equipment to exchange all types of data requires high bandwidth, greater dependability, greater support, and better interoperability. With growth expected to continue at ten percent per year, the system must have a capacity of one thousand five hundred users per node now, with plans for capacity expansion to three thousand users per node. The upgraded system will be capable of supporting the expected daily bandwidth demand as shown in Figure 6.

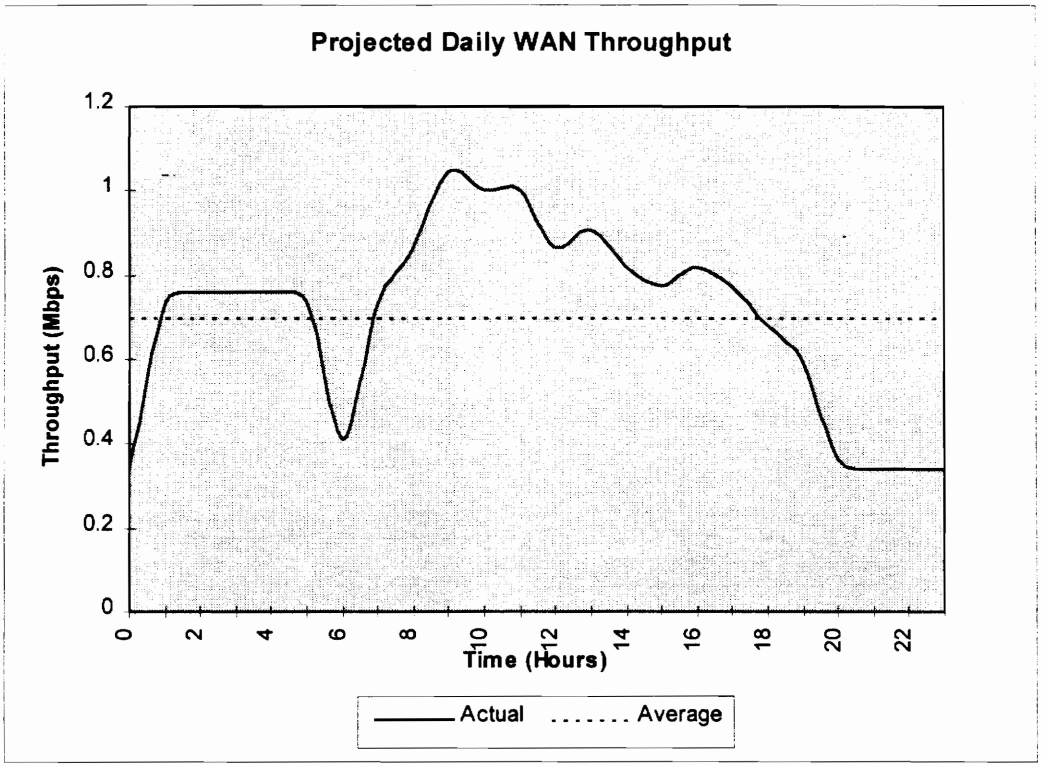

Figure 6: Estimated Daily WAN Demand in 2000

Due to the critical nature of fabrication data, the system cannot have a failure rate greater than one failure per year and any system failures must be corrected within six hours. The upgraded infrastructure's total life-cycle cost must fall within the six million dollar budget stated by SMC.

3.6 Goals and Objectives

SMC wants a communications facility to ensure fast, reliable data transfer. The objective of this project is to apply the systems engineering process to develop an upgrade design of the existing SMC network infrastructure to support information dissemination and communications for future CIM capabilities. There are five associated goals to be achieved.

1. To improve the uniformity and reliability of data communications

   Increase node availability from seventy percent to ninety-five percent.

   Remove the numerous, non-interoperable directory services and recommend a standard, interoperable network directory service.

Implement centralized services to support CAD, database, CIM, and office automation applications.

Implement a standard backbone architecture.

Reduce network failures from one per day to one per year.

Change from ad-hoc maintenance procedures to structured operations and maintenance procedures.

2. To improve network management

Increase the number of network operations centers from zero to two by providing centralized primary and secondary sites.

Recommend a Simple Network Management Protocol (SNMP) management application with performance management, configuration management, accounting management, fault management, security management, and forecast capabilities.

3. To increase network speed

Reduce the number of data communication protocols to a limited, standard set.

Segment the large, bridged network into multiple subnets with broadcast isolation.

Increase node throughput rates from 56 kbps to support projected traffic levels.

Select a WAN protocol that is more efficient than X.25.

Reduce reception time for a ten megabit file from ten hours to ten seconds.

4. To reduce life-cycle costs

Upgrade unsupported equipment to current supported products and services.

Ensure that total life-cycle costs over four years are less than six million dollars.

Retain as much of the existing infrastructure in the network as possible.

Reduce maintenance cost increases from twenty percent to one percent per year.

5. To improve expansion capabilities

   Incorporate capacity planning into the designs to handle ten percent growth per year, from 1500 users per node to 3000 users per node.

   Use standards-based technologies capable of growth.

An objective of SeMCom is to add functionality to the current system to enhance the timeliness, efficiency, and effectiveness of personnel and services. SeMCom will capitalize upon existing information sources and communications channels wherever possible to create a uniform architecture. Uniformity can help to eliminate problems. Although problems will always occur, a uniform architecture and a robust management system can help prevent them.

3.7 System Concepts

System operation, maintenance, logistics and support, and test and evaluation concepts are presented in the following sections.

3.7.1 Mission Definition

The mission of the system is to replace the current communications architecture with a reliable, manageable, standardized means for communication and information dissemination in electronic format. It will be operated all year long, twenty-four hours per day, seven days per week, in a climate-controlled environment. The system will be required to operate at any time, requiring it to operate under peak utilization periods.

3.7.2 Use Requirements

Typical operation during the week will include both low, normal, and high utilization periods. The system shall be capable of functioning at eighty percent of peak utilization for an extended period of twelve hours. Actual and estimated utilization curves are plotted in Figures 5 and 6. Advanced maintenance may require portions of the system to be turned off for a short duration of time.

3.7.3 Operational Deployment and Distribution

The system is to be deployed and operated in the western United States, connecting SMC's campuses. It shall be configured and placed on-line as the old system is removed. All sites shall be upgraded by December 31, 1996.

The system shall be capable of capacity expansion to at least two times its original design. Capacity expansion will not begin until the end of year two. During year two, expansion needs will be assessed and then appropriately implemented.

Emphasis will be placed on maximizing the use of commercially available components and non-developmental items where performance, reliability and cost requirements can be achieved. Transportation of components to all sites prior to cutover will be required. Components will have to be transported to depots throughout the United States for advanced maintenance.

3.7.4 System Maintenance

SeMCom will have three tiers of maintenance: basic, organizational, and vendor. Company personnel are responsible for basic and intermediate support that will be provided at each site. Vendors shall be responsible for advanced on-site and service facility support.

3.7.4.1 Organizational

Organizational support consists of corrective and preventive maintenance. Preventive maintenance activities include routine inspections, administration, system backups, provisions for security, moves, adds, and changes. Corrective maintenance activities include limited repair, replacement, and configuration of components. Personnel shall isolate and repair faults for components, excluding the cable plant. Basic support personnel will consist of four technicians at each campus. Each technician must provide intermediate knowledge in hardware and operations. See Figure 7 for the system support flow.

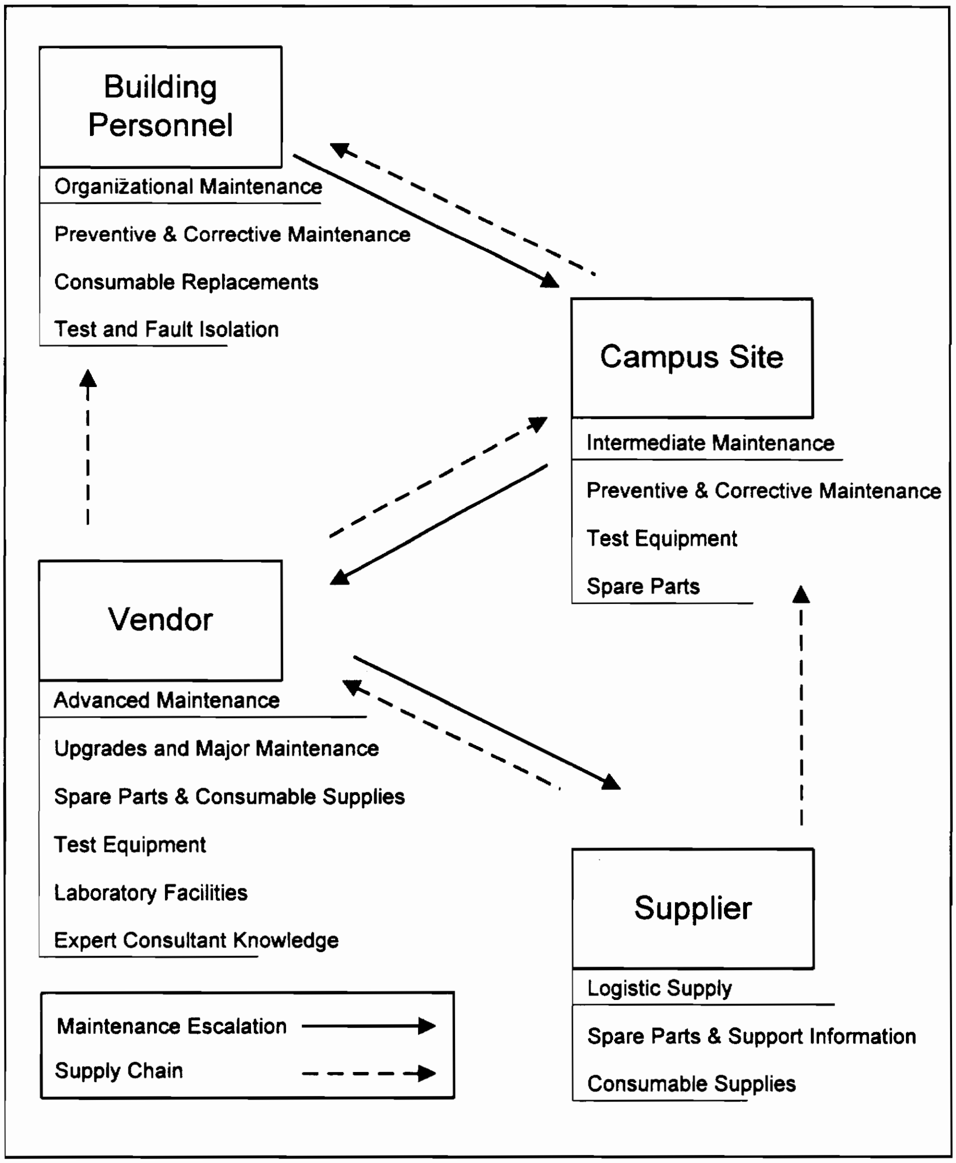

Figure 7: Support Structure for SeMCom

3.7.4.2 Intermediate

Intermediate support is responsible for preventive and corrective maintenance at all six sites. Preventive maintenance activities include optimization, security, proof of concept, and strategic planning. Corrective maintenance includes fault isolation and troubleshooting on the component level, component repair and replacement, and system analysis. Intermediate level support personnel will consist of two engineers per campus with advanced knowledge of the domain and operations.

3.7.4.3 Vendor

Vendor support is the vehicle for critical failures. Because the system will consist of diverse components, major repairs shall be conducted by the vendors and suppliers. Support from the vendor shall include, but shall not be limited to initial installation and configuration, major on-site and off-site service, and upgrades. Vendors shall offer expert level knowledge of their system components. Each campus location within SMC is responsible for arranging and maintaining service level agreements with their vendors for advanced support. See Figure 8 for the maintenance concept flow.

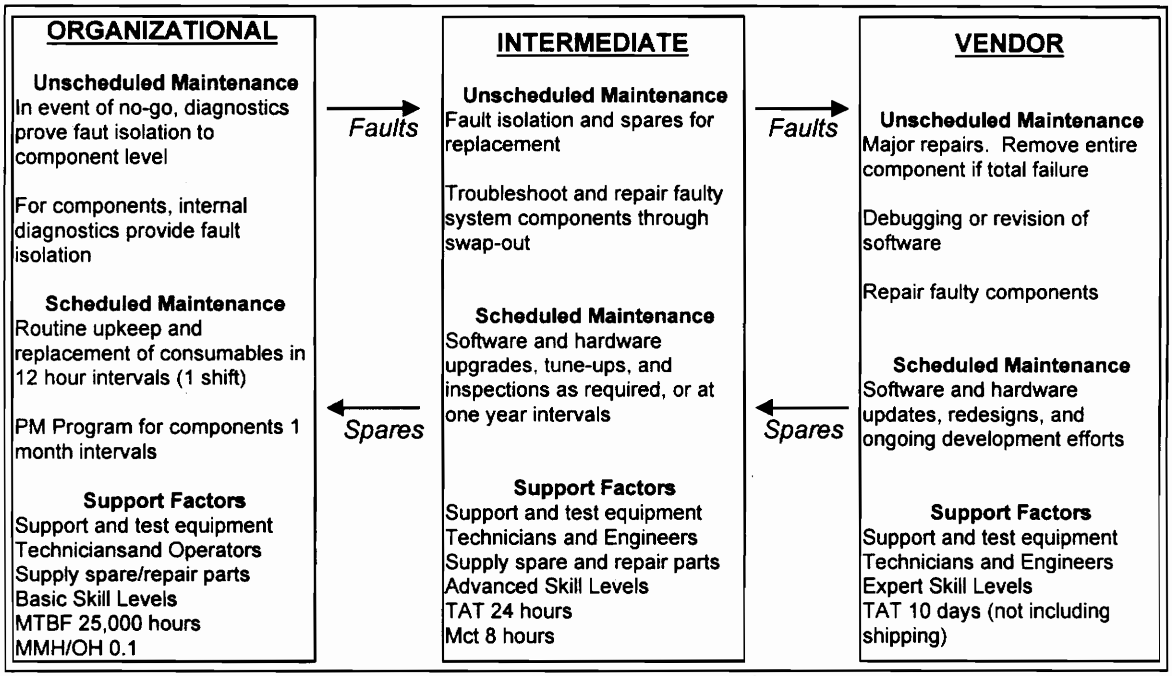

ORGANIZATIONAL

Unscheduled Maintenance
In event of no-go, diagnostics prove faut isolation to component level

For components, internal diagnostics provide fault isolation

Scheduled Maintenance
Routine upkeep and replacement of consumables in 12 hour intervals (1 shift)

PM Program for components 1 month intervals

Support Factors
Support and test equipment
Techniciansand Operators
Supply spare/repair parts
Basic Skill Levels
MTBF 25,000 hours
MMH/OH 0.1

INTERMEDIATE

Unscheduled Maintenance
Fault isolation and spares for replacement

Troubleshoot and repair faulty system components through swap-out

Scheduled Maintenance
Software and hardware upgrades, tune-ups, and inspections as required, or at one year intervals

Support Factors
Support and test equipment
Technicians and Engineers
Supply spare and repair parts
Advanced Skill Levels
TAT 24 hours
Mct 8 hours

VENDOR

Unscheduled Maintenance
Major repairs. Remove entire component if total failure

Debugging or revision of software

Repair faulty components

Scheduled Maintenance
Software and hardware updates, redesigns, and ongoing development efforts

Support Factors
Support and test equipment
Technicians and Engineers
Expert Skill Levels
TAT 10 days (not including shipping)

Figure 8: Maintenance Concept Flow

3.7.4.4 Repair Policies

Warranties will be established for specified components and will be replaced or repaired at the expense of the vendor. All other repair parts and labor associated with system maintenance will be paid for by the system owner.

3.7.4.5 Associated Effectiveness Requirements

No customized support equipment for the system shall be necessary. All support equipment shall be available as general commercial products, easily used and maintained.

3.7.4.6 Maintenance Environment

The maintenance environment will be equal to or less strenuous than the operating environment, since all advanced level facilities are equipped with climate controls. Table 1 summarizes the maintenance environment.

Table 1: Maintenance Environment Summary

Criteria	Organizational	Intermediate	Vendor
Where	Operational Sites	Operational Sites	Vendor Site
Whom	Area Personnel (Intermediate Skills)	Semi-mobile units (Advanced Skills)	Vendor (Expert Skills)
Equipment	Company owned	Company and Vendor owned	Vendor owned
Type of Work	Visual Inspection Minor service adjustments, calibration, and replacement	Overload from organizational Upgrades Troubleshooting	Overhauls, rebuilds Major service, Complicated adjustments Detailed calibration

3.7.5 Logistics and Support

Support is often overlooked or purposely ignored either due to inexperience or resource constraints. It involves maintenance, analysis, and logistics. Logistic support includes obtaining spare parts, tools, consumables, and schedules needed to support the primary system components.

The support concept identifies routine procedures and contingency plans. Procedures must deal with standard operations, routine maintenance, network component specifications, acquisition and use of test equipment, transportation needs, personnel requirements, training, and emergency action. Network support functions are typically provided by the management system. Support is necessary for: performance, configuration, accounting, fault, and security management.

Performance management measures available aspects of system characteristics so that goals for performance can be established and maintained. Global variables that might be tracked are throughput, response time, and utilization. Managed entities are continually monitored against certain thresholds. For example, an alarm may generate when the number of re-transmissions on a certain segment exceeds a threshold on a preset time period. When performance management is correctly applied, personnel become proactive in maintaining the system as potential weak areas can be noticed, isolated, and corrected before they have impact on the rest of the system.

Configuration management facilities allow managers to exercise control over the configuration of the components. Certain configurations may be modified to alleviate congestion, isolate faults, or meet changing user needs. For example, routing tables may be changed to reduce traffic loads on heavily used segments.

Accounting management uses raw utilization, response time, and throughput to measure the effectivity of systems deployed. This information often becomes the cornerstone of maintenance plans, upgrades, and new designs to accommodate changes in usage patterns or new requirements. Accounting management is also used to feed chargeback systems that are deployed to influence usage behavior and apply costs to appropriate business departments.

Fault management detects, logs, notifies users, and when possible automatically fixes problems so as to minimize system downtime. It is this support component that is most often deployed in system environments today. Fault management coupled with redundancies built into components can assist in improving system uptime and response times to fixing certain problems.

Security management controls access to resources according to certain guidelines so that the system cannot be sabotaged or sensitive information compromised. The security management subsystems categorize resources into areas of authorization. Some resources are considered public and open; other areas are restricted and may only be accessible by a manager or administrator. Resource access is tracked and stored in log files that are reviewed by the system manager. The system may alert the manager if certain resource access is attempted.

Logistic support data are required for system planning and operation. Specific types of data include frequency, timing, and sequence of maintenance tasks, details such as quantity, type, and cost of special test equipment or environmental facilities, the skill levels of personnel and their training and labor factors, the types and numbers of supplies to keep in stock, and any other additional performance factors. All of this data must be integrated to effectively coordinate resources and manpower among the five areas.

3.8 Researching Feasible Alternatives

Alternative candidate solutions capable of satisfying a need are identified and developed using certain sets of criteria. The criteria should be supportive of, interoperable with, and endorsed by a large spectrum of the industry. Selection criteria should assess the potential value added, costs, and risks of the alternative candidate solutions. The selection criteria should always include, but not be limited to, the system's functional goals. A baseline set of criteria used to evaluate products and services is presented below. The two main categories are technical and market-based.

Technical criteria can be divided into three groups: compatibility, functionality, and stability. Compatibility considers the level of coupling or interoperation that a candidate solution has with another product or service. Functionality is a measure of the extent to which the product or service meets the functional requirements defined for the proposed system. Stability evaluates the technology base of the candidate solution itself in terms of current and potential development.

Market-based criteria can also be divided into three groups: availability, consensus, and maintainability. Availability measures the variety of products to implement the candidate solution. This is related to market share and market demand.

Consensus considers the extent to which a product or service has been adopted by industry. This is also related to market share, and indirectly to customer loyalty. Finally, maintainability evaluates the level of and history of maintenance and support for the product. The alternatives presented are considered feasible, but not necessarily optimal.

3.8.1 Directory Services

Every component within a network needs some type of directory addressing or naming. If a component has a name, it must be mapped to its corresponding address. This type of service is essential for activities involving interactions between components, the basic mission of a communications network. A number of naming technologies and approaches have been developed. Three primary examples are the Domain Name Service (DNS), Network Information Service (NIS), and the ITU X.500 Global Directory.

3.8.1.1 Domain Name Service (DNS)

DNS is part of the Transport Control Protocol Internet Protocol (TCP/IP) protocol suite application layer service. It is also known as the Berkeley Internet Name Domain (BIND) that implements an internet name server. The BIND consists of a server, or daemon, and a resolver library that provides mappings between hosts and network addresses.[1] It allows stores name resources or objects and shares this information with other objects in the network. The service is distributed amongst a group of name server hosts. Clients use a resolver to translate queries. A resolver must know the address of at least one name server to process queries and distribute data throughout the network. This in effect is a distributed database system for objects in the network.

The DNS structure is hierarchical, consisting of nested functional domains. Each domain name must be unique and properly registered within its level of the hierarchy. The root of the DNS hierarchy is controlled by the Network Information Center (NIC), located in Herndon, Virginia. Specifications for this name server are defined in RFC1034, RFC1035, and RFC974.

3.8.1.2 Network Information Service (NIS/NIS+)

NIS, formerly known as "Yellow Pages," is a hierarchical and secure network information service system designed by Sun Microsystems. In concept, NIS is a simple, flat, distributed database that allows personnel to manage the network information for complex and heterogeneous computer systems.[2] Like DNS, NIS has domains of machines that share configuration information. A domain has one master server that has the canonical version of the information and slave servers that have copies of the information stored in maps. NIS uses the NIS/YP protocol to serve client requests.

NIS+ is not an extension or enhancement of the existing NIS. It is a complete redesign and rewrite. NIS+ servers are backward compatible with the NIS protocol and can serve NIS client requests. Compatibility is provided as a mechanism to ease transition. A big difference is that NIS+ uses a hierarchical database. The major advantages of NIS+ are improved speed, availability, scaleability, and security.

3.8.1.3 Global Directory

X.500 is extensible and provides a standard method of accessing names. In the X.500 model, information objects are represented by entries in the directory information base (DIB). Entries consist of at least one attribute type of at least one value. An attribute's type defines characteristics, compares the characteristics of values, and determines if more than one value is allowed.[3] Entries are positioned hierarchically in the directory information tree with a unique, distinguished name (DN) formed by tracing the path from the root to the entry's leaf node and adding the relative DN at each branch.

Different companies, organizations, and geographic regions can maintain their own DIB that links to others to form the global X.500 Directory. A global directory agent is responsible for processing requests and updates. When a user wishes to find an address, the directory is accessed by means of a directory user agent.

3.8.1.4 Directory Services Summary and Recommendations

NIS and NIS+ require a lot of manual maintenance and a significant portion of network bandwidth, its biggest weakness for a network environment. DNS works better for mixed networks, but X.500, in theory, is infinitely scaleable. However, it does not use a distributed database, which leads to problems with synchronization. Thus, the directory will probably suffer from inconsistency, naming conflicts, and inefficient speeds.[4] At one time, X.500 had strong acceptance within the commercial industry as the future directory service standard. Now, however, that appears to be fading. The direction that many software companies are taking is not to support X.500 because of the market share of DNS. Many new products that are scheduled for release will support DNS, but not X.500. Some companies mentioned that they will support DLAP, a subset version of X.500, but not the full standard. Thus, the recommendation is to use a publicly available DNS daemon on an existing UNIX workstation.

3.8.2 Network Management Services

This is a time of continuous business change. Networks need to be as responsive as the business itself to add competitive advantage. Operations costs in network computing typically exceed the combined costs of hardware and software. In some cases, it can reach more than seventy percent of the total cost of ownership.[5]

The philosophy of management in a distributed environment is one of remote management. With this in mind, every component should be accessible and remotely managed by a centralized platform. These platforms plan changes, evaluate performance characteristics, and troubleshoot problems of the network. By using backbone protocols like TCP/IP and SNMP, it is possible to move systems quickly and introduce them to different parts of the network without disruption of service to existing nodes.

Management tools for information technology personnel have not kept up with the change in the infrastructure. Earlier, one mainframe CPU could support two hundred users. Today, two hundred users might share three hundred microprocessors in their desktop servers. The need for function-oriented management tools is apparent. Tools need to be comprehensive to manage heterogeneous devices and flexible to match the growth of the network. They should support multiple users simultaneously, but not tie them to their desks. Tools should cooperate with platforms from other vendors. The capabilities of Cabletron's Spectrum, Hewlett Packard's (HP) OpenView, and Solstice's SunNet Manager are summarized below.

3.8.2.1 Cabletron's Spectrum

Spectrum was designed as a multi-user management system. Its client/server architecture has unlimited scaleability to support multiple, existing, and emerging protocols. Its Device Communication Manager (DCM) provides flexibility by freeing up the core program from the task of translating protocols. As new management protocols and client applications such as SNMPv1 and SNMPv2 and Common Management Information Protocol (CMIP) become available, the DCM can support them. Additionally, Spectrum has a new approach to network management.

Inductive modeling technology (IMT) offers some major benefits not found in traditional element management platforms. First, a model of a device will continue to provide information about the network, computer, or other entity even when the device itself fails. Second, a model of a device understands its relationship with other devices. This allows them to infer not only their own condition, but the status of other devices as well. Through IMT, Spectrum can determine which device is at fault.

3.8.2.2 Hewlett Packard's OpenView

OpenView provides consolidated system management to allow for control and configuration of large numbers of heterogeneous systems. It provides a single interface to perform a variety of tasks centrally across different platforms. For example, a function can be invoked at a central management station for execution on a remotely managed target. In addition, information required by a management station can be generated and communicated to it by the managed target.

OpenView employs a database management system for reporting and storing historical network data. It also has a Management Information Base (MIB) compiler to create an information base for network nodes that do not already have the MIB information, as well as the capability of importing third party MIBs. In addition, traffic reports and histograms can be generated. Management functions are segmented to enable network management systems to communicate with and manage each other.[6]

3.8.2.3 Sun's Solstice SunNet Manager

Solstice SunNet Manager (SSNM) is Sun's SNMP-based platform product for network management. The product provides applications for analyzing resource performance, identifying and resolving faults, tracking configuration management, and simplifying and automating management tasks. It has a large number of third-party applications that run off the platform, addressing most key areas of systems management.

