ILLINOIS STATE
UNIVERSITY

The *ISUnet*
Core Upgrade Project

June 1998

Scott Genung
Manager of Networking Systems
Office of Telecommunications

1 - Introduction

*1.1 - What is **ISUnet**?*

ISUnet is a cohesive collection of heterogeneous network environments and architectures providing a seamless integration of services for faculty, staff, and students at Illinois State University. A variety of hosts, servers, and clients use this network for a large number of applications. This environment allows for the introduction, preservation, and enhancement of electronic services for the entire University community. Although the services, products, and audiences can vary dramatically, all of these systems coexist and often times intercommunicate with each other.

Conceptually, ***ISUnet*** can be broken into 5 major components: 1) remote access or dialup which provides access to the residence, 2) public WAN (Wide Area Network) or Internet connectivity to give users access to Internet resources, 3) private WAN to provide remote networks access to campus resources, 4) CAN (Campus Area Network) or network core to provide exchange points between building networks, the WAN, and remote access systems, and 5) LAN (Local Area Network) or building network to provide desktop systems with network access. Each of these components has specific characteristics, services, and limitations. All of these areas are interrelated and combine to form the infrastructure of ***ISUnet***. When expanding or evolving this network, all of these elements must be reviewed and addressed to determine the impact of impending changes in response to growth or other needs.

To date, ***ISUnet*** has been extended to every major and most minor buildings on campus (approximately 60 locations). In the not so distant future, residence halls, University housing, and other locations will be connected to the campus network. In short, this network has come to service every major element of the University community.

*1.2 - What Does **ISUnet** Provide for the University?*

ISUnet is the network system that provides access to Internet, local server, peer server, mainframe, and debit card resources for the University. Virtually every department on campus depends upon one or more of these services for daily operations. Many departments cannot function without this connectivity.

As a result, ***ISUnet*** has become a strategic resource for Illinois State University. Today, faculty, staff, and students depend upon it as much as they do for telephone and electric service. Their needs are diverse and sometimes perceived as being independent of each other. To meet this expectation, this network must be extremely reliable and evolve or the services that it provides to the University community will suffer. Since this process occurs constantly, the infrastructure of this network experiences constant change.

1.3 - Constant Change

Modern networks have evolved into very complex entities. As a result, a significant amount of time and resources must be devoted to research to insure functionality, interoperability, scalability, evolution, and economy of scale. Without proper research, a poor choice of products can lead network planners to introduce an inadequate infrastructure that may ultimately lead to a network redesign.

The ***ISUnet*** of yesteryear was fairly simple to build. In the beginning, token ring and ethernet networks were constructed using a few types of copper cabling systems such as STP (Shielded Twisted Pair) and UTP-3 (Unshielded Twisted Pair, Category 3). For CAN applications, MMF (Multi-Mode Fiber) cabling systems were used between buildings. Non-intelligent, passive hubs were used to provide desktop connectivity. Bridges and routers were key to providing connectivity amongst these networks and the WAN (Wide Area Network). Changes in these technologies took place over several years. Once in place, vendor standardization solidified the adoption of the technology in the marketplace. Applications were simple and used network resources sparingly. In addition, users were typically statically located on the network.

Today, ***ISUnet*** is much more complex and volatile. Token ring and ethernet continue to be in use with most of the original cabling plant still in place. However, newer cabling systems have emerged with the promise of delivering higher performance connectivity to the desktop. This includes both high performance UTP-5 (Unshielded Twisted Pair, Category 5) for desktop applications and SMF (Single Mode Fiber) for CAN applications. In addition, hubs have become intelligent and more feature rich. Using high performance switching technologies shared media domains have decreased in size to reduce competition for bandwidth locally and improving overall throughput. To provide dedicated connectivity, desktop switching has emerged. Changes in these technologies occur over a few months and offer little vendor interoperability for advanced features such as vLAN (Virtual Local Area Network), multicast, and management technologies. Users have also become much more mobile and demand access from a greater variety of locations. This includes mobility from work, to home, to the Internet. Users want consistent access to resources regardless of how and where they are attached to ***ISUnet***. The number of applications in use has also grown dramatically. The consumption of network resources by these applications has become intensive.