Its discovery mechanism allows managers to automatically load the SSNM database with information on the internetwork infrastructure using SNMP MIB and routing table information. Link management allows users to monitor the status of physical connections between two devices. It also supports SNMPv2.

3.8.2.4 Network Management Summary and Recommendations

Hewlett Packard's OpenView is an effective monitoring device that provides autodiscovery and hierarchical map drawings of network nodes. The management environment allows for the effective management of distributed systems through a central interface. For example, a function can be invoked at a central management station for execution on a remotely managed target. In addition, information required by a management station can be generated and communicated to it by the managed target. OpenView runs on both UNIX and DOS machines, although UNIX is the preferred operating system. Given that SMC has many Sun workstations in place, it would be easy to leverage some of the existing infrastructure for network management.

OpenView provides a management environment that is scaleable, extensible, and reliable. It allows for the integration of a wide variety of component management applications. The application tools perform tasks in the necessary functional areas of operations, fault management, security management, performance management, configuration management, and accounting management.

3.8.3 Campus Connectivity

Creating a system that is uniform eliminates unique situations arising from architectural problems created by non-standard devices. Although problems can occur in any system, a uniform architecture can help prevent them.

3.8.3.1 Routers

Like bridges, routers make forwarding decisions based on information contained in unicast packets. Because of their ability to use and apply OSI layer three protocol information, they also manage broadcasts. The replacement of bridges with routers provides broadcast isolation and a reduction in the number of devices within a domain. Broadcast isolation capabilities allow the creation of routed networks that scale further.

Routers support redundant links and routes. Further, multiple routers may be connected to the same segments to spread the traffic load for redundancy and diversity. This distributes the load across all available routes, using all available bandwidth. Routers with more advanced routing protocols will load-share automatically.

3.8.3.2 Switches

Switches are new to the marketplace, but the integration of switches into the network represents a shift backward to bridged networks. Most switches can operate as either a level two or level three device. When operating as a layer two device, all segments on a switch must operate within the same layer three protocol network. In the case of IP, all hosts must participate in the same subnet; for IPX, all hosts have the same external network number. Thus, the switching device provides no isolation of protocol problems. When operating as a layer three device, the switch provides protocol isolation for layer three protocols.

One of the primary features offered by switching products is the extremely low latency offered. Cut-through switches make a forwarding decision immediately upon receipt of the destination address at the beginning of the packet, unlike store and forward switches that read the entire packet before forwarding. The disadvantage to cut-through switching is that it cannot prevent the forwarding of error frames. However, most modern devices are fairly tolerant, so many network architects have discounted this problem when considering switching applications.

3.8.3.3 Campus Communication Summary and Recommendation

Routers are tried and true technologies. To solve utilization and broadcast problems and centralize services at SMC, it is recommended that routers in a collapsed backbone architecture be the standard for campus connections.

Based on the examination of the actions of layer two switches and the lack of resolution of the bridging related to broadcast problems, deployment of these devices as backbone connection technology on networks with over five hundred workstations is not recommended. However, the use of switching has the potential to provide significant performance gains when used in high data and low latency demand environments. Therefore, when business demands, it is recommended that a hybrid routed and switched environment be considered.

3.8.4 Wide Area Communications

A number of technologies are presently available in the marketplace. For this project, emphasis will be placed on circuit-switched and packet-switched technologies.

Circuit-switched technology is predominantly the domain of the inter and local exchange carriers, although it is possible to have circuit-switched private networks. Software Defined Network (SDN), Software Defined Data Network (SDDN), and Integrated Service Digital Network (ISDN) are examples of circuit-switched technology.

Fast-packet switching technology is today mostly found in the private network environment, but is migrating toward tariff-based public service. Frame relay and cell relay systems are examples of networks based on fast-packet switching technologies.

3.8.4.1 Software Defined Network (SDN)

SDN is an AT&T Communications service that allows customers to build corporate networks using an AT&T switched network. This type of service is also available from both Sprint and MCI under different names. SDN provides premise-to-premise voice and data networking along with a powerful set of call management and control capabilities that enable customers to define uniform numbering plans for voice and data calls and to perform originating and terminating call screening functions. SDN calls can be made to on-network or off-network locations.

The key features of this service are dedicated access lines from each customer location to exchange, long distance networks, connectivity through the facilities and of the long distance network, customized call-processing instructions stored in network control points, management and control by the customer's telecommunications manager using the service management system, and performance and reliability assurance through the SDN control center.

3.8.4.2 Software Defined Data Network (SDDN)

SDDN introduces high-speed data transport to the SDN concept. It enables customers to transmit data on channels at speeds of 56 kbps and 64 kbps, extended to 1.544 Mbps. Higher speeds are obtained by stacking contiguous 64 kbps clear channels that appear to function as a single channel. Users can access SDDN service through modems, T1 lines, and the primary rate interface (PRI) of the ISDN.

Two unique qualities that SDDN brings are restoration of service after the network has been disrupted and quality performance. Some of the SDDN performance parameters are service availability, reliability, call blocking, and restoration speed.

Advantages of the SDDN include vendor management, access to new technology without capital outlay, peak traffic handling, private line backup service, access to widely dispersed locations, and dial-up capability.

Disadvantages include its foundation on circuit-switching technology, making it difficult to share bandwidth. Also, bandwidth allocation is only in 64 kbps increments, there are usage- and distance-sensitive tariffs, and limited management capabilities.

3.8.4.3 Integrated Services Digital Network (ISDN)

The central concept behind ISDN is to provide an integrated access interface whose characteristics are service-independent. Its objectives are to ensure product and service family compatibility, enhance interoperability with other products and services, unify private and public ISDN networks, and provide terminal portability. Essential ISDN characteristics include common interfaces and protocols for voice and data, functional out-of-band signaling, multiple logical channels, and end-to-end digital connectivity. Two ISDN user-to-network interfaces are used: the basic rate interface (BRI) consisting of two 64 kbps B-channels and one 16 kbps D-channel (2B+D), and the PRI, consisting of either twenty-three or thirty B-channels (23B+D or 30B+D).

Advantages of ISDN include access to new technology without capital outlay, peak traffic overflow, end-to-end digital connectivity, and integrated voice and data access. It also offers on-demand services that provide cost-effective backup solutions.

Disadvantages associated with ISDN include limited bandwidth sharing because the underlying technology is still circuit-switching, bandwidth allocation in increments of 64 kbps, distance-sensitive tariff, limited network management capabilities, separate subnetworks for voice, data, and signaling, and limited availability.

3.8.4.4 Frame Relay

Frame relay describes a switching technology and an interface standard. It is a fast-packet technology that has been streamlined for speed. Thus, its ability to statistically multiplex provides the same bandwidth sharing and efficiency as X.25 while eliminating much of the processing performed by the network, reducing delays.[7] This simplification is dependent upon the elimination of error recovery mechanisms. The frame relay interface allows bandwidth to be shared among multiple users, creating bandwidth allocation on demand. Each frame contains header information that influences the routing of the data. This enables each end point to communicate with multiple destinations via a single access link.

Frame relay has the advantage of becoming an increasingly mature technology with low overhead. It supports dynamic bandwidth allocation, logical mesh connectivity, and variable frame size.

The main disadvantage is that it is better for data-only applications due to variable delay, although standards are being developed for video and switched virtual circuits.

3.8.4.5 Multiprotocol Routers

Multiprotocol routers (MPRs) have been developed as one answer to network integration involving different protocol stacks. Routers either encapsulate a given protocol into a uniform protocol or convert the protocol data unit to the uniform one. Routing protocols are typically a variant of IP. Standards for MPRs are being developed in certain areas, such as the DLSw standard for SNA encapsulation into IP.

Advantages of multiprotocol routers include dynamic allocation and aggregation of bandwidth, minimization of application conversions, handling of diverse protocols, smaller variable delay, and logical mesh connectivity.

Disadvantages of multiprotocol routers include proprietary routing protocols, lack of complete interoperability and mature prioritization schemes, and higher overhead and delay than frame relay.

3.8.4.6 ATM

ATM is a connection-oriented technology. The connection is specified either during the call setup in the case of SVCs or provisioning in case of PVCs. ATM provides a means to integrate voice, video, and data to connect networks and combine their services into a seamless whole. Unlike most data communication technologies, ATM uses short, fixed-length cells that consist of a five byte header and a forty-eight byte information field. The header field contains information about the cell, including the address, used to switch the cell over a pre-established path.

ATM is a layered architecture. The high layer supports ATM services whereas the ATM adaptation layer (AAL) segments and reassembles the service information into cells.[8] Generally, the physical transport for the cells is provided by fiber, but other transport techniques and transmission facilities are being defined and implemented.

An ATM network contains three types of components; the transport, statistical multiplexers, and ATM switching equipment. The transport maps ATM cells into payloads to concentrate signals for more efficient bandwidth sharing. ATM switches do not use the shared medium bus architecture that most switches, multiplexers, and routers use today. Shared medium will become too impractical to implement as backplane speeds surpass one Gbps.

An ATM network has the advantages of seamless end-to-end paths, dynamic bandwidth allocation, logical mesh connectivity, quality of service negotiation, usage-based tariff, application and transmission medium independence, line speeds over fifty Gbps, and scaleable architectures.

The disadvantages of ATM include technology immaturity, relatively high cost, lack of standards and technical knowledge, and no error detection in the data field.

3.8.4.7 Simplified Delay Analysis for Packet Technologies

Networking technologies such as frame relay and ATM achieve high throughput and speed by streamlining the protocols compared to the X.25 network. The fundamental assumption in streamlining the protocols is that the transmission facilities used to support networks deliver higher transmission quality than the analog facilities for which the X.25 protocol was developed. Table 1 shows end-to-end delay under the hypothesized near maximum network load, calculated based on an M/M/1 queuing model.

Table 1: End-to-End Delay Performance

Type	Access Delay	Switch Delay	Trunk Delay	Propagation	Total
X.25	116.67 msec	300 msec	43.75 msec	15 msec	475.42 msec
MPR	60 msec	15 msec	9.91 msec	15 msec	91.91 msec
FR	40 msec	10 msec	7.07 msec	15 msec	72.07 msec
ATM	0.06 msec	0.2 msec	0.008 msec	15 msec	15.27 msec

3.8.4.8 Comparison of Packet Technologies

Packet technologies provide a wide range of options to support data and multimedia applications. One must select a solution that meets present requirements in a cost-effective way and allows a graceful transition to meet future requirements. Table 2 compares the features supported by four packet network technologies.

Table 2. Packet-Mode Technology comparison

Feature	X.25	Frame Relay	MPR	ATM
PVC	Yes	Yes	Yes	Yes
SVC	Yes	In-progress	Coming	In-progress
Speed	kbps	Mbps	Mbps	Gbps
Voice	No	Yes	Yes	Yes
Video	No	Coming	Partial	Yes
Dynamic Bandwidth	Yes	Yes	Yes	Yes
Switching	Yes	Yes	Yes	Yes
Signaling	Inband	Outband	Inband	Outband
Transparent	No	Yes	Yes	Yes
Delay	High	Medium	Medium	Low
Alternate Routing	Yes	Yes	Yes	Yes
Management		SNMP	SNMP	SNMP
Scaleability	Limited	Medium	Medium	High
Error Correction	Yes	No	Limited	Header only

3.8.5 Wide Area Summary and Recommendation

A number of parameters such as the topology of the present network, traffic profile, future applications and trends, and political realities should be evaluated before a definitive strategy can be reached.

Even though the emerging technologies such as SDN and ISDN offer a number of advantages, they are better suited for enterprises whose network traffic is predominantly voice and is geographically distributed. In such an environment, small amounts of bursty data traffic can be supported at little or no incremental cost. However, cost implications of the inherently bursty nature of an intensive data communications environment can only be exploited to a limited degree with circuit-switched technologies. ISDN might be used as a backup service to increase availability, to support dispersed locations, or as an access technology for carrier backbones.

With systems becoming more intelligent as well as more powerful, frame relay and MPRs are superior to X.25 and circuit switches in terms of performance and interface speed. In a data-only environment, pure frame relay has higher performance than MPRs. MPRs, on the other hand, offer less application conversion, superior packet prioritization, and more certifiable delay characteristics.

The advantage to ATM is a simple and flexible design. A single switch interface can handle voice, data, and video traffic and will forward cells to the correct processors based on cell service type. Minimal processing is needed for ATM routing, which reduces delays, which is particularly important for real-time voice and video transmission. The fixed cell size also reduces delay by simplifying processing at nodes and by reducing the number and size of buffers. A second advantage of the ATM approach is that the use of small cells and the ability to multiplex different types and rates of traffic over the same channel promote efficient bandwidth use. A third feature of the ATM protocol that enhances its flexibility is its ability to carry a variety of non-standard data rates. The protocol responds to different data rates and bursts of data.

ATM is still very expensive and lacking implementation standards. It can meet all of the requirements and provide a migration path to future applications and eventual integration of voice, data, and video, but it is not available in all areas and it does not offer any advantage other than better delay performance for data-only applications.

For the short-term, the recommendation is to feed a frame relay backbone utilizing MPRs. This will sustain the network cheaply for two years while allowing ample time for ATM and other technologies and solutions to develop. If a better solution emerges as the premier technology with availability through many networks supporting numerous interfaces, it can be implemented after two years as part of the capacity expansion to SeMCom.

3.9 Life-Cycle Cost Evaluations

The considerations in forming a total cost of ownership model for deploying networks include not only the capital costs, but also operations, support, and retirement costs. A high level cost analysis provides estimated net present value costs required to implement SeMCom. The expense data enables the company to consider current and future financial commitments for the amount of the extended system's implementation and life-cycle costs. Data were generated through analysis of similar networks with comparable architectures and also through present network costs. Costs are divided into design, acquisition, installation, operations, maintenance, and retirement expenses.

Design costs include an initial two month, four person task for system analysis and evaluation. Acquisition costs are based on price point estimates of the infrastructure components. Installation includes the cost of labor and any opportunity costs suffered by the business. Operations and maintenance costs include regular support and system administration, program management, and training costs. Retirement costs also include the cost of labor, disposal, and down time.

Basically, cost of ownership can be abstracted to a three category model. According to Strategic Networks Consulting, people represent forty-two percent of technology ownership costs, facilities include thirty-three percent of the costs, and equipment represents twenty-five percent of the costs.[9] Total life-cycle cost then adds in the variable operations costs. One life-cycle cost of ownership model is represented in the following relationship:

$$\text{Cost} = \frac{F + O}{L * U}$$

where F is the lifetime fixed cost, O is the lifetime operational (variable) cost, L is the useful life of the system, and U is the utilization. Fixed costs, also referred to as equipment costs, are a function of purchase price, transportation, installation, and facilities overhead. For this project, typical operational costs, or variable costs, were broken into five areas: consumables, materials, maintenance, personnel, and retirement. The simple cost equation is more complex than it appears, since inputs to the basic categories may have many comprehensive cost models. For example, the equation can be made more complex by expanding the utilization component. The utilization model can be conceptualized as:

$$\text{Utilization} = \frac{1 - (M + D + S + E)}{H}$$

where M is scheduled maintenance, D is downtime for unscheduled maintenance, S is the standby time, E is the actual usage, and H is testing time, all expressed in hours. However, for system estimates, a simple model is preferred. The budget estimate summary for SeMCom's four year life-cycle is shown in Table 3, below.

Table 3: SeMCom Budget Estimate Summary

Service	Cost
Research, Development, Oversight	**300,000**
Planning and Management	110,000
Feasibility and Documentation	50,000
Design and Test	140,000
Acquisition and Installation	**800,000**
Procurement	750,000
Personnel	50,000
Operations and Management	**3,900,000**
Personnel	650,000
Telecommunications	3,000,000
Maintenance and Supplies	170,000
Training	80,000
Capacity Expansion	**800,000**
Equipment	750,000
Personnel	50,000
Retirement	**70,000**
Contingency	**150,000**
TOTAL	**6,020,000**

3.10 Risk Assessments

For any development, an assessment and an evaluation of associated risks should be completed. Any consequences or uncertainties involved should be documented and communicated to project management, the customer, and the stakeholders.

3.10.1 Schedule

There is always a risk of design discrepancies and development delays. Management should be aware that schedules can slip. It is also important to know that because of slippage, actual costs tend to be larger than projected life-cycle costs in distributed, complex systems. Thus, they must recognize the need for insurance in the project schedule so that unexpected delays have less impact. Simply stated, there will most always be hidden costs and unexpected conditions.

3.10.2 Directory Services

Perhaps the largest potential downfall of DNS is its lack of true standardization. The growing popularity of open systems has shifted many administrators toward X.500 because of its standards-based nature. It is infinitely scaleable and allows public directory services to join a global directory. These capabilities are enhanced when coupled with other standards, such as X.400. DNS does not have this luxury.

3.10.3 Network Management Systems

There is a trend in the information systems industry towards open and distributed systems. This trend places a pervasive requirement on the ability to properly manage distributed, diverse, and legacy systems. For this case, new systems will be introduced and old ones will still have to be managed. New devices will have management capability, whereas older ones do not. Legacy devices must reviewed and a solution to manage them designed. Older systems generally must be evaluated, on a case-by-case basis. Case-by-case reviews and subsequent solutions pose additional time and cost resources, mainly in the operation and maintenance areas.

Additionally, network management systems add traffic to the network. It is theoretically possible to inflict more damage upon an already damaged network with extra management traffic. For example, a failing component on a non-redundant link could cause enough management traffic to be generated to completely fault the segment or the node. Thus, the additional traffic must be carefully calibrated and closely watched to ensure that it does not significantly degrade network performance.

3.10.4 Campus Connectivity

In the recommended campus environment, switched fast Ethernet and collapsed router backbones are frequently used to provide network connectivity. This method adds additional complexities such as configuration control between interconnecting segments, complex routing requirements, addressing structure and physical plant issues, and resource placement to meet performance, security, and growth issues. Some of these complexities can be reduced with the implementation of ATM.

3.10.5 Wide Area Communications

There are advantages and disadvantages to any wide area communications strategy. By using the recommended multiprotocol routers, more delay is accrued than by just using pure frame relay. Additionally, running multiple protocols on the same links can be a confusing situation for the operators, as well as a non-optimal situation for the routers' processors.

A partial solution to this would be to use the ATM format to accommodate different types of traffic, but it is not as efficient as a protocol designed to handle one specific kind of traffic because considerable bandwidth can be wasted by fixed length cells. For example, messages between processors in distributed computing systems are typically very short. Consequently, a forty-three byte cell must be used to transmit a one or two byte message,[10] wasting up to ninety percent of the payload.

3.11 Recommended Design

The approach that appears to be the most beneficial incorporates a strategy based on standards dealing with the various aspects of data communications systems to provide for scaleability. SMC will utilize a wide area backbone of multiprotocol routers running TCP/IP over frame relay. Network operations centers shall be established at Campuses One and Five on nodes A and F. Campus networks shall use a redundant collapsed backbone architecture for connectivity made up of dual-homed routers and fast Ethernet connections. DNS/BIND shall be used for addressing and directory synchronization, while the entire operation shall be managed using HP OpenView. The design goal is to make a uniform infrastructure to minimize costs, maximize throughput, and improve management.

SeMCom shall be operational by January 1997, with an expected life of at least four years, capable of capacity expansion in an incremental fashion after year two. The physical backbone infrastructure shall have an expected useful life of at least six years. Management shall be upgradeable for the life of the system.

3.12 Advanced System Planning

Design reviews will be conducted at discrete points in the program where there is a transition in emphasis and type of activity. They will assess the current status, determine ability to proceed into the next phase of activity, and modify plans as appropriate. Design reviews will be conducted to evaluate those requirements allocated or impacting the system to ensure its overall integration and interoperability. Design reviews are conducted to achieve the following technical and management objectives:

- To ensure that requirements have been properly allocated
- To evaluate the design approach and identify defects, problems, and risks
- To ensure effective follow-up of tasks and ensure adherence to schedule
- To evaluate consistency in requirements and design
- To enhance communication between organizations

3.12.1 Project Management Plan and Policies

A conceptual system design shall be developed that provides a baseline architecture for testing and that accommodates all required functionality of the system. An initial design review shall be held. If the outcome is favorable, a simulation will be developed for a proof of concept to find constraints and limitations. Concerns will be documented and addressed in the preliminary design review (PDR). If the PDR is passed, a prototype will be developed. Limitations and concerns noted during the prototype demonstration will be documented and addressed at the critical design review (CDR). Critical, show-stopping design issues shall be reviewed at the CDR and the decision to proceed with full system implementation shall be made.

To accomplish formal acceptance, a preliminary test plan shall be developed and a review shall be conducted after the completion of the simulation. A final test plan and review will be conducted. Initial validation, verification, and acceptance (VV&A) testing shall be conducted after initial system delivery and system acceptance. Final VV&A testing shall be conducted after final system delivery; this is the formal mechanism by which SeMCom's full operational capability is accepted.

An engineering review board (ERB) shall be established in support of and convened at the request of the project manager. It shall consist of senior technical personnel from the engineering organization to provide technical oversight. The ERB will convene as necessary to provide recommendations on technical changes and invite other interested parties to attend so as to contribute to the appropriate planning phases.

A project board shall be established to which the program manager shall report. Its members will be made up of the business department heads and the CIO. The project board shall monitor the progress of the effort, approve any major changes to budget or schedule, set the project implementation priorities, and ensure that all long term plans that could affect the project schedule are communicated throughout the organization.

3.12.2 Project Management Reviews

A project management review (PMR) shall convene at the request of senior management. The PMR approves any high level project decisions or appropriations. It provides incremental and full authorizations to proceed to a fully operational system.

This project accomplishes the project initiation design review and addresses the conceptual, preliminary, and detail design reviews. Guidelines for the management plan, operational concept, systems requirements document, and the configuration management document are presented following the listing of project reviews.

3.12.2.1 SeMCom Life-Cycle Project Review Listing

Project reviews shall be held in accordance with the program management plan, or when the project board issues a formal request. The scheduled project reviews and the requirements to meet the reviews are listed below.

1. Program Plan Review

The program plan review shall review the initial system acquisition plan and strategy to ensure the correct definition of system concepts, the initial requirements statement, management support staff, and the initial trade study.

2. Project Initiation Review

The project initiation review examines the proposed leadership approach. The project board is briefed on activities to approve the schedules, resource requirements, and critical path. From the critical path, risk assessments are scheduled. This review identifies the management and technical team and modifies the work breakdown structure and budget allocations. Additionally, it establishes the risk management approach by consensus and develops subsequent control systems to establish the verification and validation of the systems engineering process. The initiation review also ensures that the development of project work agreements and service level agreements has begun.

3. System Requirements Review

The system requirements review examines user requirements surveys and operational environment studies. From this, it creates is own requirements models and compares them to the actual design through analysis and derivation. The concept selection criteria are taken from the system requirements document to validate any trade studies that were performed.

4. Preliminary Design Review

The preliminary design review takes the system level performance specifications and compares them to the requirements traceability matrix to ensure a robust system architecture definition. Additional time is spent validating the risk assessments based on simulations and prototypes. The review also approves the project schedule, verification plan, implementation plan, and the initial set of manuals for engineers and users.

5. Critical Design Review

To be ready for the critical design review, all final design discrepancies must be closed and final design documentation developed in full draft form. If any item does not pass this review, it cannot be included in the design until it is fixed.

6. Simulation, Test, and Validation Readiness Review

The readiness review ensures that quality assurance methods have been practiced by examining discrepancies, action items, and their proposed fixes. To be prepared for the readiness review, the tests and simulations must be defined completely with the inputs, expected results, acceptable results, actual results, methods, and summaries of failures and solutions. Also during the review, justification for the design must be given and proved through the requirements traceability matrix. Validation procedures verify hardware, simulation, and test configurations. Any risk, feasibility, or trade-off analyses must be documented and presented. Supplementary materials for the review include the implementation plans, procedures, and configurations.

7. Factory Acceptance Review

For factory acceptance, all requirements must be met, verified, and validated. All demonstration, inspection, simulation, and test results must be documented.

8. Deployment and Operational Readiness Review

In order to begin the readiness review, all discrepancies since system verification must be resolved and documented. It is imperative to deliver all evidence that the system is capable of being deployed and operated as specified. All operation, and maintenance procedures must be defined in their entirety.

9. Final Implementation Test and Review

The final implementation test and review requires all operational and maintenance procedures to be complete and documented. Demonstrations shall be given to show compliance with the requirements and specifications. All test results, corrective actions, and discrepancy reports shall be delivered in documented format. If everything is complete, the project board accepts it and the system is approved with a final signoff.

3.12.2.2 Project Management Plan

SeMCom development will be managed according to a project management plan. It shall be prepared by the project manager during the conceptual design phase and used to define the management environment throughout the project life-cycle.

The project management plan is used to specify the project life-cycle and associated reviews. Management organization, policies, and procedures to be used on the project are specified in this document. The main function of the document is to ensure that the management aspects are used to coordinate the interrelated activities of various organizations providing support to the project. The baseline of this document can be made as early as the preliminary requirements review and updated as required.

3.12.2.2.1 WBS

Reporting to the program manager shall be any necessary administrative personnel and the two senior systems engineers responsible for the technical aspect of the project. The two systems engineers shall have twelve staff members reporting to them that include network engineers, network managers, configuration managers, and hardware technicians. Network engineers and managers shall be responsible for the research, design, planning, testing, and documentation of their portions of the system. Technicians shall help wherever possible, particularly during component configuration and integration. Configuration managers shall assist these two groups, along with the systems engineers. The project manager will be responsible for reporting progress to the CIO and will communicate between all parties involved in the planning and implementation exercise. Upon completion of the final phase of the implementation plan, the project board will state that the project is completed to their satisfaction. The work breakdown structure, allocated by the project manager, is shown in Figure 9.

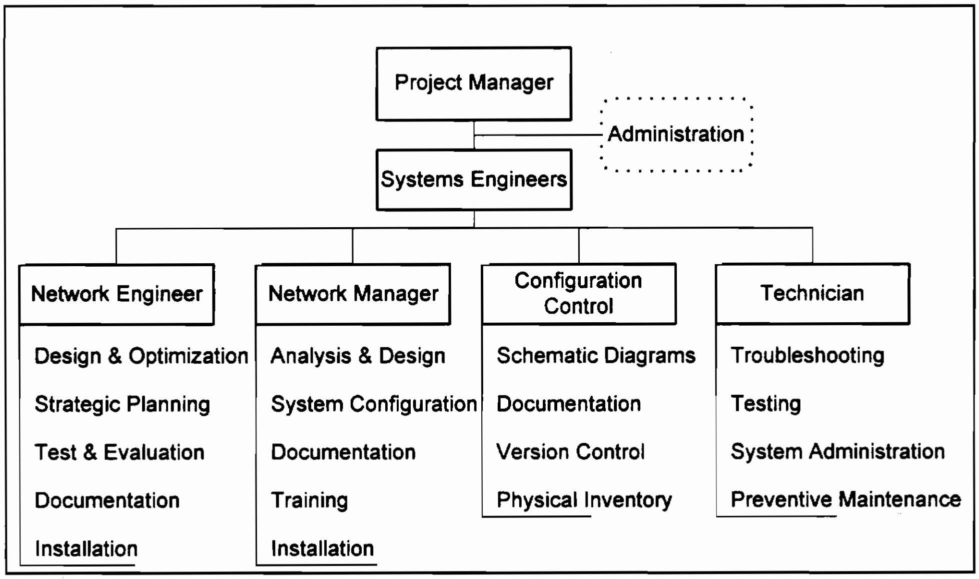

Figure 9: High-Level WBS

3.12.2.2.2 Schedule

Schedules are to be considered provisional, dependent upon supplier delivery commitments and change control procedures. Wherever possible, provision has been made for the effects of known relocation exercises and other planned activities.

A structured project implementation methodology shall be employed. The project team will report to the systems engineers and the project board. No changes to the baseline plan will be allowed without agreement from both the systems engineers and the ERB. Employing a systems engineering methodology, the project team will implement the upgrade design in complete understanding of the business unit's activities. It is the aim of the project team to achieve total transparency in the upgrade process.

Responsibility for delivery of the project will rest with the designated project manager. However, the highly technical nature of the project and the criticality of the business units to be affected means that overall success depends heavily upon the systems engineers and the cooperation of the involved parties. The project manager will coordinate all aspects of the project delivery using network engineers, business liaison personnel, vendors, and appropriate resources from the network operations center.

3.12.2.2.3 Support Documents

It is recognized that accurate documentation during design and at handover will help to ensure that the project is implemented successfully and that the value of the investment is realized over the life-cycle of the system. The project will be fully proceduralized with the production of various key working papers. These shall include:

Concept of Operations
Configuration Management Plan and Document Control Procedures
Design Specification
Product and Service Evaluation
Interface Control Document
Project Management Plan
Operator and User Manuals

Minutes of Meetings

Project Proposal and Initiation Request

Systems Requirements Document

Requests for Proposal for Support Tasks

Requirements Traceability and Verification Matrix

Master Schedule

Final Acceptance and Signoff

Service Level Agreement Statements

Test Plans, Procedures, and Reports

Training Plan

3.12.2.3 Concept of Operations Document

The concept of operations document is used to describe major operational activities, sequences of events, timelines, and interfaces. It describes the planned environment that will exist when the project transits to the operational mode. The baseline concept of operations is formed at the system requirements review and is maintained throughout the life-cycle. It helps the audience – the users, the customers, and the support personnel -- understand what the developers envision the system doing when final operational capability is achieved.

3.12.2.4 Configuration Management Plan

Configuration management is the process of ensuring the integrity and traceability of products and services. This is done by identifying the configuration of a document of a component at discrete intervals and controlling all subsequent changes. A product assurance program is the planned process of establishing requirements, standards, and procedures to ensure quality, reliability, and maintainability. Functional aspects of product assurance are configuration management, quality control, and verification and validation. The project manager is the manager for configuration control. Specific configuration management duties and processes that apply include version control, hardware baseline control, requirements traceability, change control, and document control. A baseline of the configuration management plan can be set as early as the preliminary requirements review and maintained through system implementation.

3.12.2.5 System Requirements Document

The system requirements document specifies the complete and detailed requirements of all system segments and components. Its purpose is to provide a definition of the products, operational aspects, and all requirements imposed on the project, including functional requirements and design standards or manufacturing practices to be followed in developing the system. The system requirements document is developed during conceptual design, presented at the system requirements review by the project manager, briefed to the customer, and put under formal configuration control. It is maintained through the factory acceptance test.

3.12.3 System Simulation and Test Plan

System test and evaluation measures should be developed early in the design and refined throughout the development process. Early planning helps to ensure more exhaustive testing procedures and subsequently a better design.

Simulation and testing is broken down into three areas: research and development, testing and analysis, and documentation. Research and development involves determining proper data parameters, setting up hardware, software, and workstations. Testing and analysis involves performing actual tests and mathematical computations to define expected bandwidth requirements per baseline millisecond response times. Documentation involves writing a white paper on the analysis. Standard test forms will be created to facilitate documentation.

Each test shall include the following categories, as appropriate to the specific subsystem or component. First, power supply test. Second, test transmit and receive capability. Third, test to see if system or component is manageable. Fourth, test for synchronization.

3.12.3.1 Training Plan

The training plan outlines guidelines for the instruction provided to the users and operators to ensure that the system is used effectively. Providing true end-to-end service in any network environment requires technical capability in the personnel entrusted with the responsibility to perform the task. A well-trained team of staff members oriented towards each specific area of the domain is a prerequisite for the successful delivery of quality service. To make the training efforts specific and manageable, the training approach needs to target broadly defined user communities. The training plan will be used in testing and will be finalized before the system is on-line.

4. Preliminary Design

The preliminary design uses the technical baseline as defined in the conceptual design. Subsequently, the refinement of requirements and operational concepts is accomplished through requirement allocation and functional analyses. The allocation process assigns and verifies requirement specifications to system components and begins the search for specific component solutions. Subsystem solutions are derived during system design when options are evaluated and compared against performance measures.

4.1 Detailed System Requirements

The business units are active participants in the requirements development process. They provide subject area (domain) experts to review the specifications and models and clarify requirements to ensure that SMC's operational needs are satisfied by the SeMCom design architecture.

4.1.1 Availability

Availability refers to the percentage of time that the system is available for usage. Although engineered to operate continuously, availability can be affected by component failures and environmental conditions such as power outages. For campus connectivity, a service level agreement shall allow the system to be engineered with sufficient power protection and redundancy to tolerate the failure of one or more components and continue to operate at a level that meets requirements.

SeMCom shall be available 24 hours per day, 7 days per week.

SeMCom shall have an availability of 0.995 or higher.

4.1.2 Capacity

Capacity refers to the number of simultaneous sessions (users) that the system can support. It does not refer to utilization. On a node-by-node basis, each campus must have the capability to expand by ten percent on an annual basis.

Campus nodes shall initially support 1500 simultaneous users.

Campus nodes shall be upgradeable to 3000 simultaneous users.

4.1.3 Cost

The cost of the system refers to its total price tag, in dollars, for research, development, installation, operation, maintenance, and retirement.

The total life-cycle cost of SeMCom shall be less than $6,000,000 for a four year period.

4.1.4 Environment

The environment refers to the physical surroundings and location of the equipment, and thus the conditions in which it must operate.

System components shall operate in climate-controlled and cleanroom environments. The system shall not be expected to operate reliably if the environment reaches 95% of OSHA specified cautionary ranges. The system shall be placed within fire-suppression environments with filtered power, redundant power supplies, and separate power sources.

Redundant system components shall be raised off of the floor and physically separated by a minimum lateral distance of twenty-five feet.

Electrical wiring and cabling shall be designed to specification as required by building codes.

4.1.5 Maintainability

Maintainability refers to the amount of time that must be spent on maintenance activities, either preventive or corrective.

The mean preventive maintenance time (Mpt) shall be less than two hours.

The maximum MTTR shall be six hours.

The MTBM of any system component shall be one year.

Mean maintenance time (MMH/OH) shall not exceed 0.1% of operational time.

4.1.6 Management

Management refers to system and network management. SeMCom shall provide performance management, configuration management, accounting management, fault management, and security management.

There shall be two nodes that house the primary and secondary NMS.

The NMS shall support remote configuration and administration capabilities.

The NMS shall support fault detection for primary system components.

The NMS shall have SNMP, RMON MIB support, and a GUI.

4.1.7 Performance

Performance is defined by benchmark statistics such as delay, throughput, percentage overhead, and speed.

The maximum allowable amount of management traffic shall not increase the average packet delay by more than five percent.

SeMCom shall support transmission rates higher than 56 kbps and deliver one megabit files from node to node in less than ten seconds.

4.1.8 Reliability

Reliability is the statistical measure of availability and maintainability. It can be thought of as a level of accuracy, or how much the system complies with its specifications.

The system shall operate at least 99.5% of non-scheduled maintenance time.

The MTBF of the system shall be 25,000 hours.

4.1.9 Retirement

Retirement is the last phase of the system life-cycle. When a piece of equipment has reached the end of its life-cycle, it shall be removed.

System components shall be donated to charitable organizations, educational institutions, or sold as used equipment after their utilization period.

4.1.10 Schedule

The schedule is the amount of time that it takes to design and rollout the system, from project initiation request to final signoff and acceptance.

The implementation schedule shall be less than eight weeks (forty working days).

The actual system cutover time (green zone) shall be less than twelve hours.

4.1.11 Security

Security refers to the physical and logical access controls that are placed on the system. Although security is typically handled through the facilities and physical plant organizations, basic security must be put in place on the portions of the system that come into contact with the "public" domain.

Physical access controls to the campus data centers shall be implemented.

Each campus data center and facilities management crew shall be responsible for implementing physical security as dictated by corporate headquarters.

Logical access controls to files and directories shall be implemented via authentication systems.

Each campus data center shall be responsible for implementing logical security standards as dictated by corporate headquarters.

Encryption mechanisms shall be employed for all confidential material.

4.1.12 Utilization

Utilization is a measure of the actual usage of the system in terms of the percentage of total bandwidth. Like capacity, utilization must be capable of expanding by ten percent per year.

SeMCom shall be capable of supporting utilization demands as presented in Figure 5 and Figure 6.

The system shall be designed to accommodate two hundred percent (200) of the projected average utilization.

SeMCom shall sustain peak utilization periods at eight percent (80) capacity.

SeMCom shall accept bursts of one hundred fifty percent (150) of the projected utilization for one hour.

Management traffic shall not increase packet delay by more than five percent.

4.1.13 System Simulation, Test, and Evaluation

Simulation and test procedures shall define appropriate line speeds, verify connectivity, and demonstrate robustness and reliability. The goal of simulation is to ensure that appropriate bandwidth is allocated for various configurations, correct addressing is assigned, and all components are functional and interoperable. These procedures exercise the interaction of clients using link bandwidths of 56 kbps to 1.544 Mbps to find peak and steady-state throughput and response times.

Completion will yield three important results. First, the operation of the indicated number of clients over the link running at speeds from 56 kbps to 1.544 Mbps will be verified. This will empirically indicate whether a link of that bandwidth is adequate for the indicated number of clients. Second, the collection of data may be used to analytically calculate the required data rate from the indicated number of workstations. Third, logical connectivity and operability will be verified. All results will be presented in a summary document.

For WAN simulation, ComNet III from CACI Software was chosen. ComNet III allows for the simulation of many network architectures and protocols and can run on a desktop PC with either Windows NT or Windows 95. It also provides some amount of traceability so that one may see how the traffic flow is distributed throughout the network. Results are mathematically verified through analysis.

4.1.14 Summary

The detailed system requirements can be broken down into individual indices that measure performance. Table 2 lists the technical performance measures (TPM).

Table 2: SeMCom Technical Performance Measures

TPM	Requirement	Benchmark	Weight
Utilization (Mbps)	200*Average Demand	0.8*Average Demand	13
Availability	0.995	0.8929	12
Packet Loss (%)	< 0.01	25	11
Reliability (%)	99.5	75	10
Management Station	2	0	9
File Transfer (Minutes)	<= 0.2	<= 240	9
Added Delay (%)	<=5	N/A	8
MTBF (Hours)	> 25,000	100	7
Redundancy	Fault Tolerant	None	6
MTTR (Hours)	< 6	12	5
Users	< 3000	< 1200	4
Mpt (Hours)	< 2	5	3
MMH/OH (%)	<= 0.1	0.9	2
Security	Encryption	Clear text	1
			100

4.2 Functional Analysis

Functional analysis takes the statement of need and transforms it into a working product or service. It facilitates the design process through the use of operations and maintenance concepts. Flow diagrams map requirements to functional segments and components by defining requirements, constraints, indices of performance, and life-cycle.

SeMCom operational and management functional flow diagrams to three levels are shown in Figures 10 through 12. Figure 10 shows the top-level of SeMCom's systems engineering process including conception, implementation, and retirement. Figure 11 describes the main function, operation, with its components of routing and forwarding detailed to another level of granularity. Functional analysis and section 3.9 are used to classify the different components of the system into packages.

Packaging the system stems from the functional analysis and breaks the system into its components based upon geography, similar equipment, or a common environment. The intent is to have high independence and cohesiveness to simplify the model. It also minimizes interaction to help increase understanding of the system, improve the model's representation of the real world, and to facilitate changes. Thus, packaging is critical for reliability, modifiability, and maintainability.

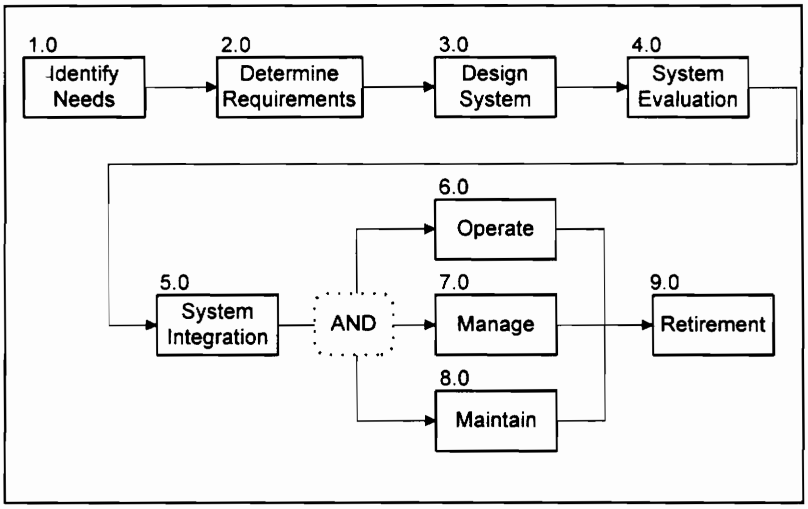

Figure 10: System Functional Flow

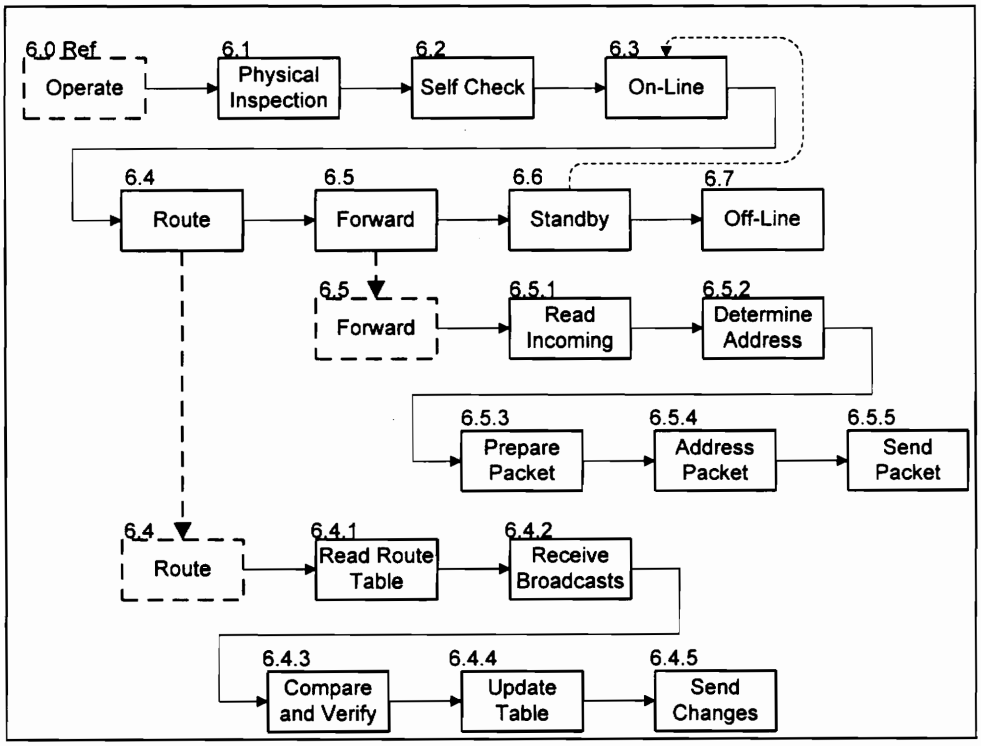

Figure 11: Operational Functional Flow

In Figure 12, the management portion is described in full. Notice that the system design is object-based so that the second-level functional flow that includes maintenance actions can be transferred to the third-level, with the same definition.

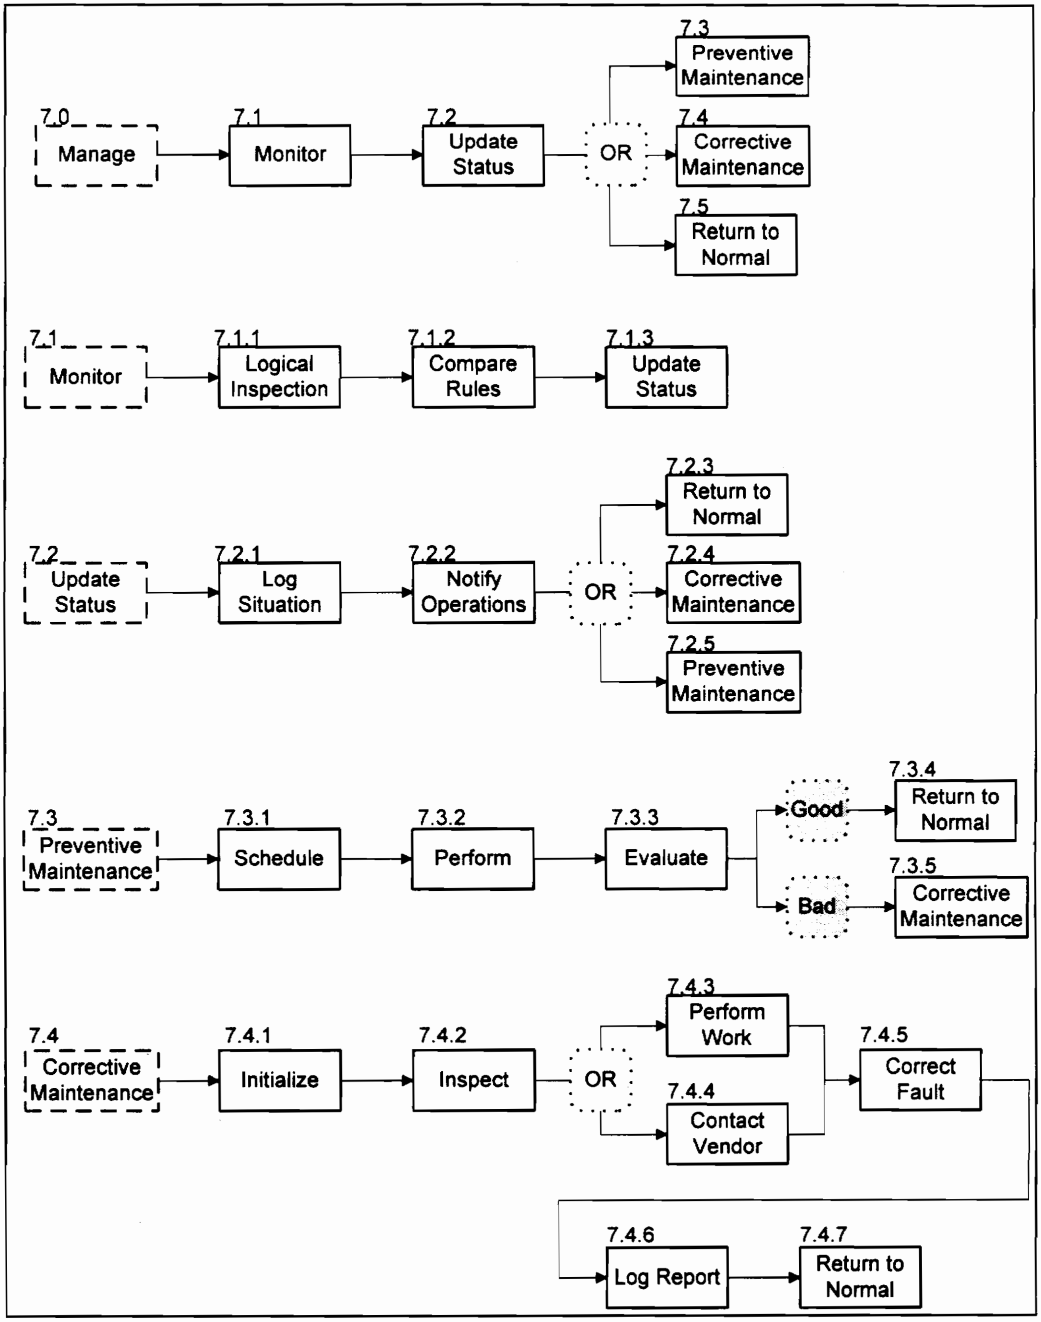

Figure 12: Management Functional Flow

4.3 Requirements Allocation

With the aid of functional analysis, specific requirements can be assigned to the different system component areas. As the requirements are allocated, it can be determined if the logic and solution are viable. Figure 13 shows requirements allocated to key components. Continued analysis would further refine package items, such as personnel and activities underneath the Management Operations category.

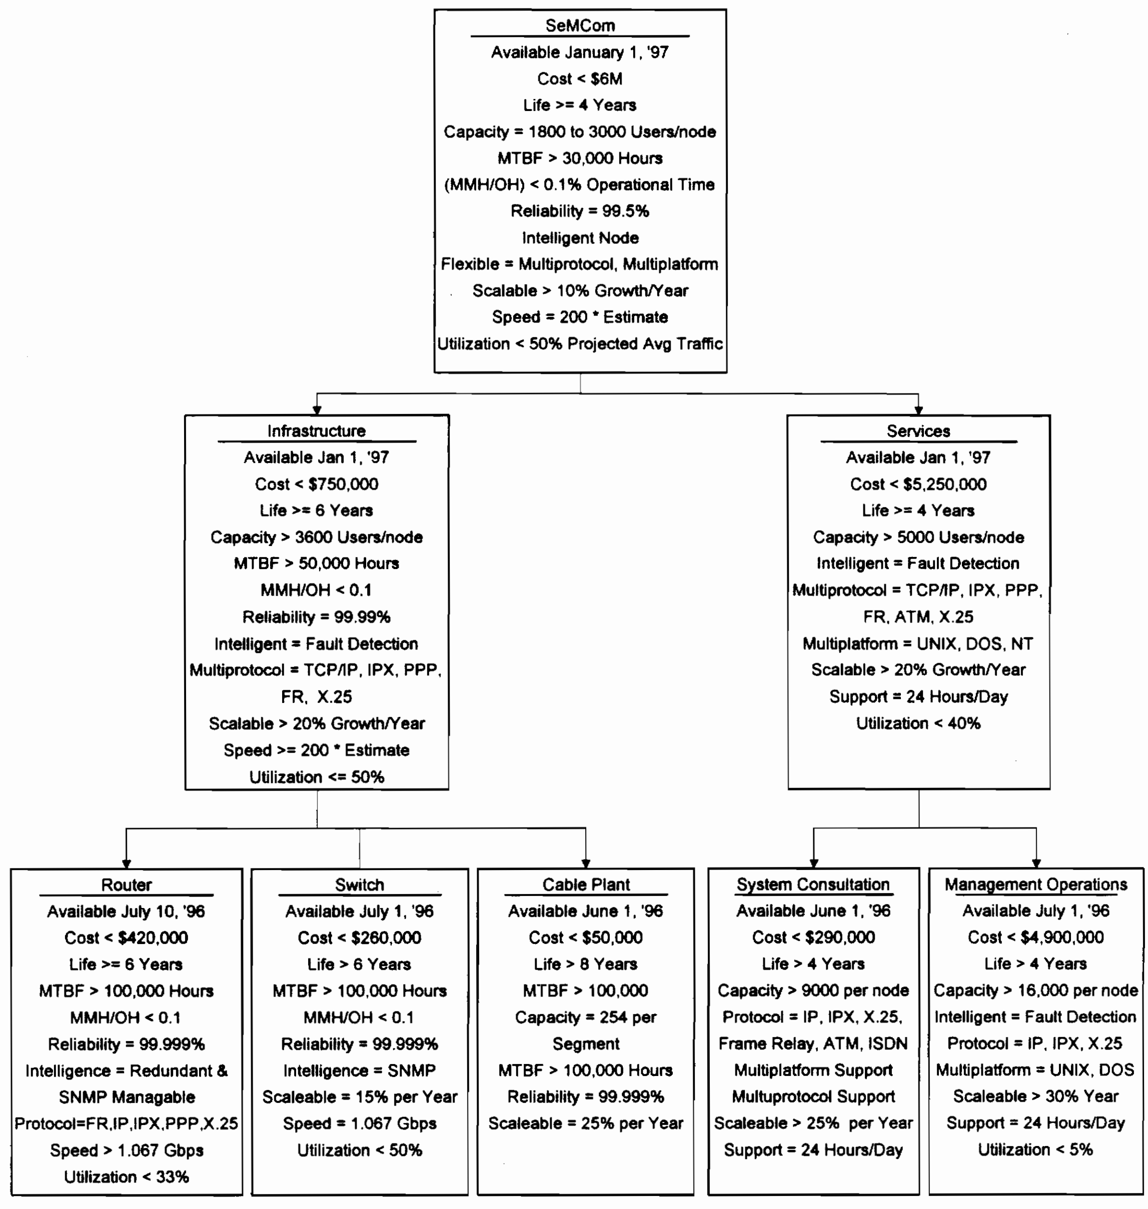

Figure 13: High-Level Requirements Allocation

4.4 Evaluation of Subsystem Alternatives

Based upon the allocation of system requirements, specific candidate solutions can be researched and evaluated. Items such as cost, benchmark performance statistics, and maintainability are used as indices to compare the candidates.

4.4.1 Routers

The two routers that rose to the top of their class were the Cisco 7500 series and the Wellfleet family.