The future of networking at Illinois State will become even more complex. Again, much of the original architectures will continue to exist but will start to diminish over time in favor of emerging connectivity solutions. Newer technologies such as desktop ethernet switching will become predominant. At the same time, a greater percentage of the cabling plant will migrate towards UTP-5 and SMF cabling systems. Newer switched media technologies will begin to enter into service on ***ISUnet*** such as switched 100BaseT (fast ethernet), gigabit ethernet, and ATM (Asynchronous Transfer Mode). Switch interoperability and the convergence of LAN switching, routing, and ATM makes research very critical to the evolution of networking within ***ISUnet***. Changes in these technologies can literally take place over a few months. Vendor interoperability doesn't truly exist yet with many of these technologies since the marketplace has moved too fast for standards bodies such as the IEEE, ANSI, and ITU to keep up. However, emerging applications demand vast amounts of network resources. Throughout the industry, network administrators have been struggling to find products that meet these requirements without deploying products that fall short of the need. Mobile computing has also become paramount. Some experts believe that the proportion of desktop computers will decrease in comparison to the rising numbers of portable platforms. This means more mobility on campus (classrooms, residence halls, labs, and offices) and off campus (analog dialup and the Internet). Applications will become even greater consumers of network resources.

1.4 - Objectives

The objective of this document is to describe the existing network infrastructure of ***ISUnet***, address the need to evolve the existing core environment, and define what the new architecture would look like. The process used for selecting specific products and the implementation procedures have also been defined.

To meet the challenges defined within this document, substantial research was required to provide direction and to translate needs into technologies. Without such efforts, ***ISUnet*** would not evolve and will eventually become an ineffective resource for the University. Now more than ever, meaningful dialogue between those who provide network resources and those who consume them are paramount to accomplish these objectives.

2 - The Existing ***ISUnet*** Core Architecture

2.1 - Technologies

Although other architectures such as localtalk and ethernet have emerged within building networks over the years, the core of ***ISUnet*** has been predominantly based upon token ring products. Currently, token ring can be implemented using two different data rates. The initial implementation in the late 1980's was based upon a rate of 4Mb/s. In subsequent years, 16Mb/s token ring emerged. All token ring within the ***ISUnet*** core is based upon 16Mb/s token ring technologies.

Token ring uses a scheme called token passing to manage bandwidth contention. In this architecture, only the station that possesses a free token may transmit at any given time. On a heavily contended network, token passing is an excellent scheme of bandwidth management. However, the existing implementation of token ring within the core is based upon shared media. This concept describes the notion of bandwidth moderation between peers on a common network. As such, bandwidth consumption will be bursty and difficult to predict. If extended periods of high utilization occur, applications will experience diminished performance and possible failures.

2.2 - The Original Core Network

The original core was composed of seven token ring fiber trunks that attached to a common ring using PC based bridges. Each trunk was based upon 16Mb/s shared media token ring that was constructed using IBM 8220 Fiber Optic Repeaters and MAUs (Multistation Access Unit). As many as ten building networks were attached to each trunk using PC based bridges.

2.3 - The Current Core Network

As ***ISUnet*** evolved, the need for replacing bridges within the core with high performance, wire speed switches became apparent. By the fall of 1995, the core bridges were overwhelmed with traffic volume and replaced with a pair of token ring switches. When this occurred, an eighth trunk was created. Over time, two more switches were added to provide dedicated connectivity for the central servers in Julian.

During the summer of 1996, a high performance core router was installed within the network core. This router improved frame propagation, enhanced the scalability of the campus subnet architecture, provided for the aggregation of Internet circuits, served as the instrument of security implementation within the networks of ***ISUnet***, and allowed the migration from vector based to link state routing protocols to occur.

By the summer of 1997, the core again became saturated. To resolve this issue, a 100Mb/s FDDI (Fiber Distributed Data Interface) ring was constructed among the 4 token ring core switches. A second core router was added to distribute core routing services for the building networks, WAN, and dialup. The number of trunks also increased from 8 to 12 to further distribute bandwidth contention.