The serial interfaces for the Wellfleet BCN, part of their family of products, can provide the performance capabilities for the mission-critical wide area network. They operate at speeds up to fifty-two Mbps, supporting leased line, dial-up, and a broad range WAN services. Numerous protocols can be supported, including PPP, frame relay, X.25, and ATM. The routers provide comprehensive traffic management capabilities through multicircuit support, traffic filters, prioritization, and compression. Each interface can be configured through the node management application, Site Manager. Its configuration editor provides a graphical interface that streamlines the configuration process. Additionally, interfaces can be added to any other Wellfleet router to allow for expansion and to mix and match technologies to satisfy varied network requirements.

The Cisco 7507 router is part of the 7500 series family. It contains seven -slots that can support multiprotocol, multimedia routing and bridging with a wide variety of protocols and any combination of ATM, Ethernet, Fast Ethernet, token ring, FDDI, serial, High-Speed Serial Interface (HSSI), frame relay, and channel attachment. The router backplane can operate at speeds exceeding one Gbps, with serial and ATM interfaces at over 622 Mbps. Network interfaces reside on interface processors that provide a direct connection between the Cisco extended buses and external networks. The 7507 has interface processor capability in slots zero and one, Route Switch Processor capability (RSP) in slots two and three, and more interface processor slots in four through six. There are bays for up to two AC-input or DC-input power supplies, but the chassis needs only one power supply to operate. Although a second power supply is not required, it allows load sharing and increased system availability. A comparison evaluation using industry-supplied parameters is shown in Table 6.[11]

Table 6: Comparison of Bay Networks and Cisco routers

Criteria	Cisco Systems	Bay Networks
Available	Yes	Yes
Compatible	Yes	Yes
Latency	< 1.39 E-05	< 1.41 E-05
Maintainability	Very High	Very High
Market Share	75%	16%
Processing Speed	1.067 Gbps	1.067 Gbps
Protocols	FDDI, HSSI, Serial, ATM, X.25, T1, ISDN PRI, IP, IPX	X.25, ATM, SMDS, PPP, ISDN
Redundancy	Unit Level	Card Level
Reliability	99.999%	99.999%
Reputation	Outstanding	Excellent
Risks	New technologies forthcoming	Dominated by Cisco
Scaleable	To five interfaces	To four interfaces
Support	Extra Cost; Good	Extra Cost, Very Good

4.4.2 Switches

The two switches that rose to the top of their class were the Cisco Catalyst 5000 and the Bay Networks 28115.

The Bay Networks 28115 provides sixteen available ports that support both ten and 100Mbps 10Base-T and 100Base-TX standards in full and half duplex modes. Two Media Independent Interface (MII) standard expansion ports provide interconnection of multiple switches or can serve as downlink ports using an external 100Base-TX or 100Base-FX standard transceiver that provides cable to fiber media conversion. All ports can be individually configured for speed, duplex support, and desired segment assignment. Individual connections or collision domains can also be mirrored to any port for diagnostic or monitoring purposes using a Sniffer. A console port is provided for direct switch configuration or remote configuration and management can be provided via HP OpenView or Bay Networks Optivity software.

The Cisco Catalyst 5000 has five slots. Its first slot is used for the supervisor engine that provides layer two switching and local and remote network management. The supervisor engine has dual, full duplex fast Ethernet interfaces that enable connections to other Catalyst 5000, routers, and servers at up to 400 Mbps. The remaining four slots can be used for any combination of switching modules. Only one power supply is required to run a fully configured system. Options include dual, hot-swappable power supplies and modules to deliver fault tolerance and repair flexibility. The Catalyst 5000 supports up to ninety-eight switched 10-Mbps Ethernet interface ports, or up to fifty switched 100-Mbps Fast Ethernet interface ports. Additional advanced features include support for three priority queues on the backplane allowing users to configure a higher or lower priority level on a switched interface to accommodate bursty traffic; mirroring of any port for diagnostic or monitoring purposes, and support for future virtual LAN applications. A comparative evaluation is shown in Table 7, with a comparison of management options shown in Table 8.

Table 7: Comparison of Cisco and Bay Networks Switches

Criteria	Cisco Systems	Bay Networks
Available	Yes	Yes
Buffer Type	Individual	Shared
Compatible	Yes	Yes
Latency	< 1.02 E-05	< 1.02 E -05
Maintainability	Very High	Very High
Market Share	34%	20%
Processing Speed	1 Gbps	1 Gbps
Protocols	FDDI, Ethernet, TR	FDDI, Ethernet, TR
Reliability	99.99%	99.99%
Reputation	Excellent	Excellent
Risks	Layer 3 Switching Unstable	Smaller end switch
Scaleable	To 98 Ports	To 16 Ports
Service Support	Good	Very Good
Virtual LAN	Yes	Yes

Table 8: Switch Management Features

Network Management Feature	Catalyst 5000	Bay Networks 28115
SNMP agent v1 (RFC 1155-1157)	Yes	Yes
SNMP MIB II (RFC 1213)	Yes	Yes
Ethernet MIB (RFC 1398)	Yes	Yes
Interface Table (RFC 1573)	Yes	Yes
Bridge MIB (1493)	Yes	Yes
RMON v1 (RFC)	No	Yes
RMON v2 (RFC)	No	Yes
FDDI MIB (RFC 1512)	Yes	N/A
SMT 7.3 (RFC 1285)	Yes	N/A
AToMIC MIB (RFC 1695)	Yes	N/A
ILMI MIB (ATM Forum UNI 3.0)	Yes	N/A
Cisco Discovery Protocol	Yes	N/A

Both the Catalyst 5000 and Bay Networks 28115 support local and remote configuration, reload, and restart management functions via either direct console attachment or separately purchased vendor management software. Only the Bay Networks switch utilizes Optivity management software; the Catalyst 5000 utilizes CiscoWorks and CiscoView management and configuration software. The recommendation is to use the Cisco Catalyst 5000.

4.4.3 ATM

Hardware prices have dropped approximately fifty percent since initial introductions and continue to decrease. The cost of ATM switches will go down at much faster rates as production increases and more vendors enter the market. Table 9 shows a comparison evaluation of two ATM products from GTE and Fore Systems.

Due to the developing nature of ATM, there is no design recommendation at this time. However, the Fore switch would be recommended if ATM were to be implemented at this time. The market should be re-evaluated in another year to determine if an ATM solution would be best at the time of the year two upgrade.

Table 9: Selected WAN ATM Switch Products

Criteria	GTE	Fore
Available	Yes	Yes
Compatible	Yes	Yes
Latency	< 1.1E-05	< 1.1E-05
Maintainability	Difficult	Difficult
Market Share	4%	22%
Processing Speed	5 to 10 Gbps	2.5 to 10 Gbps
Protocols	ATM, Frame Relay, SMDS,	ATM, Frame Relay, SMDS, Serial
Reliability	99.999%	99.999%
Reputation	Very Good	Excellent
Risks	Standards and buy-in	Standards and buy-in
Scaleable	To sixty-four ports	To ninety-six ports
Service Support	Very Good	Very Good

4.4.4 WAN Protocol

Migrating from X.25 network to a multiprotocol network is not terribly complicated, but it does take some amount of planning to ensure for improved and more reliable communications. The following subsections show how the internal workings of X.25 and TCP/IP to explain why the migration will help solve some of SMC's problems.

4.4.4.1 Traffic Simulation

The traffic simulation and analysis for network protocols involves determining the optimal packet sizes while minimizing overhead and packet delay, and maximizing throughput across a channel. Results presented in this project showed the analysis of X.25, TCP/IP, X.25 encapsulated in IP, and TCP/IP encapsulated in X.25 traffic. Some assumptions were made to simplify the analysis and to protect the true identity of SMC's network. First, all given and determined values are averages. Second, minimum header sizes for network protocols are used. And third, data placed into the X.25 and TCP data fields is pure data and contains no additional upper layer overhead. This directly affects the effective throughput.

The actual numbers for loads and percentages on communications links, throughput, and geographic distances have been changed from the original findings to properly disguise SMC's real network architecture. None of the changes assist or hinder the final results in any way. The disguised information offers the same information as the actual trace files, but it does not reveal any sensitive parameters of SMC's network. Similarly, IP network numbers and addresses were also changed.

It is also important to note that in the original network, routes were static. That is, traffic from Node A to Node B always traveled across the same link, regardless of whether it was the best pathway at the time. In the new network, routes are dynamic and load-balanced to distribute the traffic evenly across the network. For the simulation, traffic followed the pattern presented below.

4.4.4.2 Traffic Load

The traffic load for the network involves source to destination traffic as shown in Table 10. This traffic is statically routed between the nodes following the algorithm shown in Table 11. Using the algorithm, the number of hops from source node to destination node are determined and shown in Table 12. The average number of hops from source to destination is 1.32. These values are used to determine the packet delay and throughput of the network.

Table 10: Traffic Load for Source to Destination in kbps.

Node	A	B	C	D	E	F
A	0	9	4	1	7	4
B	9	0	8	3	2	4
C	4	8	0	3	3	2
D	1	3	3	0	3	4
E	7	2	3	3	0	5
F	4	4	2	4	5	0

Table 11: Routing Algorithm.

Node	A	B	C	D	E	F
A		AB	ABC	ABFD	AE	AEF
B	BA		BC	BFD	BFE	BF
C	CBA	CB		CD	CE	CEF
D	DFBA	DFB	DC		DCE	DF
E	EA	EFB	EC	ECD		EF
F	FEA	FB	FEC	FD	FE	

Table 12: Number of Hops from Source to Destination.

Node	A	B	C	D	E	F
A	0	1	2	3	1	2
B	1	0	1	2	2	1
C	2	1	0	1	1	2
D	3	2	1	0	2	1
E	1	2	1	2	0	1
F	2	1	2	1	1	0

4.4.4.3

The transfer of data over X.25 networks involves placing user data into network layer packets. These packets are then placed into the information field of a LAPB data link layer packet that is sent across a physical channel with X.121 as the protocol. The overhead of X.25 and LAPB packets are twenty-four bits (three bytes) and forty-eight bits (six bytes), respectively, for a total of seventy-two bits (nine bytes) per packet.

Using this and a bit error rate on the channel of 10-7, the overhead for each data packet can be determined. For X.25, the user data field allows a maximum of 2048 bits. Table 13 and Figure 14 show the overhead percentage for various X.25 packet sizes. The optimal size is 2120 bits, with overhead of 3.42 percent for the selective reject windowing approach and 3.28 percent for the Go-Back-N windowing approach. Selective reject calls for retransmitting one frame for each bit error, and Go-Back-N calls for retransmitting N frames. In addition, the window size for LAPB should be set to seven, and the window size for X.25 should be set to at least three.

The assumption that only data packets are sent slightly affects this analysis. In reality, for sending data, there are control packets that need to be sent over the network. Involving this overhead changes the optimal data packet size desired. Additional overhead introduced for control increases the desired packet size. Since the maximum packet size is used already, additional overhead reinforces its use.

Table 13: X.25 Overhead Percentage.

Data (bits)	X25 Header (bits)	LAPB Overhead (bits)	Total Packet Size (bits)	Selective Reject (1)	Go-Back-N (4)
0	24	48	72	100.00	100.00
8	24	48	80	90.00	90.00
16	24	48	88	81.82	81.82
32	24	48	104	69.23	69.23
64	24	48	136	52.94	52.95
128	24	48	200	36.00	36.01
256	24	48	328	21.95	21.96
512	24	48	584	12.33	12.35
1024	24	48	1096	6.58	6.61
2048	24	48	2120	3.42	3.48

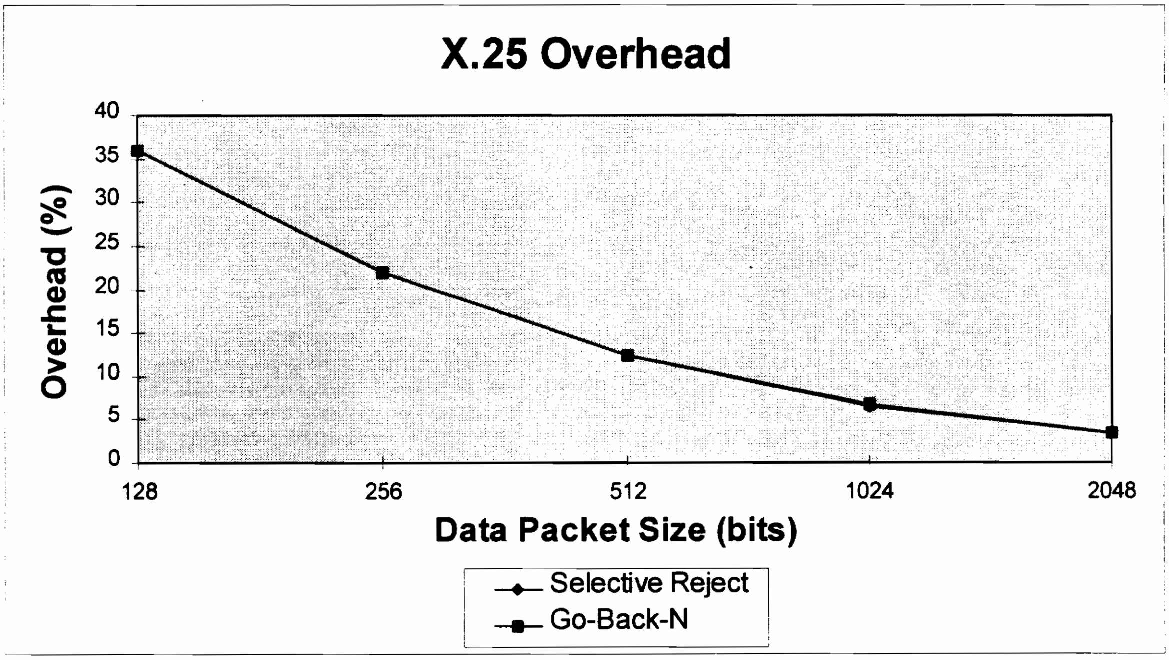

Figure 14: Overhead Percentage for X.25

Using the selective reject windowing approach and the determined overhead for various packet sizes, packet delay and throughput are determined. Table 14 shows the optimal packet size of 2048 bits gives an average packet delay of 83.70 milliseconds with an effective throughput of 23.92 kbps. These outcomes are plotted in Figure 15 and Figure 16.

Table 14: X.25 Packet Delay and Throughput.

Packet Size (bits)	Overhead Percentage (%)	Packet Delay (ms)	Effective Throughput (kbps)
136	52.94	8.60	10.04
200	36.00	10.40	13.76
328	21.95	15.00	17.52
584	12.33	24.70	20.59
1096	6.58	44.30	22.67
2120	3.42	83.70	23.92

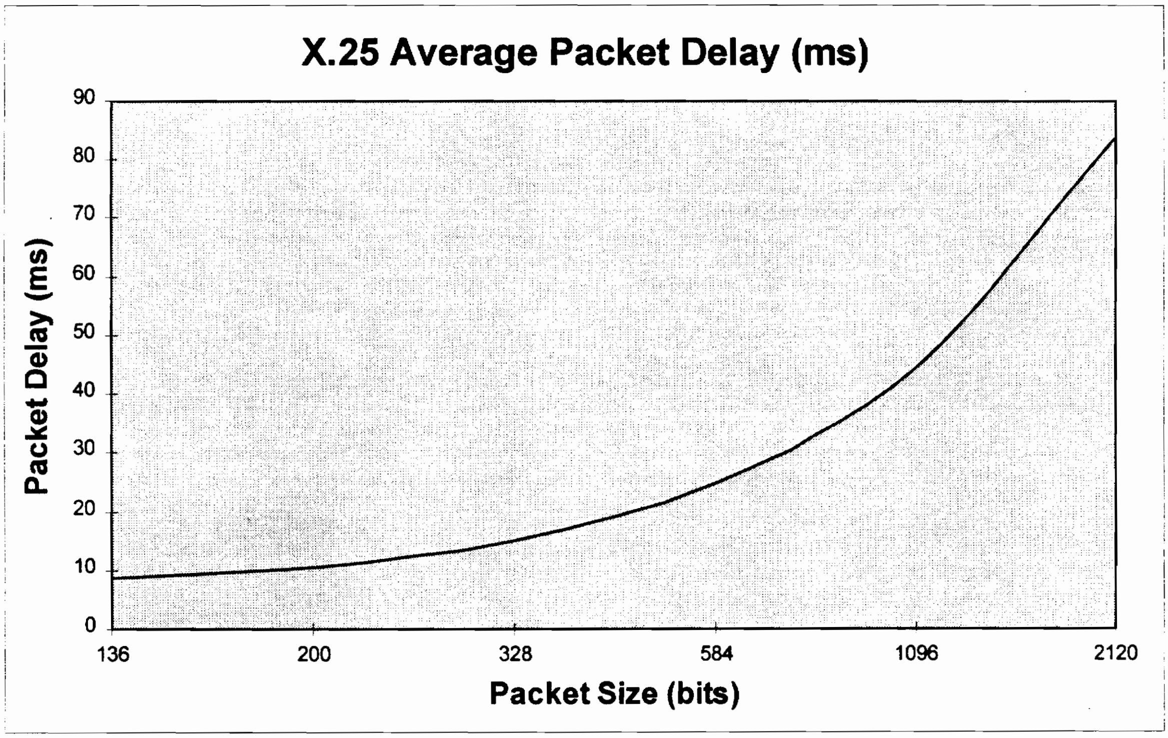

Figure 15: X.25 Packet Delay

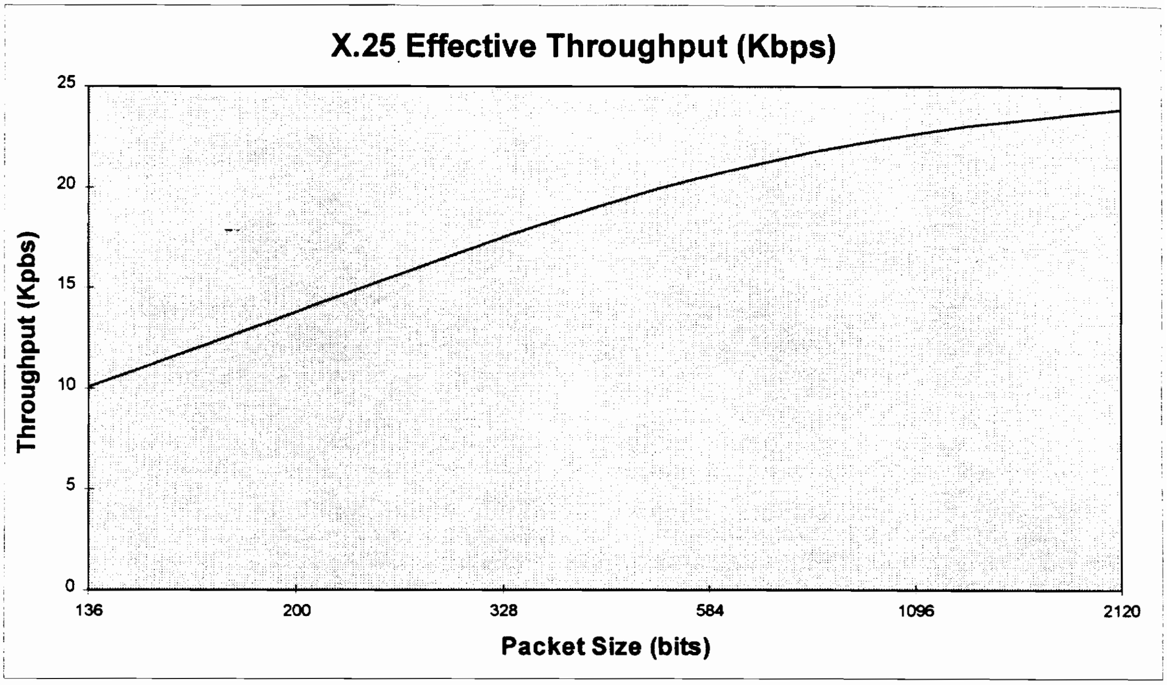

Figure 16: X.25 Throughput

4.4.4.4 TCP/IP Traffic Analysis

The transfer of data over TCP/IP networks involves placing user data into TCP transport layer packets. TCP packets are then placed into the data field of an IP network layer packet, and finally into the information field of an HDLC data link layer packet. This packet is sent across a physical channel.

TCP, IP, and HDLC data packet formats have overhead associated with them. The overhead associated with these protocols are 160 bits (twenty bytes), 160 bits (twenty bytes) and forty-eight bits (six bytes), respectively, for a 368 bits (forty-six bytes) total.

Using this and a bit error rate on the channel of 10-7, the overhead percentage for each data packet can be determined. For TCP/IP, the user data field allows a maximum of 524,288 bits. The optimal packet size is 61,440 bits for the selective reject windowing approach, which has an overhead percentage of 1.21 percent. For the Go-Back-N approach, the optimal packet size is 32,768, which has an overhead percentage of 2.42 percent. The window size for HDLC should be set to seven.

Involving control overhead changes the optimal data packet size desired. Additional small overhead packets introduced for control reduces the average packet size. To compensate for this, the desired data packet size increases, establishes a higher average size, and reduces the overhead percentage. The maximum desired packet size for TCP/IP after adding overhead is related to the types of upper layer applications involved. An overhead comparison is shown in Figure 17.

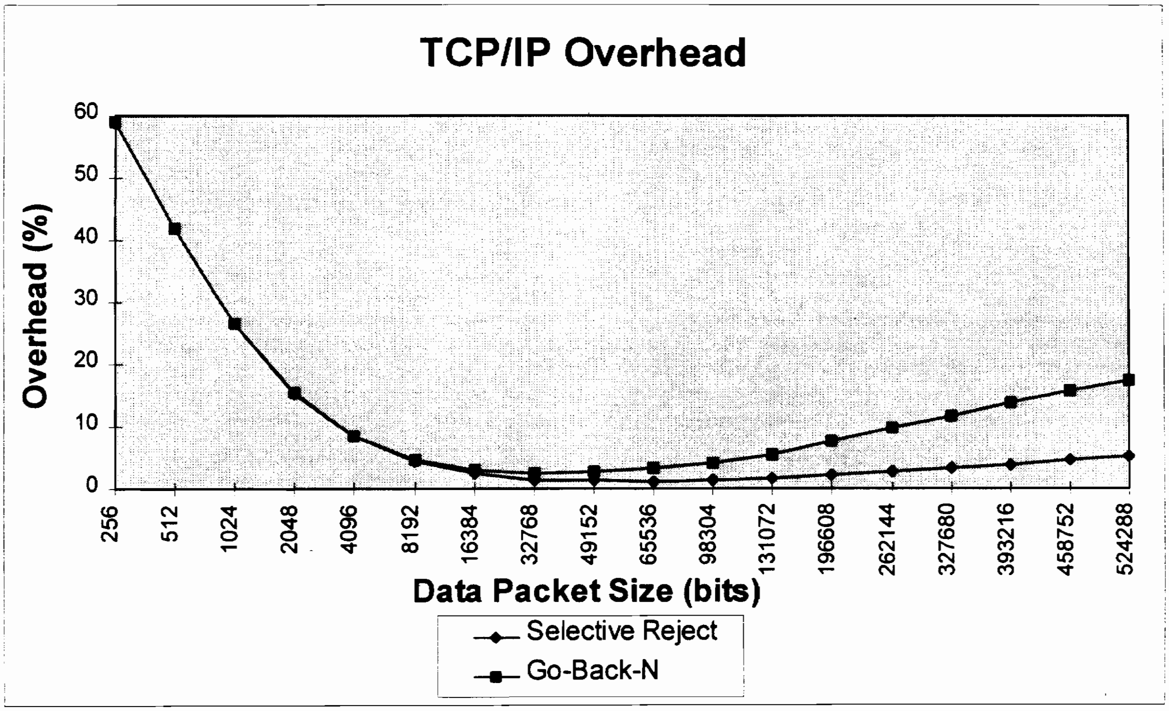

Figure 17: TCP/IP Overhead

Using the selective reject approach and the overhead determined for various packet sizes, the packet delay and throughput for TCP/IP are determined. Figure 18 and Figure 19 show the packet delay and throughput. Optimal packets of 61,440 bits correspond with an average delay of 2.40 seconds and an effective throughput of 24.83 kilobits per second.

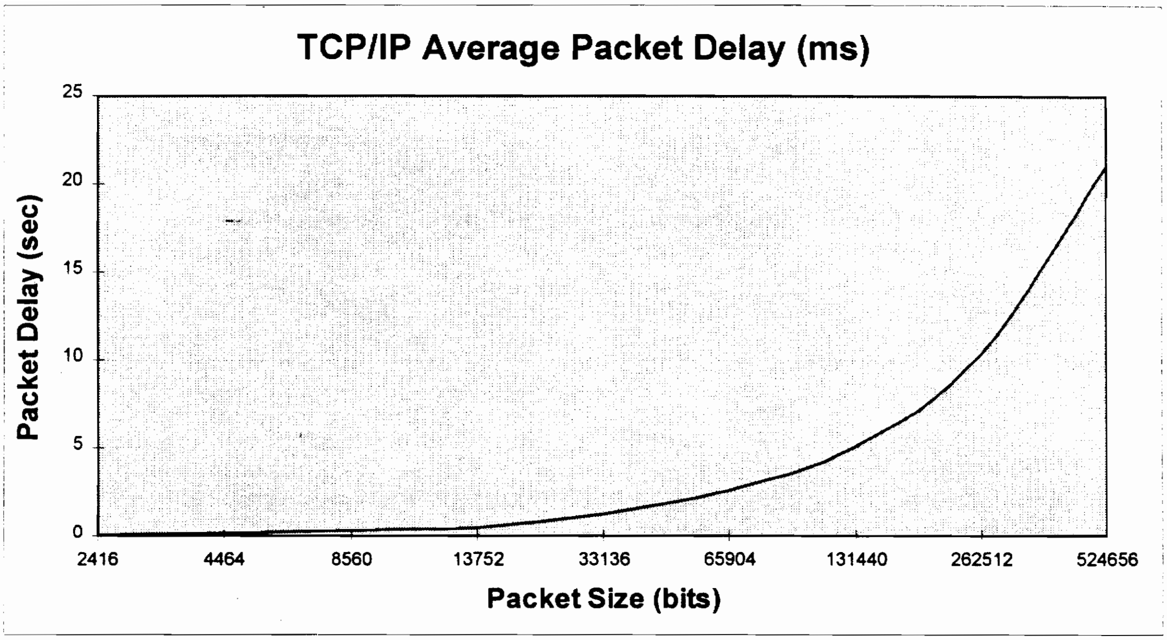

Figure 18: Packet Delay

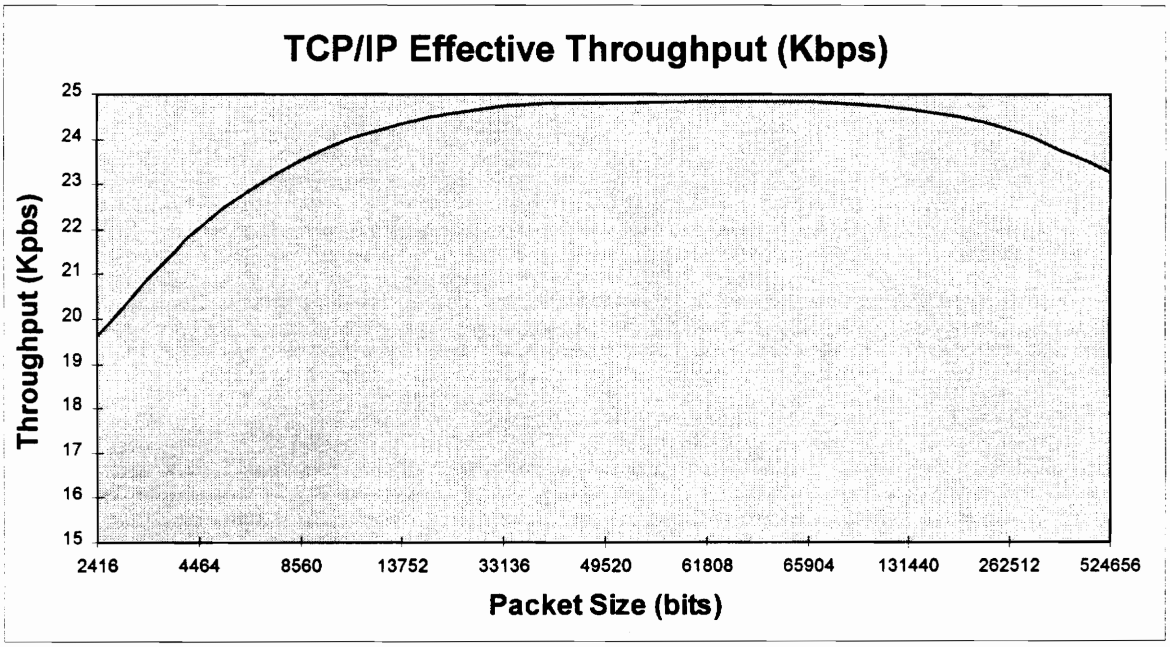

Figure 19: TCP/IP Throughput

4.4.5 Encapsulation

It is necessary to determine ways of having multiple protocols such as X.25 and TCP/IP coexist on the same network channels. Figure 20 and Figure 21 show the effects of encapsulating X.25 into IP with respect to delay and throughput.

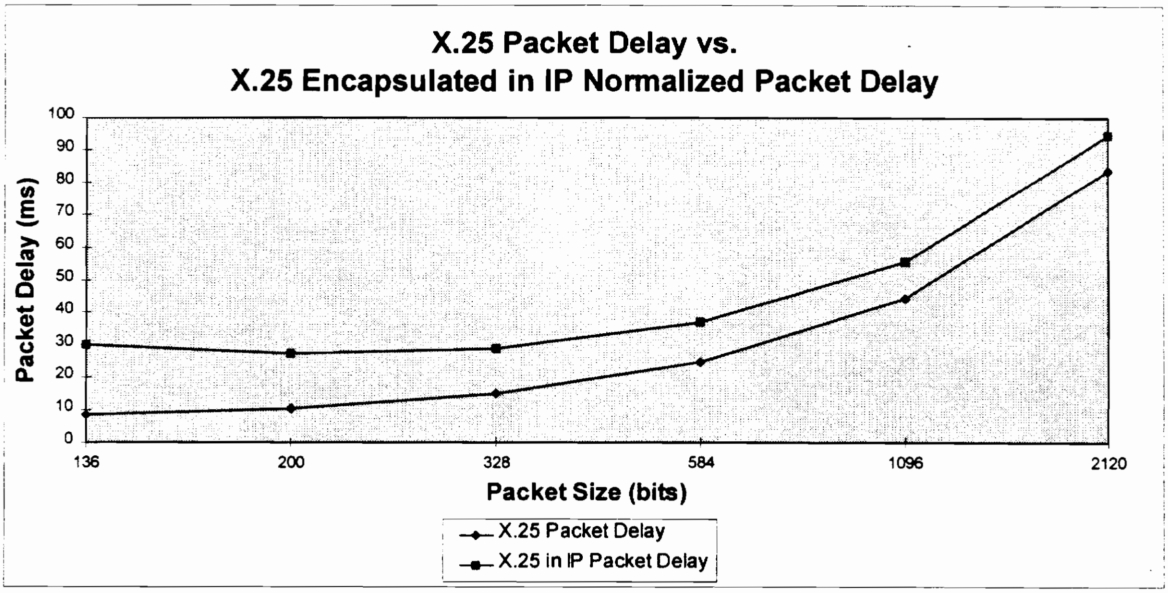

Figure 20: Encapsulation Delay

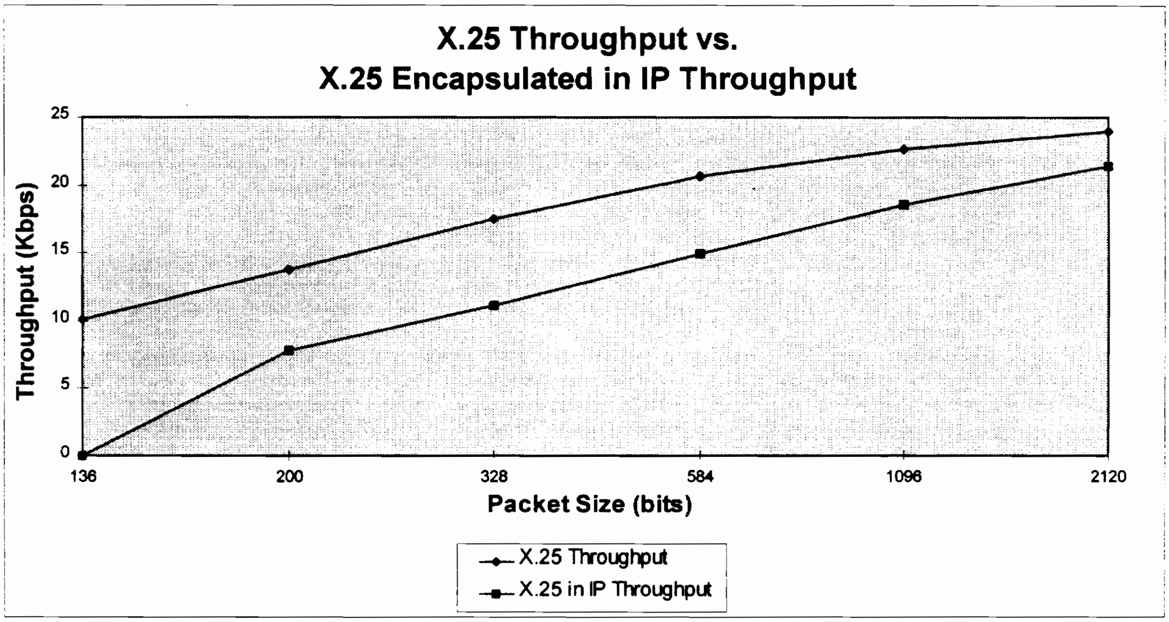

Figure 21: Throughput Comparison

The increased overhead due to encapsulation for realistic average packet sizes, about 584 bits, has significant negative impacts on packet delay and throughput, not counting processor utilization. The average packet delay increases by fifty percent and the throughput decreases by twenty-seven percent. Adding processor cycles for translational operations will increase the delay significantly, so therefore encapsulating X.25 into IP is not recommended.

Encapsulation of TCP/IP into X.25 is shown in Figure 22 and Figure 23. The impacts due to increased overhead are not as severe as with X.25 encapsulation into IP. Packet delay increases by eleven percent and the throughput decreases by five percent, for realistic average packet sizes of 33,136 bits. Therefore, when encapsulation is required, the method of encapsulated TCP/IP into X.25 is recommended. For the given network, this type of encapsulation is required when TCP/IP traffic must flow through a packet switch. However, in the final IP routed network, a multiprotocol router implementation will alleviate the need for encapsulation.

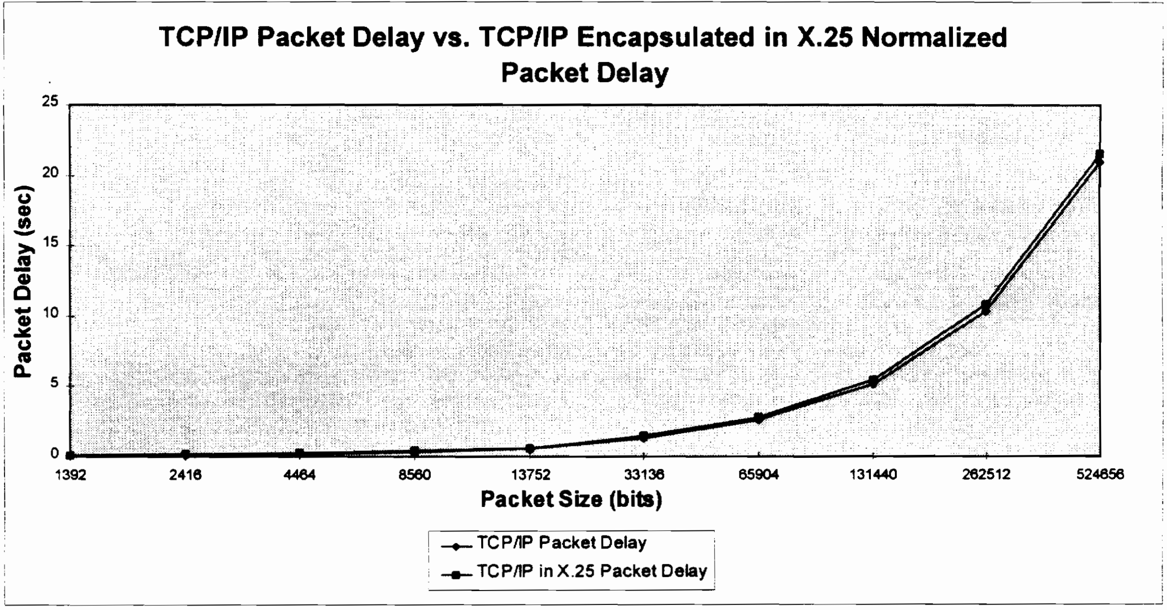

Figure 22: Encapsulated Delay Comparison

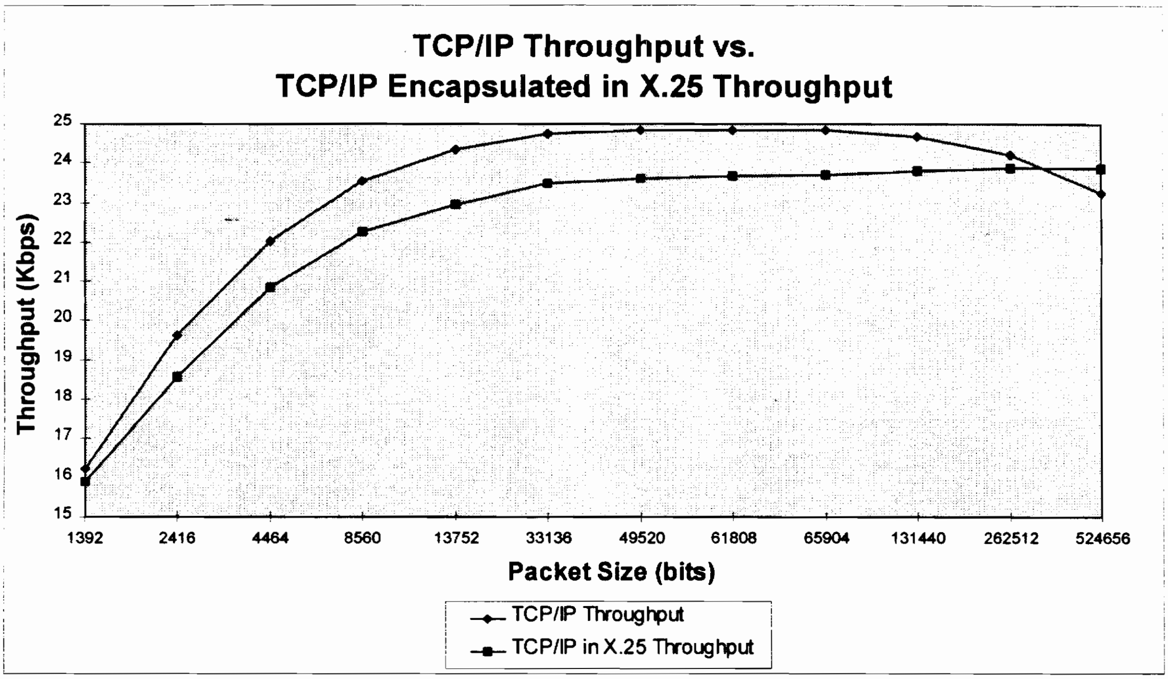

Figure 23: Throughput Encapsulation

4.4.6 "Typical" Traffic Scenario

The "typical" traffic scenario creates some realistic analysis and results for the network. In this scenario, the given traffic loads are assumed to be 200 kilobyte file transfers. It is assumed that permanent virtual circuits are already established.

The coexistence of X.25 and TCP/IP traffic is affected by this analysis. In general, the average X.25 packet is smaller than the average TCP/IP packet. With the coexistence of traffic, the average packet size for the network decreases. The affect of the coexistence is also relevant in the overhead, packet delay, and throughput, all of which will have average values between the that of X.25 and TCP/IP. For TCP/IP, the smaller average packet size has a positive impact in that the average size of retransmissions.

4.4.6.1 "Typical" Traffic Analysis

For the X.25 "typical" traffic analysis, several approximations are made for the amount of upper layer control packets needed and the size associated with each. These estimates are shown in Table 15 and were derived based on Sniffer traces of similar file transfers. The highest achievable average packet size using this type of traffic is 1093 bits with an average overhead of seven percent The average packet delay for this scenario is 86 milliseconds, and the average throughput for the network is 22.46 kbps.

Table 15: Data and Control Information for X.25 200 kB File Transfers.

Category	Number of Bits	Number of Packets
Overhead		
Session Establishment	3200	10
Session Close	640	4
Acknowledgments	64000	800
Data Packets	57600	
Data	1638400	800
Total	1763840	1614
Average Packet Size	1092.84	
Percent Overhead	7.11 %	

If interactive traffic, such as that of a Telnet session were used, results would be somewhat different. Estimates for interactive traffic are shown in Table 16 and derived based on Sniffer traces of similar sessions. The highest achievable average packet size using interactive traffic is 244 bits with an average overhead of thirty-five percent.

Table 16: Data and Control Information for X.25 Interactive Traffic.

Category	Number of Bits	Number of Packets
Overhead		
Session Establishment	3200	10
Session Close	640	4
Acknowledgments	9600	120
Data Packets	8640	
Data	40000	120
Total	62080	254
Average Packet Size	244.41	
Percent Overhead	35.57 %	

4.4.6.2 TCP/IP "Typical" Traffic Analysis

For the TCP/IP "typical" traffic analysis, several approximations are made for the amount of upper layer control packets needed and the size associated with each. These estimates are shown in Table 17 and were derived based on Sniffer traces of similar file transfers. The optimal achievable average packet size using this type of traffic is 63,478 bits with an average overhead of 1.36 percent. This calls for the data packets to have up to 273,072 bits of data in them. The average packet delay for this scenario is 2.47 seconds and the average throughput for the network is 24.78 kbps.

Table 17: Data and Control Information for TCP/IP 200 kB File Transfers.

Category	Number of Bits	Number of Packets
Overhead		
Session Establishment	5632	10
Session Close	1584	4
Acknowledgments	2880	6
Data Packets	1920	
Data	1638400	6
Total	1650416	27
Average Packet Size	63477.54	
Percent Overhead	1.36 %	

Estimates for interactive traffic are shown in Table 18 and were derived based on Sniffer traces of similar sessions. The highest achievable average packet size using interactive traffic is 564 bits with an average overhead of 72.07 percent.

Table 18: Data and Control Information for TCP/IP Interactive Traffic.

Category	Number of Bits	Number of Packets
Overhead		
Session Establishment	5632	10
Session Close	1584	4
Acknowledgments	57600	120
Data Packets	38400	
Data	40000	120
Total	143216	254
Average Packet Size	563.84	
Percent Overhead	72.07 %	

4.4.6.3 TCP/IP Encapsulated into X.25 "Typical" Traffic Analysis

For the TCP/IP encapsulated into X.25 "typical" traffic analysis, the same values used for the TCP/IP analysis were used. The TCP/IP packets were placed into 2048 bit data fields of X.25 packets to show the results of encapsulation. The optimal achievable average packet size using this encapsulated traffic is estimated at 2120 bits for the X.25 packet with an average overhead of 3.9 percent. The average normalized packet delay for this scenario is 2.60 seconds for the transfer of an entire TCP/IP packet, and the average throughput for the network is 23.72 kbps.

4.4.7 Simulated Current Network Failure

Failures occur on any network. Component failures happen due to a number of causes, including fire, flood, power outages, human error, and software failures. Whatever the cause, the management system and the network must be reliable enough to deal with the situation effectively. Likewise, when a network management system is employed, it should not take a bad situation and make it worse. In the past, this case has not been uncommon, especially for SNMP-based management.

However, the recommended design will ensure that there will not be SNMP storms that flood the network with management overhead and information whenever there is a critical failure. There are two types of critical failure: node and link.

4.4.7.1 Node Failure

In a node failure, the access point and all of the devices behind it are cut off. Therefore, the network is denying service. Although the denial of service is not good, it is not a terminal condition for the network backbone.

When a node does fail, there will be an initial traffic spike on the rest of the network caused by rejection and time-out notifications, routing updates, and SNMP alarms and traps. Alarms will first be generated to notify the network management center of a problem. The management center will respond with a series of traps. Network management traffic will increase until the problem can be diagnosed. Meanwhile, the active nodes continue to send traffic to the failed node. If HDLC is in place, only a limited amount of traffic (the window size) can be sent, otherwise, traffic will be continually sent, and will continually time out until the nodes get updated.

Using EIGRP, in the worst case it should take less than one second to notify the other nodes that a node is not responding. Once the routes converge, the node will be flagged unusable. At this point, traffic will no longer be forwarded to that node, but rejected and returned to the source station. Now, there is less traffic on the backbone. The situation will remain this way until the node comes back on line, broadcasting that it is back on line. Nodes will then update their routing tables and all traffic will continue normally.

Figures 24 and 25 show the average delay and average throughput in the case of a failure in Node C with respect to events over time.

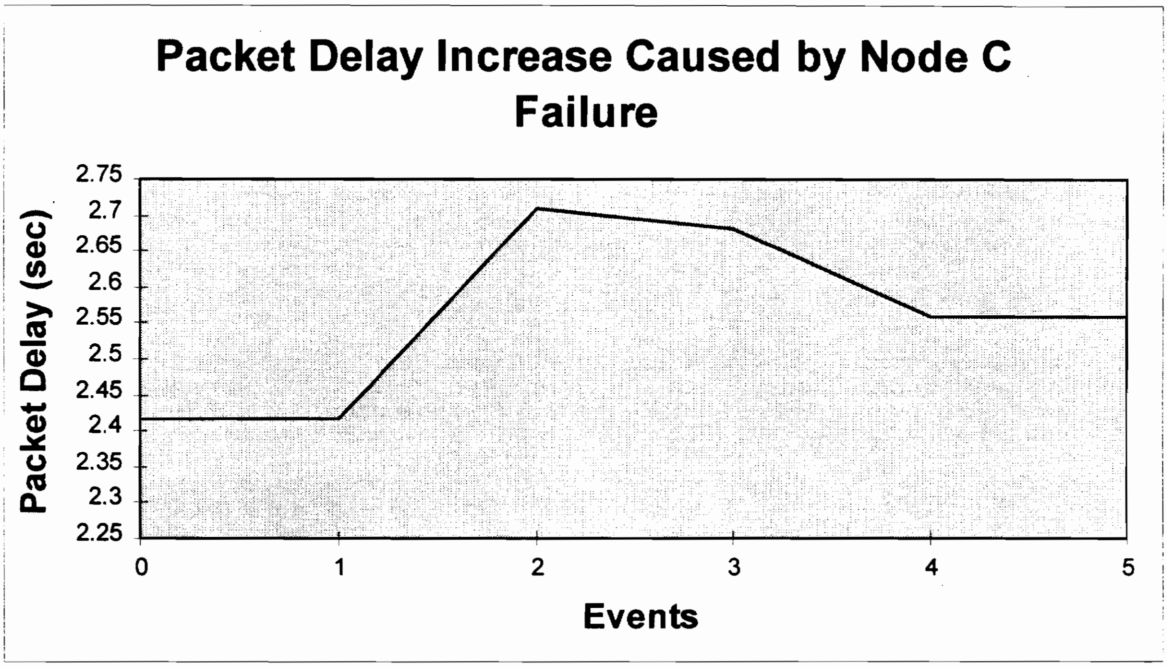

Figure 24 : Packet Delay for a Failure of Node C

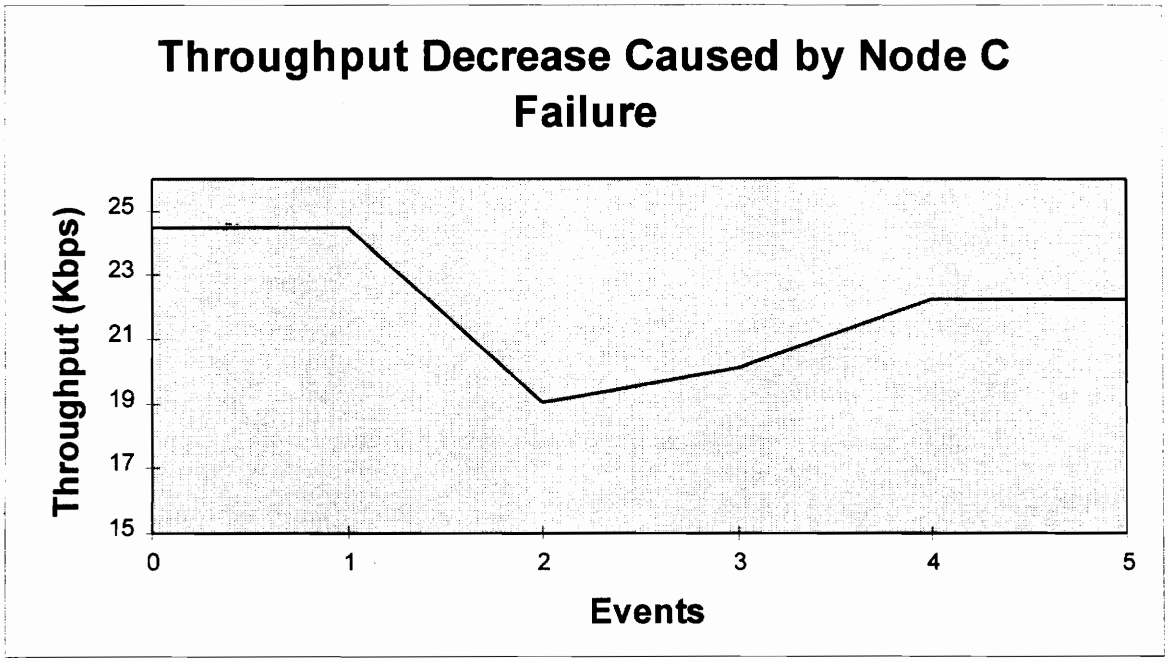

Figure 25: Effective Throughput for a Failure of Node C

4.4.7.2 Link Failure

The other type of failure is a link failure. This is actually much worse than a node failure for the network as a whole. Although a link failure may not deny service, because of multiple link connections, it can overload another link segment.

When a link does fail, there will be an initial traffic spike caused by rejection and time-out notifications, routing updates, and SNMP alarms and traps. Alarms will first be generated to notify the network management center of a problem. The management center will respond with a series of traps. Network management traffic will increase until the problem can be diagnosed. Meanwhile, the connecting active nodes will notice that the line is down.

Using EIGRP, in the worst case it should take less than one second for the network to recognize that a link is not responding. Once the network converges, the link will be flagged unusable by all nodes and traffic will be re-routed. With the re-route, however, it is possible to overload a link. This is especially true if you have two links in parallel, such as on Nodes A and F. Here, if one link fails, all of the offered traffic will be placed on one link. In the case of A, the offered load is greater than the supported throughput.

The re-routed traffic must be handled by the remaining links and nodes. Re-routing can cause a cascading effect that could be very detrimental to the network. The example of losing link AB is shown below. Figure 26 depicts the average throughput for the entire network, with a degradation of seventy-five percent.

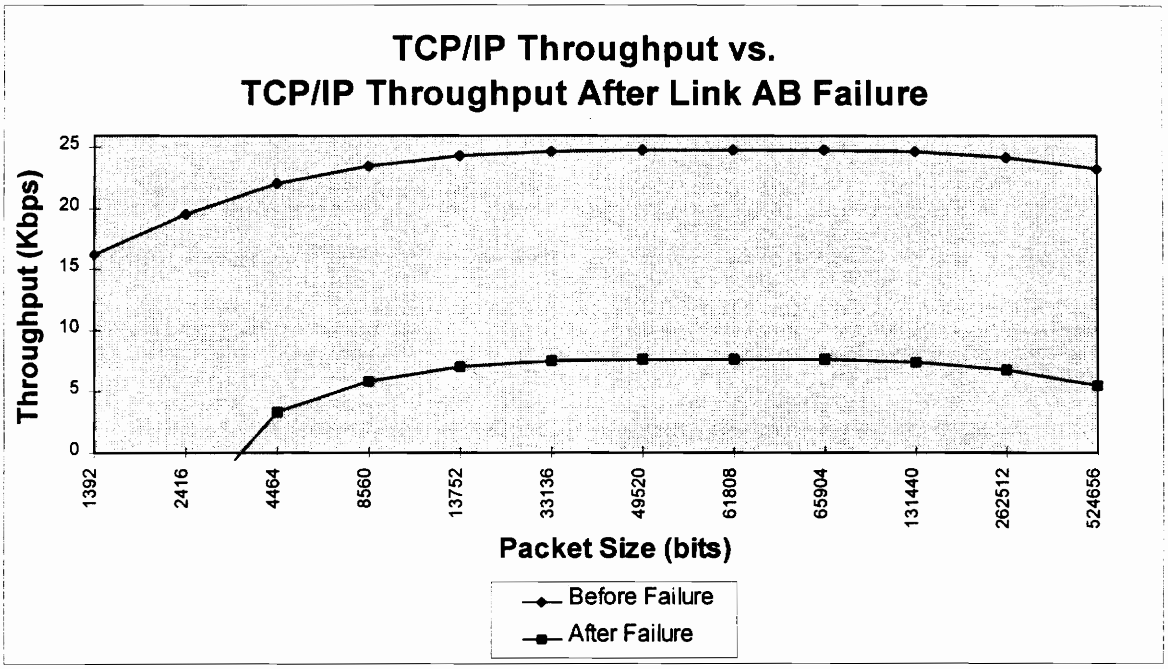

Figure 26: Effective Throughput for a Failure of Link AB

The situation will remain this way until the link comes back on line and the routers broadcast that it is useable. Nodes will then update their routing tables and all traffic will continue normally.