2.4 - Building Network Integration

By 1993, ***ISUnet*** was experiencing broadcast traffic levels that consumed as much as 50% of the total traffic volume. As a result, the need for replacing source-route and translational bridges used to provide connectivity between the building networks and the core with routing products emerged. These routers provided more control than the displaced bridges through broadcast management, caching, and proxy services as well as supporting bridging methods for non-routable protocols. This architecture reduced the volume of broadcast traffic propagation within the building networks and within the core which resulted in reclaimed bandwidth and enhanced scalability and throughput.

By the fall of 1995, projects were initiated to collapse shared media token ring networks into multiple segments using token ring building switches. This architecture reduced local bandwidth contention while providing high performance frame propagation between multiple segments. This model also reduced the impact of network faults. Shared media ethernet networks were collapsed in a similar manner. However, these environments were not integrated directly. Connectivity between these networks was typically provided through the network core using building routers.

2.5 - Internet Connectivity

During the summer of 1997, a second core router was deployed to enhance IP based connectivity within ***ISUnet.*** Although all Internet connectivity continues to be provided by a single core router, enough redundancy is now available to allow either router to be used for Internet services in the event of a critical failure of any core router.

By January of 1998, four circuits provided ***ISUnet*** with Internet connectivity from three different ISPs (Internet Service Provider). All of these circuits were attached to a single core router with each device attached directly to a network trunk interface on separate switches within the core. A second FDDI network was constructed between the two core routers in an effort to enhance to provide an optimal path for the exchange of traffic between these routers and provide a means of survivability in the event of a partial outage.

In order to offer Internet connectivity, a process for resolving Internet host names to IP addresses is required. This service is called DNS (Domain Name Server). By May of 1997, the primary DNS was moved to a dedicated mid-range Unix platform with the secondary DNS on a high-end Unix server in Julian.

2.6 - Remote Access Connectivity

The current dialup system on campus marks the third generation of products for Illinois State University. With more than 330 channels in service, bandwidth consumption can reach high levels during peak periods. In an effort to address this problem, the dialup networks are attached to a core router that minimally impacts overall bandwidth consumption within the core. A dialup user who wants to connect to an Internet resource will never impact the central core directly. This traffic passes through an FDDI ring between the two core routers.

2.7 - Deficiencies of the Existing Core

Since the beginning of ***ISUnet***, the network core was based upon token ring technologies. This created three major problems: 1) integration, 2) scalability, and 3) longevity

The integration of dissimilar technologies within the existing core is a problem due to the inherent differences between token ring and ethernet technologies (ie: bit ordering, framing formats, frame lengths, etc). The most obvious problems appear when translational bridging is employed. Currently, ***ISUnet*** experiences periods of high volume broadcasting as the resulting of chatty NetBIOS/NetBEUI based applications such as Windows 95 and NT. When translationally bridged, transparent caches within a local router, ethernet building switch, or ethernet access switch can be quickly overwhelmed as this traffic is propagated from the interior of the network to the edge. Since the MTU (Maximum Transmit Unit) on token ring is larger than ethernet, frame fragmentation can occur which decreases overall network performance because large frames must be split into several smaller ones.

Scalability within the core also poses challenges. There has never been any technology developed for token ring environments beyond FDX (Full DupleX) token ring (32Mb/s). As a result, scaling a large core infrastructure has been challenging. FDDI becomes an option only if products consistently implement frame formats and bridging schemes. Since no standard for token ring framing over FDDI exists, this is a major issue. Few options exist for mixing these implementations since solutions are vendor specific.

Translational bridging issues and core network technology integration difficulties also severely limit the integration of large, high performance ethernet networks within the current core. These problems dictate the need to evolve ***ISUnet*** beyond its current form. As bandwidth demands increase at the desktop, so too will the demand increase proportionately within the backbone. With few options for enhancing the backbone, little flexibility exists for evolving the core using current products. The only real opportunity for improvement appears to be through a redesign of the core.