5. Detail Design

The detail design starts with the configuration specified in the preliminary design. It further provides system specifications, detailed documentation, and development of the prototype and as-built system.

5.1 Project Schedule

The project schedule to be followed is shown in Appendix A. It shows the minimum amount of time and resources necessary to complete the project within budget. All design specifications except encryption can be implemented according to schedule.

Members of the technical staff are responsible for research, development, implementation planning, capability assessment and testing, and technical documentation. Each member of the team shall be responsible for development in his or her area of expertise. The two senior systems engineers shall coordinate resources, initiate design review activities, and direct the formal systems approach.

5.2 Quick Fix Recommendations

From the descriptive scenario presented in section three, there are two methods of action that should be taken immediately, to be followed by the detailed design specifications. The first action is to disable IPX packet processing on the AIX servers.

IPX packet processing represents one of the largest consumers of processing time on the CPU of the routers and servers. Any reduction in the amount of IPX packet processing is reflected by a reduction in processing time on the CPUs. Similarly, elimination of other protocols on unnecessary interfaces reduces the burden on the processors as well. It is recommended that all unnecessary protocols be removed from any interfaces where they are not required. In particular, it has been identified that the AIX servers are likely candidates to have IPX processing removed. This change should reduce broadcast traffic and lower CPU utilization.

The second action is to investigate and implement UDP broadcast spanning tree support. The current UDP forwarding mechanism employs helper address statements to provide for redundancy by forwarding multiple copies of the UDP broadcast throughout the campus environment. Implementing a UDP broadcast spanning tree would allow the management portion of a single copy of the necessary UDP broadcast to move through the network without causing increased traffic. This change will help reduce broadcast traffic and utilization through the elimination of multiple helper address statements.

5.3 Design Specifications

An open system for real-time communications is one in which a third party also adheres to the standard upon which the system is based. Architectures and standards provide structure for many levels of interconnection. Standards support the passing of information between the various elements of an enterprise information system. A good paradigm for open communications architectures is represented by the standards that characterize UNIX and TCP/IP. Standardization provides the benefits of common systems for cost savings, improved integration, and information flow. Creating an open, uniform network eliminates situations arising from architectural problems, such as congestion, data loss, and component failures. Although problems can occur in any network, standardization and uniformity can help prevent them.

5.3.1 Address Space

When constructing uniform architectures, uniform addressing is necessary. If a server is relocated, addresses contained within its clients must also be changed. By referring to servers by name rather than address, however, the actual address of servers becomes irrelevant. Hence, the reason to use DNS. The relocation of a host requires only that the DNS server be updated to provide the new address, through either tables or rules. Table-based addressing requires that the system administrator maintain a table of records that consists of names and associated aliases. Each machine may have multiple aliases. Rule-based addressing requires that the system administrator maintain a correspondence between a name and an alias according to a pre-defined set of rules. SeMCom shall use table-based addressing for local users due to configuration management concerns. Remote users shall use rule-based addressing.

SMC applied for and received a Class "B" IP network number of 172.22.0.0. Class B addresses use sixteen bits (the first two decimal places) for the network portion and sixteen bits (the last two decimal places) for the host portion of the address. To successfully maintain addresses, the use of subnetworks is recommended.

5.3.1.1 Subnetworks

Subnetworks organize hosts into logical groups. Subnet routing allows numerous subnetworks to exist within a given network. If subnet routing is utilized, then the bits in the host field (the last sixteen bits for Class B) are divided into two groups, subnet and host. Thus, the final address has three parts: network, subnetwork, and host. Since the system does not know which bits in the host field are to be interpreted as the subnetwork portion, a subnet mask is required.

The subnet mask informs the system as to which bits should be interpreted for each portion. A subnet mask is a thirty-two bit number with one-to-one correspondence between each of the thirty-two bits in the address. The first field of the mask is always 255 so the system can interpret the network number. Generally, the entire eight bit field is either turned on (255) or off (0). However, for added flexibility, variable length masks can be employed. This allows a system to use a variety of configurations.

For this design, an eight bit subnet mask is recommended. An eight bit subnet mask appears as 255.255.255.0 in dotted decimal form. In this scheme, ranges of host numbers are reserved for similar devices; there can be 254 subnets per node with 254 hosts per subnet. Limiting the number of hosts and subnets helps to reduce congestion on separate segments and to prevent broadcast storms that SMC was experiencing previously. Subnetworks divided among the campuses is shown in Table 19. If, during capacity expansion, address space and broadcast domains are an issue, a variable length mask could be employed.

Table 19: IP Address Assignment

Node	Campus	Addresses	Total Networks
A	1	172.22.1.0 through 172.22.36.0	36
Internet Access	1	172.22.1.0 through 172.22.2.0	(2)
B	2	172.22.37.0 through 172.22.70.0	34
C	3	172.22.71.0 through 172.22.104.0	34
D	4	172.22.105.0 through 172.22.138.0	34
E	5	172.22.139.0 through 172.22.172.0	34
F	6	172.22.173.0 through 172.22.206.0	34
Future Use	TBD	172.22.207.0 through 172.22.240.0	34
Serial Interfaces	All	172.22.241.0 through 172.22.254.0	14

This scheme allows for at least 9144 hosts per node. Even if half of the hosts are not user platforms, there would still be well in excess of the required 3000 users per node. Node A (Campus 1) has the capacity for over 9652 hosts, including the two networks reserved for Internet and firewall access. The thirty-four networks designated for future use can be placed within an existing campus or assigned to a new campus. Additionally, serial interfaces can be reassigned with variable subnet masks if necessary.

Each subnetwork also follows a standard for numbering attached hosts. Host addresses 1 through 16 are reserved for routers while addresses 17 through 48 are reserved for hubs, switches, or other SNMP devices. Additionally, 49 through 76 are reserved for servers and gateways while 77 through 254 are given to departments.

5.3.2 Connectivity and Availability

Campus backbone connections shall have critical availability. Critical availability environments provide almost immediate recovery by preventing the loss of business continuity caused by multiple failures. Recovery brings users back to operational status, but may not be a permanent fix. After recovery, any repairs necessary should be completed in the MTTR time frame. Service level agreements will determine the accepted maximum and mean time to recovery. The mean time to repair (MTTR) expected values will vary on a case-by-case basis, but the maximum allowable MTTR shall be six hours. An example configuration is shown in Figure 27.

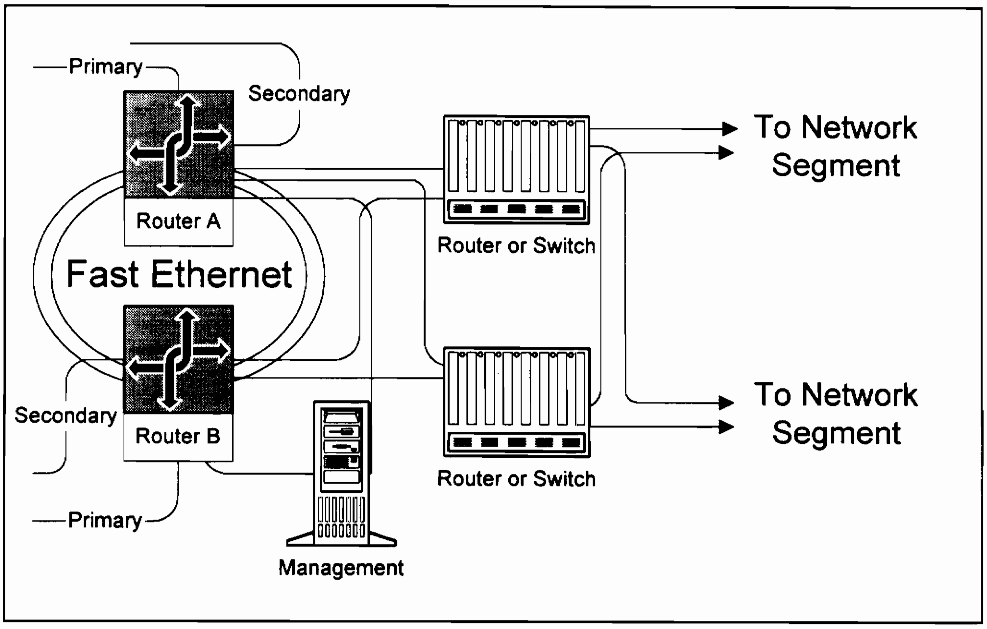

Figure 27: Critical Availability Architecture for CIM Campus Connectivity

5.3.3 The Collapsed Backbone

The key distinction between the distributed backbone and the collapsed backbone is that the once geographically distributed backbone has been centralized. In the collapsed backbone architecture, the router and the switch play a key role in traffic management. The router provides traffic isolation between the segments and the backbone. In addition, wide-area connectivity is provided by interfaces on the routers, isolating traffic to only those segments where it belongs. Switches perform layer two services and extend the geography of the network. There are two main types of backbone components, backbone access and subnet service.

Backbone access components provide concentrated access to the network backbone network via connections with subnet service components. Access components may be used to route high traffic volume directly onto the backbone if upgrading the service is not feasible. Generally, though, access components are fed by service components. Access routers will be high-end, backplane-based routing platforms supporting the standard routing and transport protocols described above.

Service components provide traffic concentration and filtering from larger, campus connectivity sites. Service routers will route between multiple local segments, subnetworks, and between a limited number of remote segments if necessary. A remote link will also use a service router, but it will technically be a form of an access router. Service routers support full routing and transport protocols.

Using redundant fast Ethernet links for the campus connection offers a high-bandwidth path, allowing users to be distributed among Ethernet segments coupled to the WAN. This enables segment size to be scaled to the various bandwidth requirements. Each campus site varies in size and requirements for data access and must be specifically designed to meet those requirements. The architecture presented here illustrates a number of features, including redundant components and dual-homing of hosts.

Dual-homing protects the network from a single component failure. Each component has a connection to both primary and secondary links. Redundant components protect from a single point of failure that would isolate the backbone. The risers to the primary and backup components are run in separate conduits and cross-connected, ensuring access to the backbone if a cable is cut or if a component fails.

It is also recommended to enable the Hot Standby Router Protocol (HSRP) with redundant connections to all segments. Enabling the HSRP will allow workstations to employ a common default gateway address that is shared by both routers. This is done by assigning two different IP addresses as the default gateway for each router. In this proposed configuration, the client is oblivious to a router failure as the backup router assumes the same MAC address with which the client had been communicating.

The integration of switching technology into the network requires careful analysis to ensure that the many diverse products are adequately understood and managed. If the existing infrastructure offers insufficient bandwidth, switching offers a possible alternative. The careful review of problems in SMC's network has demonstrated the need for some layer three intelligence, preferably routing. Current layer three switching products and their immature technology base are not recommended at this time. Switching shall be confined to layer two switching only. Networks in excess of three hundred active connections shall be separated via layer three devices.

The basic problem of having no layer three devices, besides broadcast isolation, is management. Management software works well with routed networks if a hierarchical subnet structure is employed. The software can delineate routers and locate hosts by the subnet portion of their address. Because of the correlation between network subnets and physical media, the software can associate the logical location with the network address of a host. With layer two networks, however, the rules of IP network addressing require that all hosts be on the same protocol subnet. Thus the NMS can no longer correlate the logical location of the host and a network segment, making the topology then appears as one "flat" network with no protocol hierarchy, making management very difficult.

5.3.4 WAN Configuration

In the design and roll-out of a network backbone, an intrinsic portion of the design goals include simplicity, low overhead, robustness, stability, flexibility, and speed. The recommended solution is to migrate the packet switches to routers and to replace the low-speed leased lines with frame relay and ISDN links.

The frame relay network will have a Committed Information Rate (CIR) initially of 256 kbps, expandable to 1.544 Mbps and higher. Nodes shall be replaced one-by-one with the Cisco access routers. Node A shall be completed first since that is where the CIO wishes the first network management console to be, with its backup at Node F. The recommended path is A, B, C, E, F, and D. This way, there is minimal conversion between protocols from point to point. Using IP between routers, X.25 between packet switches, and IP encapsulated in X.25 between router and packet switch will result in minimal translational processing, thus increasing overall network speed. Figure 28 shows the high-level connections.

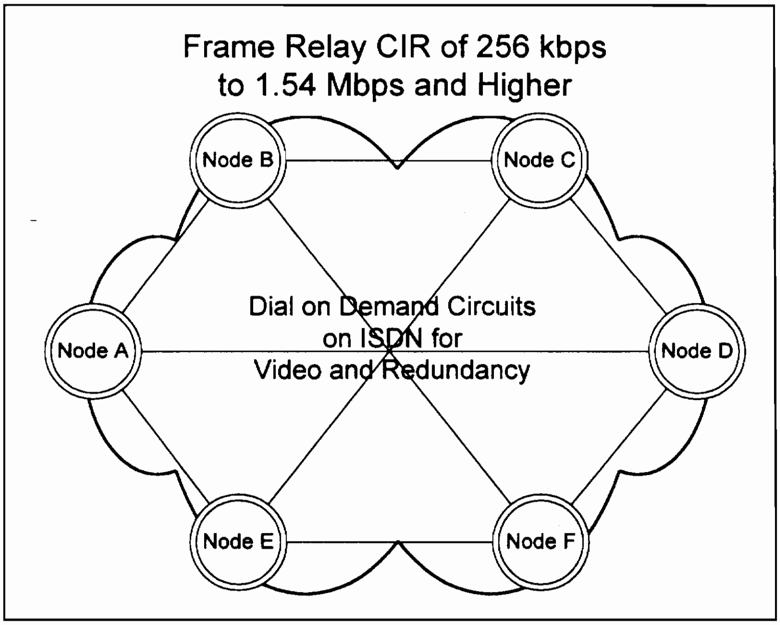

Figure 28: High-Level Connectivity

The ISDN lines shall be used for two purposes, video transfer and redundancy. Normally, only video conference traffic shall move across the ISDN links. Each node can be assigned a dial on demand route, whereby the link is not always in an active state. Instead, the link will disconnect itself when not in use, and establish itself dynamically as needed. Additionally, if the main links fail, the ISDN links will automatically activate. This greatly enhances reliability, reduces cost, and allows for exceptional quality video.

As the network grows, more bandwidth will be needed. Multiple T1 circuits will provide the best way to deliver bandwidth to the frame relay network. Deployment of T3 circuits to deliver fractional T3 bandwidth will not be economical since new cable will need to be installed. It is recommended to immediately convert to T1 connections and to add them as necessary. The risk for multiple T1 circuits will come from the service provider's ATM services. As carriers offer ATM services and local exchange carriers deploy fiber-access circuits, ATM controllers will assume much of the transport role. Data, voice, and video can be supported using high-speed ATM. Also, nodes A and F should have direct connections to provide better redundancy. WAN configuration diagrams are shown in Figure 29.

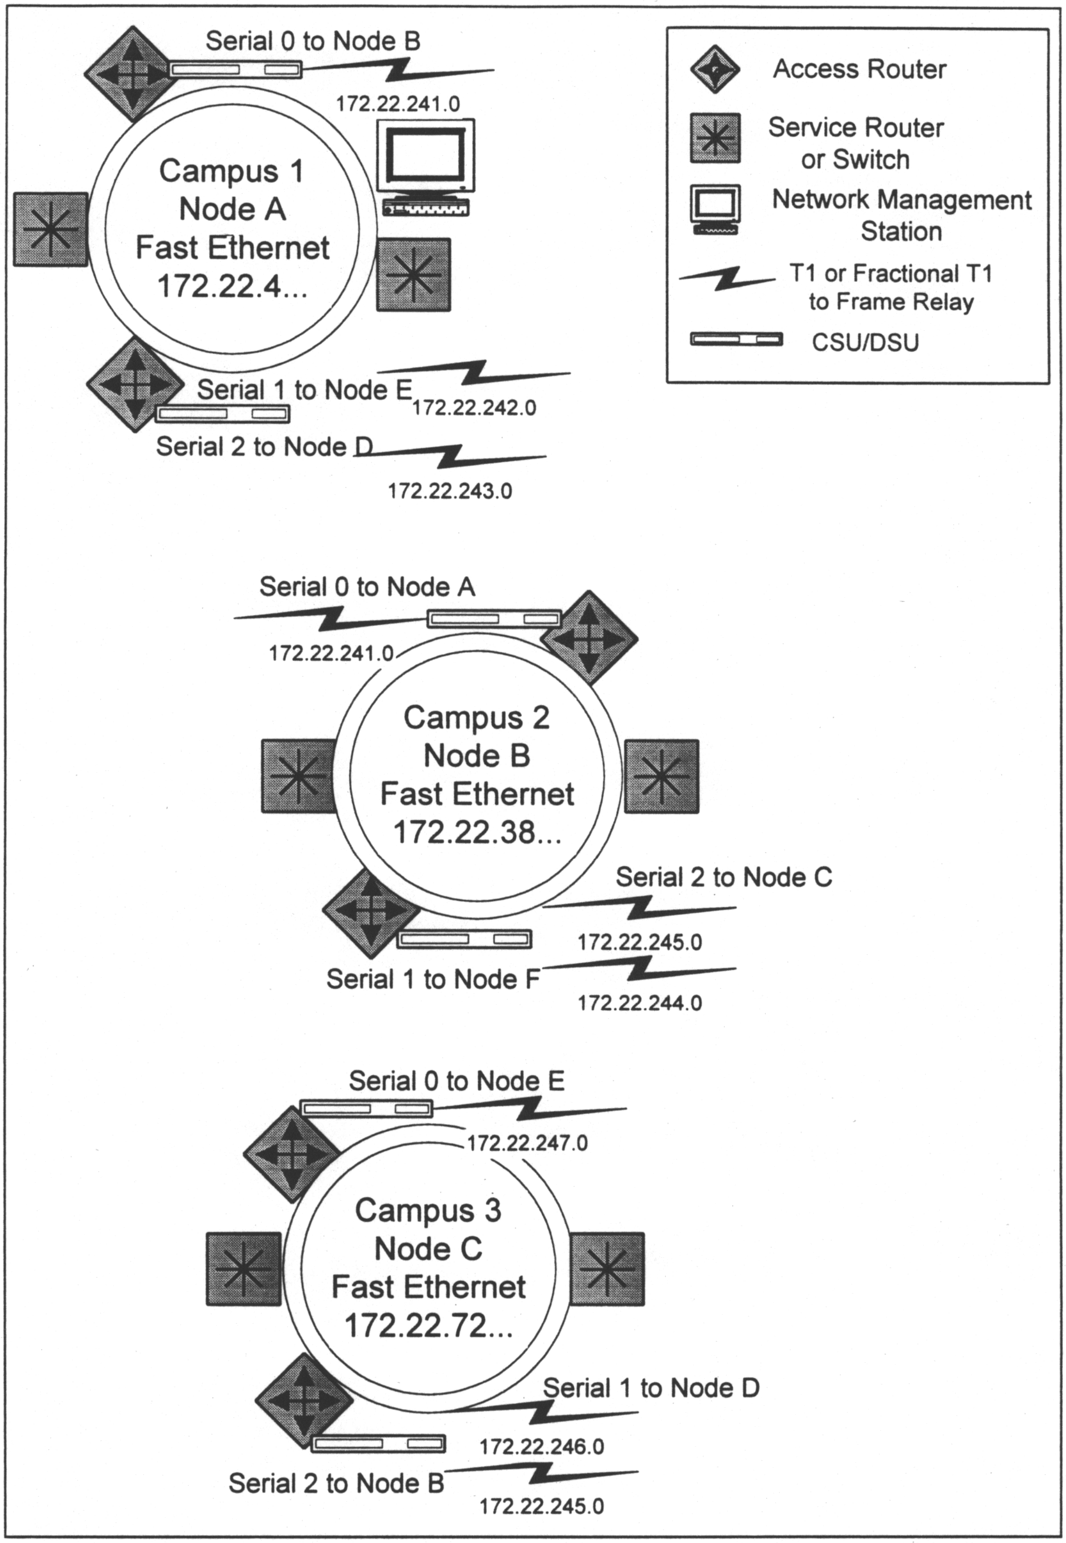

Figure 29: WAN Connection Diagrams for Campuses 1, 2, and 3

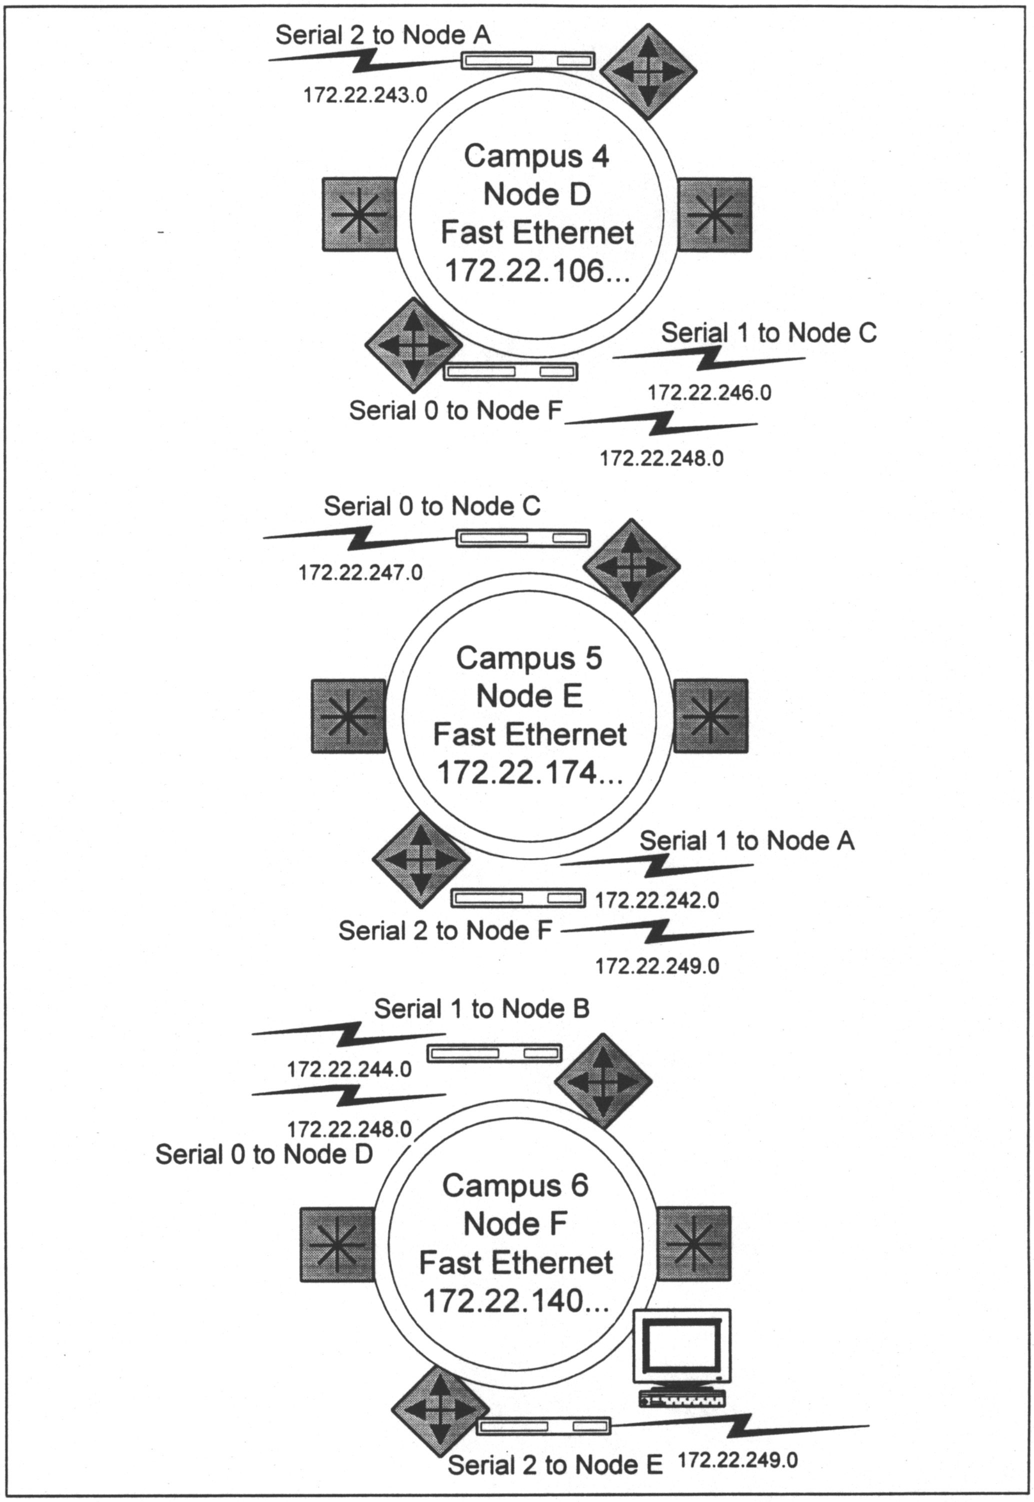

Figure 29 (cont): WAN Connection Diagrams for Campuses 4, 5, and 6

Mathematical analysis and the preliminary test plan will be used to model the architecture. Table 20 dictates the direct ISDN connections.