An area of even greater concern is that token ring is an architecture that is perceived by many within the industry as a dead technology. In the last few years, very few products have emerged within the market include token ring technologies. Most experts agree that token ring will continue to diminish in scope and importance over time. Despite half-hearted efforts from only a few vendors within the market, there will be no resurrection of this technology.

3 - Needs Dictate Change

3.1 - Background

From the beginning days of ***ISUnet***, the University required a basic set of services which included resource sharing, mainframe emulation, and basic Internet services (such as mail and file transfer). Today, additional application requirements have begun to emerge including newer, more powerful Internet services (such as the web, network news, and on-line chat sessions), peer-to-peer networking, networked multimedia (such as LAN based video), distributed services, client-server, the debit card system, and document imaging. These additional requirements have developed rapidly over the last few years and have had a substantial influence on the growth of this network.

Along with new application requirements, newer means for extending these services to the end user are being found through enhanced remote access systems, expanded Internet connectivity, residence hall networking, mobile computing, and smart classrooms. All of these ideas focus on providing greater convenience for the user to access services on ***ISUnet***.

With these new emerging applications and services, the existing network infrastructure for ***ISUnet*** is becoming overburdened. Current applications require more network resources than applications in the past. Concern has already been raised about the responsiveness of the current environment. Certainly, the introduction of more users as well as the introduction of additional applications will only exacerbate this problem.

When trying to evaluate network capacity, it is important to determine what applications will be used on the network as well as when and where these resources will reside. Since the design of ***ISUnet*** offers anyone the ability to connect to virtually anywhere at anytime, it is next to impossible to accurately predict what network resources will be needed. Therefore, this network architecture needs to be flexible enough to allow bandwidth and supporting services to be selectively introduced.

These emerging requirements dictate that ***ISUnet*** must be flexible enough to meet the challenges that these new services pose. If these issues are not considered and not acted upon, the current model of ***ISUnet*** will become inadequate for current and future service requirements. If this should occur, this network will not be able to provide the resources that the University requires. With appropriate steps, ***ISUnet*** will continue to meet the requirements of the University for many years to come.

3.2 - Growth

ISUnet has seen dramatic growth rates in the past few years. In the spring semester of 1997, the remote access system serviced more than 750,000 connections. By the end of the fall semester of 1997, over 1,000,000 connections were serviced. During the spring semester of 1998 on a busy day, the dialup supports over 21,000 connections. This is an increase of more than 100% over the previous record-breaking peak of 10,000 connections in the spring semester of 1997. As of December of 1997, there were 282 channels in the production dialup pool. By mid February of 1998, 48 additional channels were introduced using a second pilot number with a 30-minute connection limit while the original pool retained a 6-hour limit. At the same time, X2 based 56K technology was introduced into both pools. Plans are already underway to develop a new pilot dialup project to test the new V.90 (ITU standards based 56K technology) microcode. When placed in production, an additional 144 channels will be introduced. These resources will not be available until the beginning of the fall semester of 1998. Current traffic studies show that even these plans are not aggressive enough to accommodate projected growth.

Internet connectivity is provided through a series of circuits using various ISPs (Internet Service Providers) to optimize traffic management. Prior to December of 1996, a single T1 circuit (1.54Mb/s) was used to connect ***ISUnet*** to the Internet through MCI. At that time, this bandwidth was adequate for the needs of the University based upon snapshot data. By December of 1996, a second T1 circuit (aggregate of 3.08Mb/s) was installed through MCI for route diversity and improved capacity. During this period, Internet circuit utilization rose by almost 300% resulting in overconsumption of these circuits. In July of 1997, a third T1 circuit (4.62Mb/s) was introduced from Uunet. During the 1997 fall semester, Internet consumption again rose. A fourth T1 circuit (6.16Mb/s) was introduced from Sprint during December of 1997 and again Internet consumption rose in response to the newly available bandwidth. Projections for the fall semester of 1998 now suggest that this capacity will need to be doubled (12.32Mb/s) just to maintain services at the current level.

ISUnet has grown very rapidly since its inception in 1988. Unlike the rather constant 3-5% annual growth rate of the University's telephone system, the campus network has been growing at an annual rate of 25%. By January of 1998, approximately 9,000 nodes were connected to this network with approximately 5,000 of these nodes in use at any given time. By June of 1999, this growth rate is expected to reach as high as 47% due to the impending introduction of new residence hall networks. Other factors such as the introduction of the new Internet based undergraduate curriculum and continued expansion and upgrading of systems within computer labs, classrooms, and office spaces will likely continue to drive these growth rates in the years to come.

Within the ***ISUnet*** core network (CAN), traffic volumes during September of 1997 were more than double that measured during the same period of the previous year. During the fall semester of 1997, up to 50GB (GigaBytes) of data traversed the core network each day with peaks as high as 70GB. By the end of the 1997 fall semester, new records were being set weekly. As a result, there were times when the existing core infrastructure would become overburdened and applications would fail. Although this growth has dictated the need for major enhancements to the CAN, technologies have only recently emerged that can be implemented to allow the network core to scale to accommodate anticipated growth.

3.3 - Newer, Better, More

Growth is only one component of need. Newer and better applications also drive the process of evolution.

In 1993, a University academic strategic planning committee was formed to develop a long term plan for academic computing needs at Illinois State University. This committee was charged with developing a unified vision for end user computing needs for the next several years. This was accomplished by obtaining a campus consensus of these needs from campus faculty, staff, and students through questionnaires and interviews. The result of the work of this committee was to define a set of goals and recommendations to meet these needs by the year 2000. The document that was developed became known as the *Vision 2000 Strategic Plan*.

The *Vision 2000* plan describes a set of basic goals that must be met in order to accommodate the growing needs of the University for instructional purposes. Although the *Vision 2000 Strategic Plan* was produced in May of 1994, it can still be used as a tool for projecting academic needs of the University. These needs can be summarized into the following (most of these are taken directly from the Vision 2000 planning document):

1) "All faculty, staff, and students will be conveniently connected to ***ISUnet***."

2) "***ISUnet*** will be able to deliver the information that the campus community needs for teaching, learning, research, and administrative activities. This includes full digital communication including video, audio, graphics, and text transfers."

3) "State of the art research computing environments will be made available on ***ISUnet*** with connectivity to national supercomputing facilities through the Internet."

4) “The University will provide support for research which takes advantage of computer and networking technology.”

5) “Students will be strongly encouraged to own their own personal computers.”

6) “All members of the University community (faculty, staff, and students) will be able to access resources available on ***ISUnet*** and Internet from on or off campus. Faculty and staff members will have remote access to lab servers for software maintenance or file editing.”.

7) “The capacity of ***ISUnet*** will be increased to meet the demands of the end user.”

8) “Smart” classrooms will be developed with access to ***ISUnet*** resources. These classrooms will need to “accommodate multimedia technology designed to support instructional delivery and foster collaborative learning experiences.”

9) “***ISUnet*** will integrate portable computing and wireless communication.”

10) “Distance learning classrooms will be developed to provide for the maximum amount of teacher and student interaction.”

3.4 - Residence Hall Networking

In 1994, the concept of residence hall networking emerged on campus as the result of a renovation of the long-standing general education curriculum. This new curriculum emphasizes the use of computing technologies within ***ISUnet*** and the Internet. The resulting change has enhanced the role of communication technologies within the student experience.

In the spring of 1995, a pilot residence hall network was deployed in Smith House within Watterson Towers. This pilot consisted of approximately 220 connections based upon shared media ethernet. In the spring of 1997, this network was upgraded to switched ethernet and became a pilot environment for continued testing of equipment that would be used in subsequent phases of the residence hall networking project.

In the spring of 1997, funding was obtained to develop the first phase of this project. The first phase will address the residence hall network needs of the South Campus (Hamilton, Whitten, Atkin, and Colby) and the West Campus (Haynie, Wright, and Wilkins) residence hall complexes. The capacity of these two complexes totals about 2,800 network connections. With a budget of $2 million, it is uncertain if all of the buildings within these complexes will be provided with connectivity. Subsequent phases of this project in other campus residence hall complexes may follow in future years. The capacity of the five residence hall complexes on campus is more than 7,800 students. University owned apartment complexes (such as Shelbourne and Cardinal Court) are also being reviewed as potential locations for future connectivity projects.