Table 20: WAN Connection Plan

Node	Campus	Direct Connections
A	1 (Primary NMS)	B, D, E
B	2	A, C, F
C	3	B, D, E
D	4	A, C, F
E	5	A, C, F
F	6 (Secondary NMS)	B, D, E

The ICMP (Internet Control Message Protocol) will be used for error and configuration messages. ICMP supplies a number of different messages including echo, redirect, time exceeded, and router advertisement. It is an integral part of any IP implementation. EIGRP will be used for dynamic routing. Initially, static routing will have to be used because of the packet switches. In case of an emergency, it is much more effective and reliable to have a dynamic protocol than to have a static table.

With a single-protocol backbone, all routers are assumed to support a single routing protocol for a single network protocol. In this type of environment, all other routing protocols are ignored. If multiple protocols are to be communicated, unsupported protocols must be encapsulated within the supported protocol or they will be ignored. Encapsulation (tunneling) takes packets or frames from one system and places them inside frames for another network system. This provides a means for encapsulating packets inside a routable protocol via virtual interfaces. Synchronous Data Link Control (SDLC) transport is also an encapsulation of packets in a routable protocol. In addition, SDLC provides enhancements, such as local data link layer termination, broadcast avoidance, and media conversion services.

Tunneling can disguise the nature of a link, making it look slower, faster, or more or less costly than it may actually be in reality. This can cause unexpected or undesirable route selection. Routing protocols that make decisions based only on hop count will usually prefer a tunnel to a real interface. This may not always be the best routing decision because an IP cloud can comprise several different media with very disparate qualities; for example, traffic may be forwarded across both 100 Mbps Fast Ethernet lines and 9.6 kbps serial lines. Attention must be paid to the metrics used by each protocol.

Because encapsulation requires handling of the packets, it is generally faster to route protocols natively than to use tunnels. Tunneled traffic is switched at approximately half the typical process switching rates. This means approximately 1000 packets per second aggregate for each router. Tunneling is CPU intensive, and as such, should be used cautiously. Routing updates, service updates, and other administrative traffic may be sent over each tunnel interface. It is easy to saturate a physical link with administrative information if several tunnels are configured over it. Performance will depend on the passenger protocol, broadcasts, updates, and bandwidth of the physical interfaces. It is also difficult to debug the physical link if problems occur.

5.3.5 Network Management

From the management environment, all elements are centrally managed, monitored, controlled, operated, and administered from the NMS. The primary NMS is located at Campus One, with the secondary site located at Campus Five.

The management data repository as well as the manager component of all management applications are located here. All applications for the network administrators are provided in the administrative domain, like those for software management, configuration and change management, and security management. All applications for the network engineers are provided in the operative domain, for operational control, problem tracking, and performance management.

In large environments, a balance must be struck between the number of objects that a management platform directly monitor and the size of the hardware platform required to support that management function. Hierarchical implementations of management environments are possible through the configuration of event handling and topology mapping. Each distributed management station can be configured to report events to another management station. A set of criteria can be established that will guide the forwarding of events from remote to central locations. The primary management station can have knowledge of significant outages, but will not be notified of local events. This configuration allows for the central operations staff to be aware of critical events that may not be in their direct management domain. This type of configuration is referred to as "Manager of Managers." [12]

Monitoring the network with SNMP requires that certain information be queried from managed objects. Each managed object typically has a MIB associated with it that has a variety of information and statistics about the managed objects in the network. For various types of monitoring tasks, Table 21 shows the MIB objects that should be used. The polling interval is dependent upon how critical the network is and how much impact on delay the network management traffic is allowed to have. Typical TCP/IP networks poll objects every fifteen minutes.

Table 21: Recommended MIB Objects

Category	Objects	Description
Utilization	ifInOctets	Number of octets into an interface
	ifOutOctets	Number of octets out of an interface
	ifSpeed	Interface (line) speed
	sysUpTime	Time since system came on-line
Status	ifOperStatus	Device status: up, down, testing
Errors	ifInDiscards	Number of incoming discarded (balked) packets
	ifInErrors	Number of incoming packets with errors
	ifOutDiscards	Number of outgoing discarded (balked) packets
	ifOutErrors	Number of outgoing packets with errors
	ifInUnknownProtos	Number of packets with unknown protocols
Congestion	ifInDiscards	Number of incoming discarded (balked) packets
	ifOutDiscards	Number of outgoing discarded (balked) packets
	ifOutQLen	Length of outgoing queue
	tcpToAlgorithm	TCP retransmission algorithm
	tcpRetransSegs	Number of retransmissions on a segment
Traffic Patterns	ifInNUcastPkts	Number of incoming non-unicast packets
	ifOutNUcastPkts	Number of outgoing non-unicast packets

Remote Monitoring (RMON) in network management allows for the retrieval of information pertaining to entire segments or hosts on a network.[13] Table 22 shows examples of the type of MIB information that would be retrieved for remote monitoring.

Table 22: Recommended RMON Objects

Category	Objects	Description
Ethernet Status	etherStatsOctets	Statistics on Ethernet segment
	etherHistoryUtilization	Utilization history on Ethernet segment
Host Information.	HostOutPkts	Outgoing packets from a host
	hostOutErrors	Outgoing packets with errors
Status	hostOutBroadcastPkts	Outgoing broadcast packets
Source to Destination. Information	MatrixSDPackets	Number of source to destination packets
	matrixSDOctets	Number of source to destination octets
	matrixSDErrors	Number of source to destination errors

To keep costs down, the NOC can take advantage of existing hardware platforms at each campus location, but some amount of hardware and software upgrade is necessary. Additional memory and software modules are required. For each campus, this shall cost approximately $35,000. It is recommended that the management software be placed under a software maintenance contract to ensure the latest software configurations. Yearly costs for maintenance contracts will total approximately $5000 per node.

The configuration of the network management will vary depending upon preferences, management traffic impact, and network management operating procedures. As it is presumed that full network monitoring and management capabilities will be desired, management stations will be configured to poll stations on a predetermined interval. All routers and hubs will be configured to send traps to both main management platforms. This configuration will enable both network operations centers to act independently should one site's management system fail, a facility has a complete or partial outage, or the WAN links are lost.

For the impacts of network management on the average packet delay in the system, the "typical" traffic scenario considers 200 kB file transfers. In addition to the data and control packets for file transfers, ICMP and EIGRP traffic is added to the average traffic load. The ICMP traffic consists of a 196 byte message every 120 seconds, while the EIGRP traffic consists of 118 byte message every 90 seconds. This additional traffic add an average of 525 bits per second of data across the entire network.

For analysis of the network management traffic, some assumptions have been made for simplicity and clarification. First, ARP and BOOTP traffic is not considered. Second, the average packet size is not affected by the addition of management traffic.

5.3.5.1 Maximum Allowable Network Management Traffic

The maximum allowable amount of network management traffic is derived for the limitation of not increasing the average packet delay by more than five percent. For this analysis, the "typical" traffic scenario of file transfers is used. The assumption that the average packet size is not affected has minor ramifications. As the average packet size drops, overhead increases, and allowable amount of management traffic decreases.

The increase in packet delay for the three types of network traffic being analyzed, X.25, TCP/IP, and TCP/IP encapsulated in X.25, has varying results. The worst case, where the five percent increase offers the lowest additional overhead, is with X.25. Using the results from X.25, management can add an average of up to 6.2 percent traffic on the network. This translates to an average of approximately 7.7 kbps of allowable management traffic on the network. One concern with this approach is that the average throughput decreases by ten percent. Since the worst case was used, when the fully implemented IP network is complete, slightly more management traffic will be allowed.

5.3.5.2 SNMP

SNMP is the protocol by which network management information is retrieved. The SNMP packets are sent between monitoring stations and the monitored objects. The transmissions are in the form of queries to objects or traps from objects.

5.3.5.2.1 SNMP Packet Layer Analysis

The transfer of SNMP traffic over IP networks involves placing SNMP information into UDP transport layer packets. The UDP packets are then placed into the data field of an IP network layer packet, and finally into the information field of an HDLC data link layer packet. This packet is sent across a physical channel. Figure 30 shows an example of this encapsulation.[14]

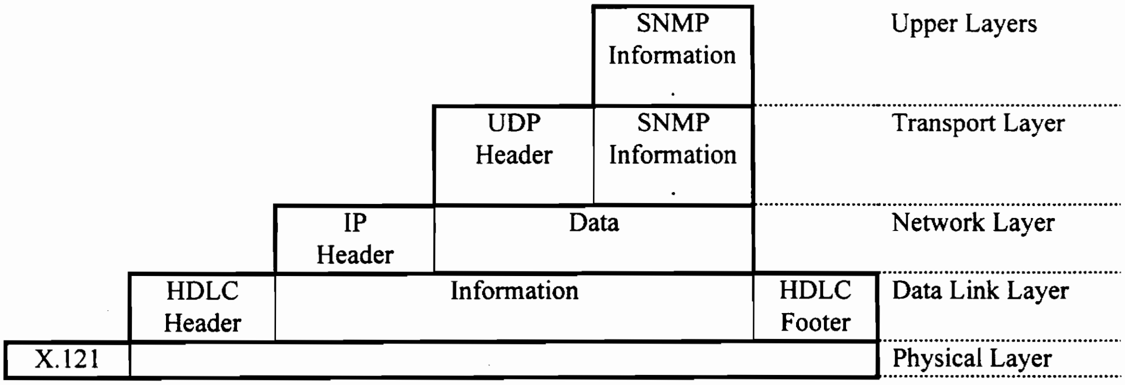

Figure 30: TCP/IP Protocol Layers.

The UDP, IP and HDLC data packet formats are used to determine the amount of overhead associated with the transfer of SNMP information during a transmission. Overhead associated with UDP, IP, and HDLC packets are sixty-four bits (eight bytes), 160 bits (twenty bytes) and forty-eight bits (six bytes), respectively.

5.3.5.3 Network Management Delay Analysis

The packet delay is analyzed based on the management information on the network for the proposed polling of recommended SNMP MIB information every five minutes. The SNMP MIB objects are broken down into two categories, interface information (if) and TCP information (tcp). The statistics gathered from interface objects are relevant for each protocol on each port of the polled device. For example, Node C has three links connected to it, which correlates to nine ports each having two protocols, X.25 and IP. Therefore eighteen polls for interface traffic is required for that node.

The interface traffic introduces 292 bytes of information for a Get command to the polled device, and 326 bytes for the GetReply from the device. TCP traffic adds 106 bytes for the Get and 110 bytes for the GetReply. The entire calculated load of traffic is doubled for the retrieval of information by the secondary management station. This is done primarily in case node A fails, the statistics will be available for analysis as to why a failure occurred.

Recommended management traffic totals to an approximate average of 1986 bits per second over the entire network. Overall, this adds about 1.56 percent overhead to the regular traffic. Again analyzing the worst case impact of additional overhead on the network for the X.25 case, the average packet delay increases by 1.18 percent, which is considerably less than the allowable amount. Likewise, the throughput on the network drops by an average of 2.62 percent due to the additional traffic introduced for management tasks.

5.3.6 Security

SeMCom shall support an entire range of unclassified and classified information, ranging from publicly available data to company confidential material. It shall operate across a collection of disjoint company networks. Each component network shall provide services with different security protection requirements. SeMCom will not be responsible for the intranetwork security of the networks that comprise the common communications environment and assumes that they are operated in accordance with their stated security policies. SeMCom will employ a collection of security devices to provide in-depth security measures. This collection of devices will work together to create an environment more robust than could be obtained using any one device.

Data confidentiality requires that encryption and firewalls be employed. SeMCom shall establish a secure, encrypted connection across the various networks at the network layer of the communications stack. There are several commercial systems currently available that provide this type of service. Since Cisco routers shall be employed, their partnered IOS/Cylink encryption mechanism shall be employed. Data integrity services provide assurance that the data received is the same the same data that was sent, or that modifications to the data are detected to a predetermined level of confidence. Since SeMCom assumes a collection of "system high" networks, there are no requirements for mandatory access controls to differentiate users based on the classification level of the information traversing the WAN links. For Internet connectivity, Checkpoint's Firewall One shall be employed because it is all ready in place in the organization. Firewall One provides access-list security so that the network cannot be probed from the outside. Security services are summarized in Table 23.

Table 23: Security Services

Security Services	Enabling Mechanism	Implementation	
		Products	Standards
Confidentiality	Encryption	Cisco IOS	DES, RSA
Integrity	Checksums		CRC-16, CRC-32
Authenticate	Firewall, Passwords, PIN	Checkpoint	DES, RSA
Access Control	Firewall	Checkpoint	
Non-Repudiation	Certificates	Cisco IOS	
Availability	Redundancy, Fault Tolerance	Infrastructure	
Management (Physical)	Cipher Locks, Key Cards, Fences, Guards	TBD	
Management (System)		HP OpenView, Tripwire	SNMPv2, RMON

In addition to deciding which security measures to employ to connect an organization's internal network to the common network, each business unit must establish and enforce a security policy. The strength of the security policies are only as strong as their effective implementation and enforcement.

5.4 System Test Plan

The system test plan shall specify the purpose, goals, and overall description of the testing to be performed, including the scope and any controls to be used. Plans shall state the requirement to be tested, the title of test, the element of the system being tested, the test method (analysis, inspection, or demonstration), inputs, expected outputs, required accuracy, collection methods, quality assurance, and acceptance criteria. It shall also specify any resource requirements.

Procedures shall detail the cases to be executed. These procedures shall include a description of the test method, conditions necessary to complete the test, an unambiguous step by step description of the process, acceptance criteria, results, logical and physical diagrams, the name and date of the specific test, the name of the test engineer, a subsequent pass or fail grade, and necessary actions to take in the event of a failure. Procedures shall be performed and results recorded for each campus location. System level tests will be conducted in a "typical" operational environment where the segment of the system being tested is not isolated, but in full connectivity as it would be in a production atmosphere. Thus, all parts of the system shall be operational, regardless if they are directly related to the procedure or not. Each campus will have a baseline test procedure along with a unique set of test procedures depending upon the numbers and types of products and services that it owns. The simulation and test plans shall be used to verify operational status and identify deficiencies. A planned review of the results shall determine if the campus site is ready to go on-line.

5.4.1 Test Goals

Test procedures shall define appropriate line speeds, verify connectivity, and demonstrate robustness and reliability. The goal of this testing is to ensure that appropriate bandwidth was allocated for various configurations, correct addressing was assigned, and all components are functional and interoperable. These procedures exercise the response time of SeMCom given the respective bandwidth required for the respective campuses. The scenarios will test operation during peak and steady state operations.

Completion will yield three important results. First, the operation of the indicated number of clients over the link running at the specified speeds and capacities will be validated. This will empirically indicate whether a link of that bandwidth is adequate for the indicated number of clients. Second, the collection of data may be used to analytically calculate the required data rate from the indicated number of workstations. Third, logical connectivity and operability will be verified. All results will be presented in a summary document.

Testing is broken down into three areas: research and development, testing and analysis, and documentation. Research and development involves determining proper data parameters, setting up hardware, software, configuring routers, and setting up workstations. Testing and analysis involves performing actual tests and mathematical computations to define expected bandwidth requirements per baseline millisecond response times. Documentation involves summarizing and publishing the analysis. Standard test forms will be created to facilitate documentation.

Inputs include the number of clients, the bandwidth of the interconnecting links, and the number and types of data packets. The number of clients parallels the number of remote workstations and drives the number of simultaneous packet streams that must be generated. Interconnecting bandwidth is the bandwidth of the telecommunications link or route to be tested. Data packets are those packets sent to the receiver that elicit a response. Each packet pair measures the time between a request and a subsequent response.

5.4.2 Configuration

Preparation for the tests begins by assembling a mock configuration of an actual campus, followed by the regular deployment network. The first tests will demonstrate operation concepts while the latter tests will test the actual network. A network analyzer is attached where it can observe the elapsed time between the generation of a request and the reception of a response in either case.

Once the configuration is assembled, the test begins with the triggering of either the simulation engine or the packet analyzer to collect data. From this point forward, all packets are collected. If the response time between all packet pairs is within given tolerances, the test is deemed a success, and the given bandwidth will support the given number of clients. Additionally, the data collected will be recorded to support an analytical calculation of the exact required bandwidth. Likewise, if total network utilization is within given tolerances, the test for the network as a whole is deemed a success, and that specific configuration is suitable for current requirements.

5.4.2.1 Test Tools

Tests will be performed with servers, routers, connecting links, clients, and a network packet analyzer. All tests shall follow the design layout and configurations set by the systems engineers as described above.

5.4.3 Router and Switch Test Plan

SeMCom will use Cisco routers to establish its WAN communications backbone. It shall also use Cisco switches for the campus backbones to establish communications from the individual nodes to the campus WAN interface at each site. Testing will ensure that connectivity has been established, all communications paths are functional, that addressing and directory schemes are working correctly, and that encrypted data can be properly recovered. Each campus will have an encryption in front of an access router. The campus infrastructure will hang off the primary and secondary access components. Each building segment will have at least one service switch or router.

The components must first be configured with the basic startup configuration that defines the communication protocols and assigns network addresses. Configuration files for each are unique and thus each must be tested independently. A complete configuration file contains parameters that permit or deny specific protocols or processes on by an address. Testing will intentionally violate routing, address, and processing rules and observe the result. If the components and configurations are working properly, any violation attempt well be unsuccessful.

Preliminary testing shall verify that data can be transmitted and received by the node components over the communication lines. An Internet Control Message Protocol request packet, better known as a "ping," is addressed to a remote object and transmitted across the communications line. If the components are able to communicate, an ICMP response will be received from the remote unit. A successful ping indicates that data can be transmitted. Each segment will be tested to ensure the requirements have been met.

5.4.4 System Level Test Plan

SMC is upgrading its communications infrastructure. The six corporate sites shall be connected via multiprotocol routers and frame relay transmission services. Each component of the system will have been tested prior to the system level test.

The system level test plan is written to test end-to-end functionality. It will verify that end users can successfully communicate with one another and provide an accurate report of any failures, discrepancies, or questions that need resolution.

At each campus, the hardware and software test environment shall consist of at least one standard workstation with standard desktop configuration, one access unit and one service unit. The routers will be configured to allow static routing, dynamic routing via EIGRP, and multiprotocol transmissions.

End-to-end testing shall be performed between each node. The test procedures are designed to test transmit and receive operations for database queries, documents, CIM commands, graphics and CAD files, electronic mail, SNMP commands, and voice and video conferences. Each function shall be tested by a user on every segment of every node. Each procedure ṣhall be repeated for each campus segment. At completion, results will be summarized and fully documented.

5.5 Training Plan

The key to providing true end-to-end service in any network management environment lies in the technical capability of the personnel entrusted with the responsibility and the available tools. Assuming that a viable set of tools are available to perform the tasks, a well-trained team of staff members is a prerequisite for proactive management that ensures quality service. The training efforts will therefore be directed with these objectives in mind.

5.5.1 Introduction

The sheer size of SMC with a vast array of requirements, unique environments, and customization needs poses a challenge to the training initiatives. The distributed nature of the departments and business units and their operations only compounds this challenge. The selection of HP OpenView as the corporate-wide standard for a management platform and the availability of several standard training modules from HP's training organization will help focus efforts on certain key areas. To make the training efforts specific and manageable, it is envisioned that the approach needs to be targeted towards broadly defined user communities.

In this context, three groups of personnel and three types of training have been identified as being necessary in the initial training offering. The target audience will be executives and managers, network engineers and technicians, and developers.

These categories of audience will be best served by executive overview training, technical training oriented toward personnel in the production environment, and technical training tailored to suit the needs of the application development staff. Advanced training may be designed in the future based on specific needs of the user community.

5.5.2 Curriculum

Executive training will be designed for managers and executives. The purpose of the program is to provide an introduction of the HP OpenView platform and associated products. HP's standard training offerings will be modified and customized. Specific attention and focus will be centered on the Operations Center and Network Node Manager products, since these two are the crucial entities that will affect network management efforts in the future. Training will cover the capabilities of the two products and how they integrate into the HP OpenView platform concept. This program will be two days in duration.

Technical training in production will be an important element of network management efforts. This program is intended for those first line users in the production environment that directly operate, manage, and design the systems and network resources. These personnel will be the engineers and technicians whose combined efforts directly affect the health of the network. Consequently, it is believed that this group will be best served in its day-to-day efforts by HP OpenView Operations Center and Network Node Manager courses. Operations Center collects management information, messages, and monitoring alerts from throughout the computing environment and presents them to the central management system. Upon receipt of the information, the Operations Center can initiate corrective actions and assist in problem resolution. Network Node Manager is a graphical management application that provides fault, configuration, and performance management capabilities for multivendor TCP/IP networks.

A preliminary examination of these courses has revealed some content overlap. Consequently, the customization of these classes will also include merging the two modules into one integrated offering. The program will be five days in duration and will introduce the users to the basic and intermediate network management capabilities of the two products and how to use them. It is also expected that the initial offerings of this course will be at local facilities. Advanced training in these modules will be designed based on input from the attendees and a review by NOC technical specialists.

Technical training in development will be targeted toward the development staff to aid in their efforts to develop management applications that are comprehensive and architecturally sound, to work efficiently in a multivendor, standards-based environment. Given the widely variant nature of business unit operations, several courses are available. These fall under the umbrella of developer products that include the SNMP Developer Kit and Platform and the Distributed Management Development Kit and Platform. The SNMP Developer Kit and Platform are designed for the developer creating applications to manage SNMP and TCP/IP devices or networks. The Distributed Management Developer Kit and Platform are for those engaged in application development to manage networks using a mix of communications protocol stacks. The type of course and the level of customization (if any) required will ultimately influence the length of these training programs. Advanced training programs in these areas can also be offered.

5.5.3 Development

The NOC staff will be in contact with business units to identify the available training. Currently, the NOC is examining the standard training offering for the OpenView family of products. Based upon this evaluation, it will identify the contents of the standard offering, the training needs of SMC, the level of customization needed, and the duration and frequency of training. Additionally, it will identify training options for other vendors' products and communicate the details to the business units.

Based upon the amount of training needs identified by the business units, the NOC will negotiate a pricing structure for the training and work with the business units and internal training organizations to identify the advanced training needs and the appropriate course contents.

5.6 Standards Plan

To ensure that forthcoming technologies and services are dealt with in a systematic manner, a standards process must be developed to introduce new or better subsystems and components to SeMCom. Figure 31 shows this process.

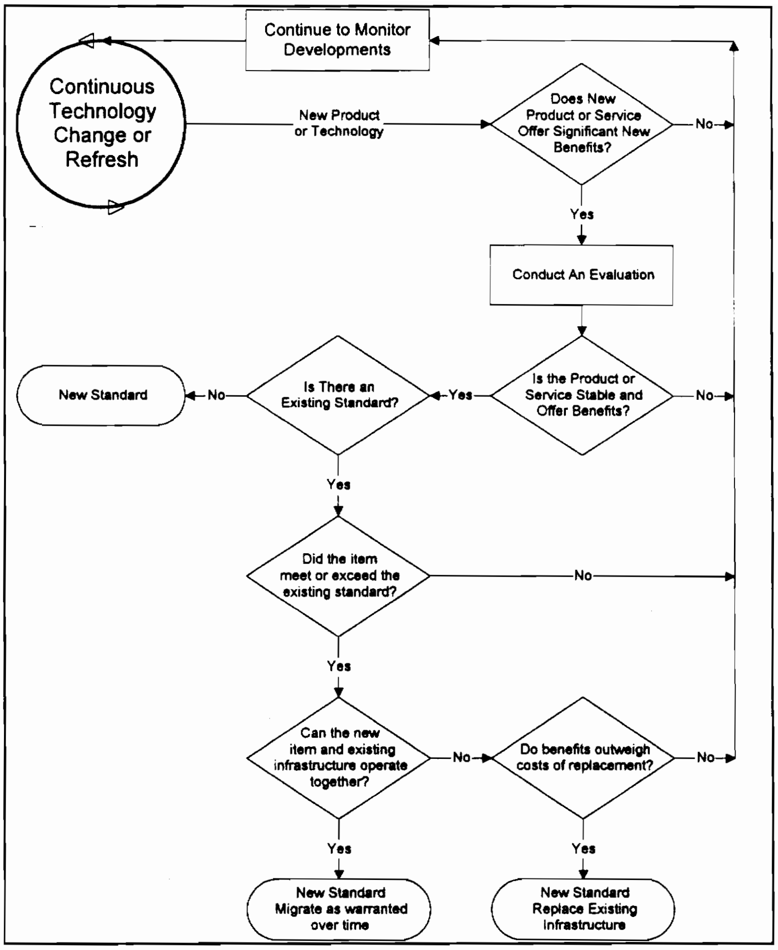
Continue to Monitor Developments
Continuous Technology Change or Refresh
New Product or Technology
Does New Product or Service Offer Significant New Benefits?
No
Yes
Conduct An Evaluation
Is the Product or Service Stable and Offer Benefits?
No
Yes
Is There an Existing Standard?
No
New Standard
Yes
Did the item meet or exceed the existing standard?
No
Yes
Can the new item and existing infrastructure operate together?
No
Do benefits outweigh costs of replacement?
No
Yes
Yes
New Standard Migrate as warranted over time
New Standard Replace Existing Infrastructure

Figure 31: Standards Process

6. Recommendations for the Future

For a complete systems engineering effort, some future work is necessary. It is recommended to continue to refine the preliminary system design and complete a formal design review with network and systems engineers. Functional flow diagrams can be developed to further refine requirements allocation to lower levels. More indices of performance could establish discriminators for component evaluation criteria and requirements. It is also necessary to develop detail design documentation such as interface control documents and implementation plans. Functionality testing should be completed and documented once the system installation is complete. A more detailed evaluation could be justified once the development for CIM software, hardware, and facilities moves to full production.

There is an emerging consensus in the industry that ATM is the best long-term choice to support a global, desktop-to-desktop network infrastructure. This new network model offers solutions to many problems in today's networks and will support new applications that can offer increased productivity. First, it is scaleable. ATM is defined to work at different speeds and on different physical media. Second, its traffic management protocols support different types of traffic. Third, it features a multiple priority system that provides the ability to integrate, prioritize, and manage different types of traffic. ATM can support a large number of subnets that do not need to be defined by physical locations. Not only does it offer higher speeds, but it can also scale to even higher levels of connectivity and bandwidth without requiring a change in architecture because it has the potential to consolidate multiple networks, meet future application requirements, utilize readily available standard media, support virtual networks, and make use of the public infrastructure.