3.5 - Distance Education

Distance education is a term that describes an environment where learning is independent of space. To accomplish this objective, technologies such as those provided by the Internet or distance learning applications have been implemented.

The level of interest in this educational methodology has been growing steadily over the last few years. Both on campus and within the State of Illinois, projects are being developed to enhance distance education applications and environments to make these services more flexible and universal.

4 - Requirements of a New Core Network

4.1 - Background

The evolution of ***ISUnet*** is dictated by the needs of those who consume network resources. As a result, requirements are very specific to current and projected needs. Those needs defined in the previous section are only a synopsis of the complete lists of variables that drive this evolution.

In order to construct a new core network, the objectives of this new infrastructure must be clearly defined and related to need. Each of these elements can then be used to define the requirements. Throughout this section, the objectives of the core upgrade project are defined along with the motivations for each element.

4.2 - Seamless Integration into the Original Core

Regardless of what technologies are used within the new core, seamless connectivity between old and new is paramount. This need exists because clients will ultimately reside on networks attached to multiple environments while major services (such as DNS) will ultimately reside within the new core network. As such, any network resource must be available to any client that resides anywhere within the ***ISUnet*** infrastructure. In addition, movement from the old core to the new should also have minimal impact on either host or client.

4.3 - Scaleable and Selective Introduction of Bandwidth

Bandwidth is an expensive commodity. As such, it should be treated as an important resource and only be introduced where and when applications demand it. High performance network architectures should not be deployed just for the purpose of upgrading the network. Instead, these technologies should be introduced in response to existing and anticipated needs of the user at reasonable intervals. As a result, a very flexible methodology of introducing bandwidth is needed.

At the same time, ***ISUnet*** must be ready to provide high performance networking on an increasingly larger scale as the demand for it grows. This means the rapid introduction of high performance concentrators and workgroup switches into building networks that attach to core network electronics where needed while concurrently providing connectivity for legacy environments. This will allow the network to evolve purely and selectively upon the basis of need.

As bandwidth demands grow, the implementation of scalable uplink architectures will be critical. Multiplexing technologies will offer an attractive means for accomplishing this objective by providing for logical bandwidth aggregation.

As always, it makes sense to provide high performance connectivity to servers and select clients. It does not make sense to provide enhanced connectivity to those who use their systems for only basic services. This concept allows the University to maximize the financial and technical network resources where and when they are needed.

4.4 - Preservation of Existing Token Ring Technologies

The preservation of existing token ring technologies is essential to implementing a successful migration plan of legacy technology into the next generation network. It is equally important to optimize this existing technology within the new infrastructure but at a minimal investment. This process allows those users who must continue to reside on token ring to do so with no degradation of service into the indefinite future. As a result, the need to move from legacy token ring to switched ethernet technologies is dictated by the user and not by the network architecture.

4.5 - Integration of Token Ring and Ethernet Technologies

It is no secret within the network industry today that ethernet has become the predominant network architecture. With shared and switched 10Base, 100Base, and gigabit ethernet, other technologies (such as switched token ring, FDDI, and fibre channel) cannot offer the richness in connectivity at the cost that ethernet can. As a result, ethernet technologies will become the network architecture of choice for connectivity within ***ISUnet*** and the world over the years to come.

To provide connectivity for these users while at the same time preserving connectivity for legacy environments, a seamless integration of these architectures is required. This allows users to mix connectivity between token ring and ethernet architectures as needs dictate without impacting the quality of connectivity for any group of users. This includes but is not limited to such architectures as policy based vLANs, DHCP, and IP multicasting. To accomplish these tasks, integration needs to occur within a common set of network electronics. This allows for the universal implementation of services and maximizes connectivity between legacy networks.

4.6 - Integration of Voice, Video, and Data Services

As a greater dependence upon Internet services has emerged, more rich and functional applications have appeared that integrates multimedia objects into a common interface. Products that are as simple as hypermedia enabled web pages and as complex as video conferencing and collaboration tools have become desirable for many users. These services consume vast amounts of network resources. However, the current network environment is not designed to provide these services.

Conventions for the integration of data, voice, and video services over a common architecture are necessary for the success of the implementation of multimedia applications on a large scale. An architecture based upon QOS (Quality of Service) is required to provide these services qualitatively. This can either be based upon cell based AAL (ATM Adaptation Layer) classes, frame based RSVP (ReSource reserVation Protocol), or a combination of the two.

To optimize video traffic over frame based data flows, an IP multicast architecture must be implemented. This architecture must be flexible so that multicast tunnels can be extended anywhere on ***ISUnet***. To provide connectivity to Internet based multicast tunnels, MBone (Multicast Backbone) support will be required from one or more ISPs (Internet Service Provider).

Current video applications are designed around room based configurations that connect to similar environments in other locations. Desktop-based video applications take advantage of network connectivity. To integrate both types of services, video gateways will be needed.

4.7 - Enhanced Internet Connectivity

As the CAN and LAN environments of ***ISUnet*** evolve, measurements must be made on the impact of the introduction of these enhanced environments upon the existing WAN Internet architecture. From a traffic management perspective, it has become quite obvious that continuing to add single T1 circuits in a vane effort to provide enhanced bandwidth aggregation falls short of the goal of adequate Internet capacity. Higher capacity connectivity must be pursued since applications can quickly consume T1 pathways.

At Illinois State University, most WAN based Internet services are not mission critical. However, some applications provide critical services for the University such as those used by Milner Library. These services rely on systems that reside at the University of Illinois campuses.

While the need for traditional Internet services (such as the web, email, and network news) continue to be reviewed, thought must also be given to the inevitable introduction of I2 (Internet 2) based services (such as distance education). Again, the current WAN model is incapable of providing I2 services. The introduction of I2 requires ATM connectivity over the WAN using physical circuits that are DS3 or better.

4.8 - Address Management and Mobility

In an effort to promote mobility, reduce the complexity of address management for the user, and provide tracking for network administration, a new address management scheme needs to be introduced. This scheme will employ policy based vLAN (Virtual Local Area Network) technologies that map statically or dynamically obtained IP addresses to virtual workgroups. Using virtual trunking technologies, virtual routes for these workgroups can be announced throughout the network. This allows users to maintain membership to defined or acquired IP subnets throughout the entire network. Traditional interior routing protocols will continue to be implemented to concurrently provide clients with access to resources located within the new and legacy core.

4.9 - Proactive Resource Management

With ***ISUnet*** evolving in size, complexity, and depth, it is more important than ever to continue developing a proactive process to managing and evolving the network. Proactively responding to issues on the network decreases the occurrence or impact of network faults and the time needed to resolve them. This is critical since the dependence upon network resources has risen to a very high level across campus and because this network has become so large and complex.

To accomplish these goals, more network tools based upon SNMP (Simple Network Management Protocol) and historical RMON (Remote MONitoring) must be introduced. To provide flexibility, web based management tools, network topology maps, and resource databases must become the product suite for those who maintain the campus network.

4.10 – IP Centricity

In an effort to provide universal access of services such as file and print sharing and mainframe emulation, these applications must migrate towards IP (Internet Protocol) based connectivity. In this environment, these services would be available to clients throughout the network and the dialup with OS (Operating System) independence. This configuration would also reduce workstation and network complexity by reducing the number of protocol suites in use.

5 - The New ***ISUnet*** Core Architecture

5.1 - Technologies

In order to develop a core network architecture that meets the requirements described in the previous section, a radically new core network needs to be constructed. This new environment must contain many new technologies to meet these objectives and to position ***ISUnet*** for the introduction of future technologies and applications.

The first issue to address when developing this new network architecture is to define what technologies will be used. Within the core itself, there are three potential technologies to choose from: 1) FDDI (Fiber Distributed Data Interface), 2) gigabit ethernet, and 3) ATM (Asynchronous Transfer Mode).