How this information infrastructure will be implemented remains a question. The plans for ATM-based network services for RBOCs are still evolving. Some will have stable, but limited availability in the first quarter of 1997. In terms of service availability, the Internet Exchange Carriers are far ahead of the RBOCs. Although ATM switching products do not currently possess the standards and public infrastructures needed to support large-scale networks, carriers, standards groups, vendors are adding features and making progress daily. It is recommended that ATM switches and other solutions be examined again for use in the SeMCom upgrade, two years from now. ATM has the potential for meeting most of the network requirements, but another player in the arena will be gigabit Ethernet.

The capital costs of introducing a completely new technology, like ATM, are quite high. Whereas taking the existing infrastructure and upgrading it, like gigabit Ethernet, are quite low. There will be issues between the two, with high bandwidth on one side and quality of service on another. All of these issues must be thoroughly researched and evaluated. Gigabit Ethernet is still only in the development stages. It is today like ATM was five years ago, much more of a concept than a reality. Regardless, there are many factors to be considered.

Once a full effort is underway to achieve complete CIM interoperability, a more detailed evaluation could be justified. At that time, it will probably be necessary to enhance and reevaluate the technical performance measures so that alternative solutions can be compared and contrasted again using the multiattribute utility function as described in Appendix B. Additionally, to ensure a complete understanding of the operations, the systems process in Appendix C must also be reviewed for modification in the future.

In two years, SMC will expand their network by keeping the same technology, enhancing it, or pursuing another course of action. It will be important to start the systems engineering process again to find the solution that best meets the company's needs. For now, the proposed frame relay and ISDN backup solution will meet and exceed the specified requirements with a relatively lower cost than many of the alternatives.

7. Conclusions

A hybrid systems engineering process works well to develop this type of network communications infrastructure upgrade design. The conceptual design phase established high level direction to implement an architecture with routers and switches, supported by DNS, SNMP, and multiprotocol frame relay infrastructures. Project management, system operation and support, and requirements were created that detailed design goals and allocated resources for the effort. The preliminary design mapped requirements to the components, establishing an evaluation of alternative candidate solutions for the preliminary design and configuration. The detail design continued the development by providing detailed specifications and test plans.

SeMCom will focus on the exchange of information across organizational boundaries. It will be a framework that will foster the integration of existing and future architectures. In time, it will define information exchange services, CIM applications, types of media to be exchanged, standards, services, responsibilities of providers, and interface requirements. Providers and users will continue to be responsible for their own internal architectures and system implementations.

Fundamental improvements in the quality of communications systems have reduced the need for the extensive error-detection processing of X.25 and have placed more emphasis on raw transmission and switching speeds. To be effective, networks must redistribute network processing and utilize intelligent end-points.[15] Market drivers have contributed significantly to the emergence of new methods like frame relay.

Frame relay networks are superior to X.25 in a number of ways, including reduced interconnecting costs, lower overhead, increased performance, reduced network complexity, and increased utilization. By their very nature, WAN services must cater to a wide range of applications, such as voice, bulk data, bursty applications, and video conferencing. Frame relay is especially suited for applications where there are high throughput rates or bursts of data. Delay-sensitive applications such as voice and video are less suited, although some video applications that have low information content or video compression algorithms have been developed. In addition, frame relay for voice is also on its way.

Corporate applications requiring high bandwidth on demand, such as those found in SMC for manufacturing and engineering, are candidates for frame relay technology, so long as required response times are assured through good network design. In general, manufacturing and engineering firms can use frame relay to transmit large data files and images, such as those found in CAD, CAM, and CIM systems, quickly and efficiently.

The HP OpenView management environment provides a standards-based platform upon which a number of management applications can be integrated. This integration results in a common, intuitive interface that functions as a single console in a distributed, multivendor computing and network environment. The management environment is characterized by a coordinated set of applications and data stores. These applications comprise a complete integrated systems and network management solution, providing system monitoring, operations support, problem management, performance management, security, and change management. The entire tool set is configurable to optimize processes and procedures.

The desired situation of a faster, more robust, multiprotocol routed network is attainable. Even with the network management overhead requirements, following a reasonable design and roll-out plan will help to ensure minimal stumbling blocks.

From this project, it has been shown that the upgrade and management of a wide area network that meets SMC's requirements can be effectively provided. Through the implementation of the OpenView management platform, a robust and flexible network and operations center can be designed. This type of system supplies the performance and management capabilities that are desired within the constraints imposed by the CIO.

WORKS CITED

1. Black, Uyless. A Model for Computer Communications Standards. Englewood Cliffs, NJ: Prentice-Hall, 1991, pg. 77-80
2. Frisch, Aeleen. Essential System Administration. Sebastopol, CA: O'Reilly & Associates, Inc., 1995, pg. 616-617
3. Steedman, Douglas. X.500: The Directory Standard and Its Application. Twickenham, UK: Technology Appraisals, Ltd., 1993, pg. 10-22
4. Ibid., pg. 3
5. Lippis, Nick. Rockland, MA: Strategic Networks Consulting, Inc., 1995
6. Enck, John. A Manager's Guide to Multivendor Networks. Horsham, PA: Professional Press Books, 1991, pg. 95-107
7. Stallings, William. Data and Computer Communications 4th Edition. New York, NY: Macmillan, 1994, pg. 424-442
8. Taylor, D. The McGraw-Hill Internetworking Handbook. New York, NY: McGraw-Hill, 1995.
9. Lippis, Nick. Rockland, MA: Strategic Networks Consulting, Inc., 1995.
10. Motorola University Press. The Basics Book of ATM. Reading, MA: Addison-Wesley, 1994
11. Bradner, Scott. "Router Tests." Harvard University, 1996
12. Rose, M. The Simple Book: An Introduction to Management of TCP/IP-based Internets. Englewood Cliffs, NJ: Prentice-Hall, 1991.
13. Stallings, William. SNMP, SNMP v2, and CMIP: The Practical Guide to Network Management Standards. Reading, MA: Addison-Wesley, 1993.
14. Black, Uyless. TCP/IP and Related Protocols. New York, NY: McGraw-Hill, 1992.

REFERENCES

1. Black, Uyless. A Model for Computer Communications Standards. Englewood Cliffs, NJ: Prentice-Hall, 1991.

2. Black, Uyless. TCP/IP and Related Protocols. New York, NY: McGraw-Hill, 1992.

3. Blanchard, Benjamin S., Wolter J. Fabrycky. Systems Engineering and Analysis. Englewood Cliffs, NJ: Prentice Hall, 1991.

4. Bradley, T., C. Brown, A. Malis. "RFC 1490: Multiprotocol Interconnect over Frame Relay." Herndon, VA: InterNIC, 1993.

5. Bradner, Scott. “Router Tests.” Harvard University, 1996.

6. Case, J., K. McCloghrie, M. Rose, S. Waldbusser. "RFC 1450: Management Information Base for version 2 of the Simple Network Management Protocol (SNMPv2)." Herndon, VA: InterNIC, 1993.

7. Datapro Educational Services. “Telecommunications Project Management Notes.” Delran, NJ: Datapro Information Services, 1994.

8. “Directions in CIM for Semiconductor Wafer Fabrication.” Solid State Technology, February, 1994.

9. Enck, John. A Manager’s Guide to Multivendor Networks. Horsham, PA: Professional Press Books, 1991.

10. Frisch, Aeleen. Essential System Administration. Sebastopol, CA: O’Reilly & Associates, Inc., 1995.

11. Interlink Network Consulting Group. “Hands on Internetworking with TCP/IP.” Cary, NC: American Research Group, 1993.

12. Jennings, B. "RFC 1943: Building an X.500 Directory Service in the US." Herndon, VA: InterNIC, 1996.

13. Lemnios, Zachary. “Beyond MMST: The Virtual Factory.” Solid State Technology, February, 1994.

14. Lippis, Nick. Rockland, MA: Strategic Networks Consulting, Inc., 1995.

15. Malis,A., D. Robinson, R. Ullmann. "RFC 1356: Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode." Herndon, VA: InterNIC, 1992.

16. McClain, Gary R. International Connectivity Standards. New York, NY: Van Norstrand Reinhold, 1990.

17. Motorola University Press. The Basics Book of Frame Relay. Reading, MA: Addison-Wesley, 1993.

18. OpenView Forum. "Members Conference." Seattle, WA: June, 1995.

19. Rose, M. The Simple Book: An Introduction to Management of TCP/IP-based Internets. Englewood Cliffs, NJ: Prentice-Hall, 1991.

20. Schwartz, Mischa. Telecommunications Networks, Protocols, Modeling, and Analysis. Reading, MA: Addison-Wesley, 1987.

21. Semiconductor Industry Association. The National Technology Roadmap for Semiconductors. San Jose, CA, 1994.

22. Spohn, David. Data Network Design. New York, NY: McGraw-Hill, 1993.

23. Stallings, William. Data and Computer Communications 4th Edition. New York, NY: Macmillan, 1994.

24. Stallings, William. SNMP, SNMP v2, and CMIP: The Practical Guide to Network Management Standards. Reading, MA: Addison-Wesley, 1993.

25. Steedman, Douglas. X.500: The Directory Standard and Its Application. Twickenham, UK: Technology Appraisals, Ltd., 1993.

26. Tanenbaum, Andrew. Computer Networks, 2nd Edition. Englewood Cliffs, NJ: Prentice-Hall, 1990.

27. Taylor, D. The McGraw-Hill Internetworking Handbook. New York, NY: McGraw-Hill, 1995.

28. Waldbusser, S. "RFC 1757: Remote Network Monitoring Management Information Base." Herndon, VA: InterNIC, 1995.

Appendix A - Schedule

ID	Task Name	Duration	Start	Finish
1	**Headquarters**	**31.75d**	**6/24/96**	**8/6/96**
2	Program Plan Review	0.5d	6/24/96	6/24/96
3	**Network Implementation**	**23.2d**	**6/24/96**	**7/25/96**
4	**Project Initiation Review**	**5d**	**6/24/96**	**6/28/96**
5	**Deliverables**	**1d**	**6/24/96**	**6/24/96**
10	**Additional Requirements**	**4d**	**6/25/96**	**6/28/96**
14	**Assessment Phase**	**23.2d**	**6/24/96**	**7/25/96**
15	Meetings with Departments	1d	6/24/96	6/24/96
16	Interim Change Process	0d	6/24/96	6/24/96
17	Technical Meetings and Wa	0.8d	6/24/96	6/24/96
18	**Scheduling and Pricing**	**1d**	**6/24/96**	**6/24/96**
31	**Site Survey**	**22.7d**	**6/24/96**	**7/24/96**
52	System Requirements Revi	3d	7/5/96	7/10/96
53	Preliminary Design Review	1d	7/24/96	7/25/96
54	**Design Phase**	**9.6d**	**6/28/96**	**7/11/96**
55	Detailed Project Plan with Price	2d	6/28/96	7/1/96

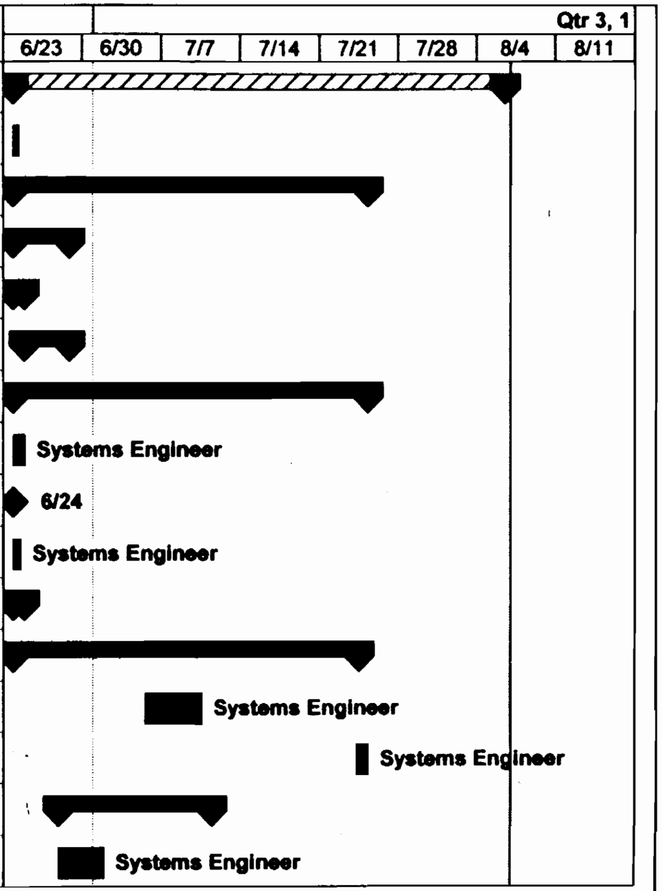

Project:
Date: 8/7/96

Task	■	Rolled Up Task	■
Progress	▬	Rolled Up Milestone	◇
Milestone	◆	Rolled Up Progress	▬
Summary	▬		

Figure 32: Master Schedule

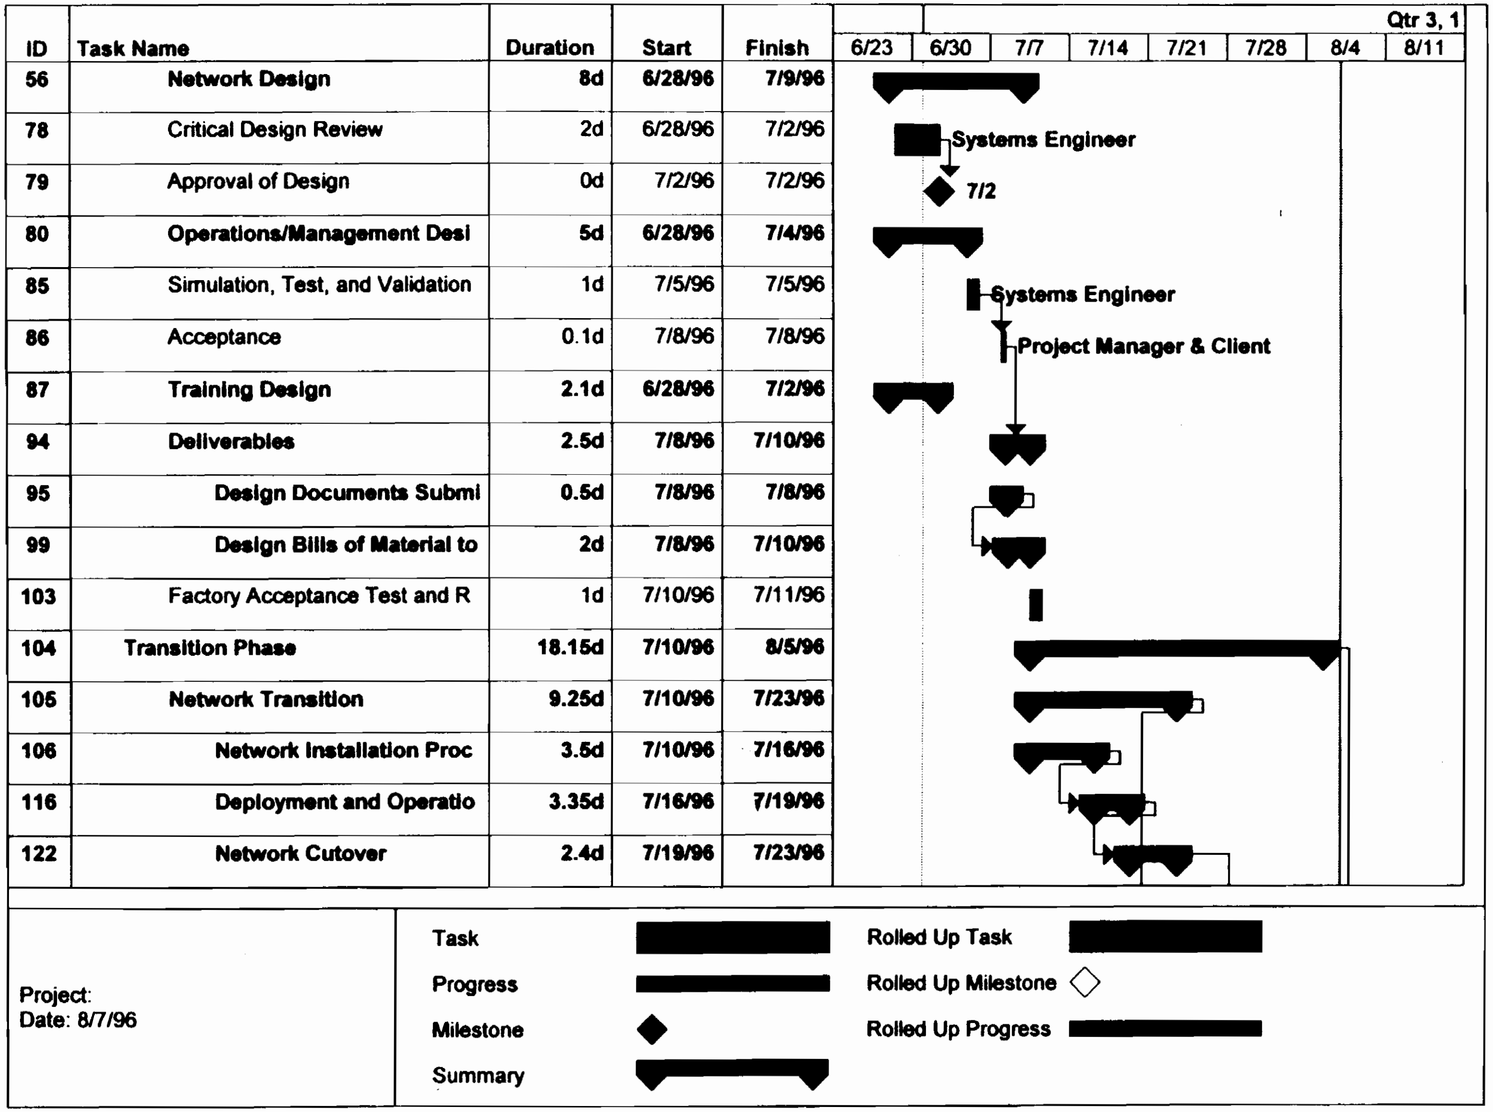

ID	Task Name	Duration	Start	Finish
56	**Network Design**	8d	6/28/96	7/9/96
78	Critical Design Review	2d	6/28/96	7/2/96
79	Approval of Design	0d	7/2/96	7/2/96
80	**Operations/Management Desi**	5d	6/28/96	7/4/96
85	Simulation, Test, and Validation	1d	7/5/96	7/5/96
86	Acceptance	0.1d	7/8/96	7/8/96
87	**Training Design**	2.1d	6/28/96	7/2/96
94	**Deliverables**	2.5d	7/8/96	7/10/96
95	**Design Documents Submi**	0.5d	7/8/96	7/8/96
99	**Design Bills of Material to**	2d	7/8/96	7/10/96
103	Factory Acceptance Test and R	1d	7/10/96	7/11/96
104	**Transition Phase**	18.15d	7/10/96	8/5/96
105	**Network Transition**	9.25d	7/10/96	7/23/96
106	**Network Installation Proc**	3.5d	7/10/96	7/16/96
116	**Deployment and Operatio**	3.35d	7/16/96	7/19/96
122	**Network Cutover**	2.4d	7/19/96	7/23/96

Figure 32 (cont): Master Schedule

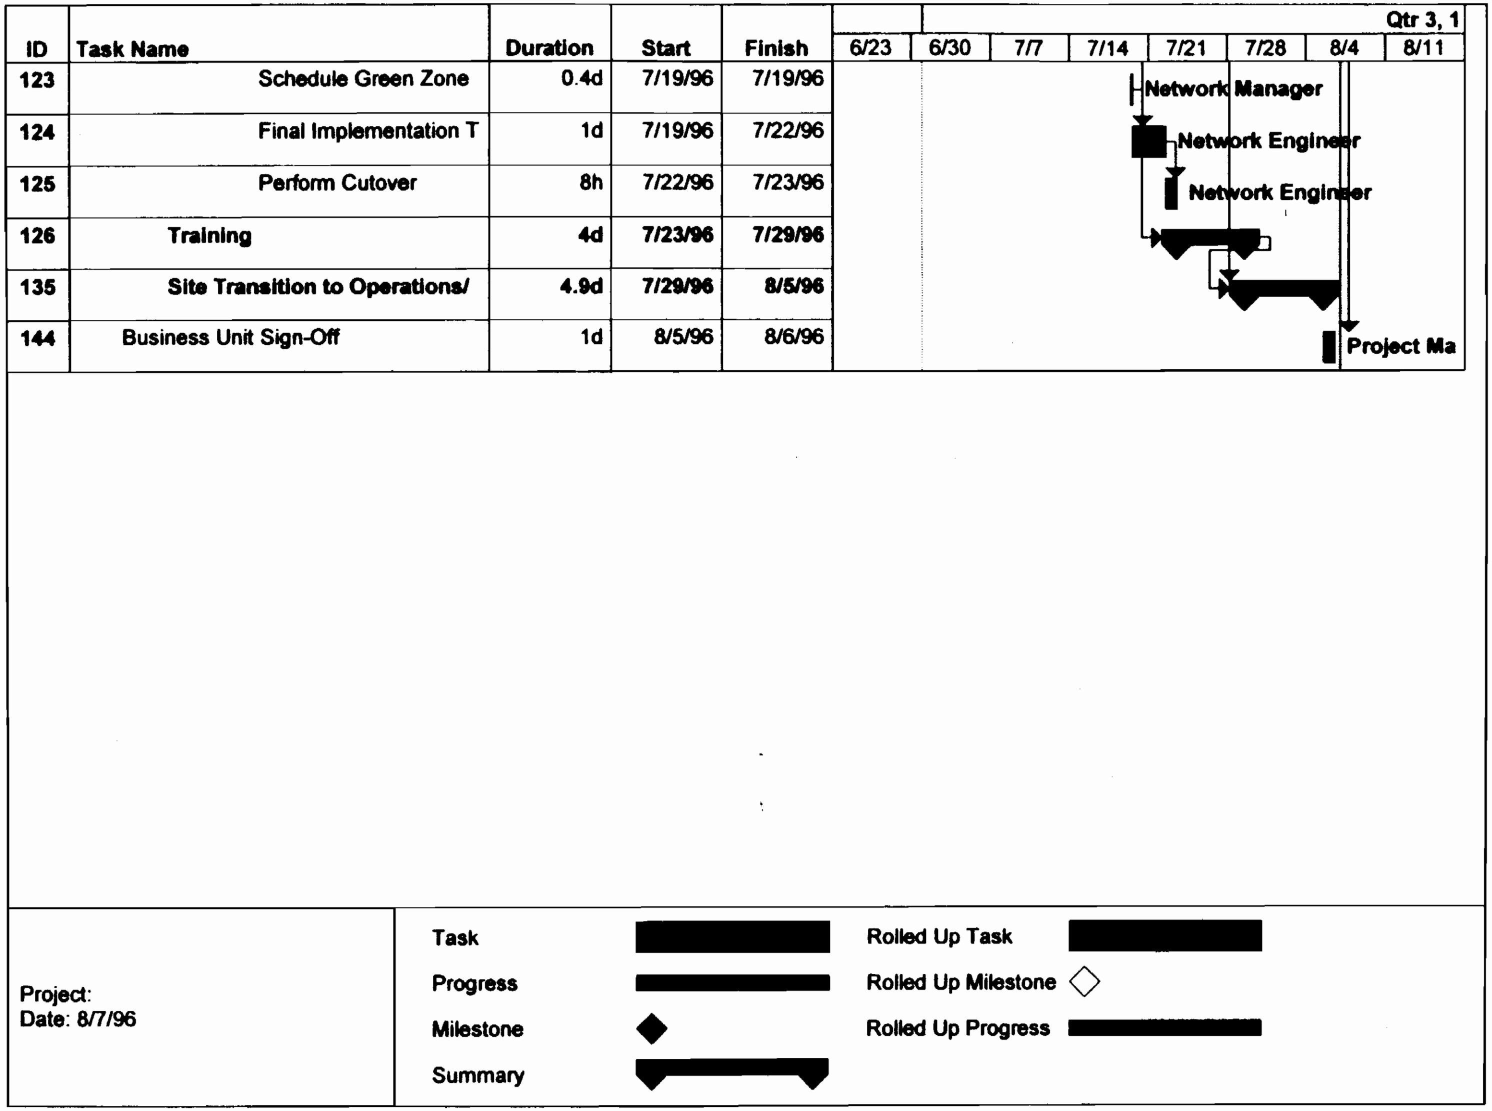

ID	Task Name	Duration	Start	Finish
123	Schedule Green Zone	0.4d	7/19/96	7/19/96
124	Final Implementation T	1d	7/19/96	7/22/96
125	Perform Cutover	8h	7/22/96	7/23/96
126	Training	4d	7/23/96	7/29/96
135	Site Transition to Operations/	4.9d	7/29/96	8/5/96
144	Business Unit Sign-Off	1d	8/5/96	8/6/96

Project:
Date: 8/7/96

Task	■	Rolled Up Task	■
Progress	▬	Rolled Up Milestone	◇
Milestone	◆	Rolled Up Progress	▬
Summary	▬		

Figure 32 (cont): Master Schedule

Appendix B - Ranking Alternative Solutions

Although the term ranking may conjure up the notion of simply comparing values, the process is more complex. Weighting indices takes into effect non-performance concerns that are not normally considered, including effects on non-users, impacts of transition on the existing system, sensitivity to variation, and acceptance levels.

In ranking the candidate solutions, it is necessary to apply the performance measures to the components. Scores were either taken from independent studies or product statistics. Most of the evaluation was done by researching different performance factors. After the values were calculated, they were normalized to a relative scale to make decision model calculations easier.

Since these are normative, multiattribute functions, the decision model used to evaluate these scores was a multiobjective utility model. The premise of the normative approach is to take the complex, multidimensional tasks and transform them into a structured sequence of simpler, unidimensional tasks. Parametric utility functions can be used to specify a utility of one given for the highest ranking candidate and a utility of zero to the lowest ranking candidate. From here, the scores are converted to utilities.

$$u_i(x_i)$$

where u is the utility function and x is the score for index i.

Once the utility has been established, tradeoffs are determined. The method of measurement can be directly and independently assessed. This is where the utility function is used. What makes this function so useful is that it can be expanded to as many values as necessary. It can be applied to two, three, four, or even more indices of performance, depending upon the requirements. Now the values of the tradeoffs can be substituted into the equation. Depending upon the number of indices, the equation can be solved in different ways. Three indices requires using the quadratic formula; four or more indices may require a software package to simplify calculations. In any case, the feasible answers will be in the range of {-1..1}. It is now possible to rank any element on the basis of these answers.