FDDI is a shared media, variable length frame-based architecture with a speed of 100Mb/s using token passing technology for bandwidth contention. As mentioned in the section discussing deficiencies within the existing core network, FDDI lacks the flexibility to scale or integrate with a mixed legacy environment. It also lacks any QOS (Quality of Service) features that would allow the core network to mix multiple applications such as data, voice, and video.

Gigabit ethernet is based upon CSMA/CD (Carrier Sense Multiple Access / Collision Detection) for bandwidth contention with a speed of 1Gb/s. Since it is based upon the IEEE 802.3 standard using variable length framing, integration into legacy ethernet environments is relatively simple. With 10BaseX, 100BaseX, and 1000BaseX, ethernet offers a high level of bandwidth scalability. However, gigabit ethernet is still immature and lacks any universal specifications for load balancing over parallel links. Like FDDI, gigabit ethernet continues to lack QOS (Quality of Service) features. However, RSVP (Resource reSerVation Protocol) is an emerging IETF standard that will be implemented in the not so distant future to partially address this deficiency.

ATM is an isochronous architecture based upon 53 byte cells. Through a combination of service classes and predictable latency with fixed length cells, QOS can be implemented for mixing multiple, dissimilar applications over a common infrastructure. ATM is also the only scalable architecture from 25Mb/s to OC-48 (2.48Gb/s) and beyond. This environment provides for the integration of voice, video, and data applications in very diverse environments. Legacy data integration is provided through LANE (LAN Emulation) and MPOA (Multi-Protocol Over ATM) services. This mechanism allows for the introduction of a cut-through routing scheme to distribute virtual routing services on multiple devices. These features along with the eventual need for ATM based QOS beyond RSVP suggest that ATM is the best architecture for a new core network since QOS is an important part of the future requirements of ***ISUnet***.

5.2 - The New Core Network

The new core network is based upon the concept of a distributed environment. This means that no single network device or geographical area represents the core itself. Instead, the core of the network actually resides in multiple locations using many devices to insure survivability in the event of a major outage. This architecture also lends itself to scalability since resource consumption can be distributed.

To distribute services within this model, this new environment is comprised of multiple layers. The new core network segregates ***ISUnet*** into 3 tiers: 1) ATM core, 2) edge switch, and 3) building network.

The ATM core is composed of multiple high performance, cell-based switches. These switches exchange traffic at very high speeds in a mesh that provides survivability in the event of a failure. To accomplish this objective, four ATM core switches will be collocated within a telephone switching center or data center. These spaces have been selected since fiber and environmentals are already provided and integration of both services over a common medium is convenient.

The edge switch is a product that provides connectivity between the ATM core and the building networks. For ATM connectivity, a high performance point-to-point link is constructed with multiple ATM core switches. This architecture provides survivability in the event of a failing core switch. Building networks are attached to an edge switch using high performance, FDX (Full DupleX) uplinks. A frame-based load-balancing scheme will also be implemented to allow for high performance, scalable connectivity of switched environments through the logical aggregation of multiple links. Routing services with silicon based cut-through routing logic will displace existing packet-by-packet routers. Using data flows, this architecture will improve the scalability of IP subnet routing on ***ISUnet***.

5.3 – DNS, DHCP, and vLANs

In select environments (such as the residence hall networks), the use of DHCP (Dynamic Host Configuration Protocol) allows users to request addresses dynamically to prevent address duplication, protect reserved address pools, provide audit logging, and simplify workstation definitions. Using extended lease periods, the client can maintain this address for long periods of time such as a semester or more. If the client is disconnected from the network for an extended period, the served address is returned to the pool and made available for other clients to request.

To provide these services, the existing campus DNS (Domain Name Server) architecture must be reengineered because DHCP must coexist with this service to provide dynamic hostname updates when needed. The other limitation of the existing DNS architecture is that all hostname entries are static.Using DDNS (Dynamic DNS), the primary DNS tables can be simplified to include only those hostnames that must be statically defined. The remaining clients that do not need static entries can use dynamic hostnames with DHCP.

To provide a scalable and survivable DNS/DDNS/DHCP architecture for ***ISUnet***, two Unix based systems need to be constructed. The first system would serve as the primary DNS and would displace the existing primary (138.87.1.10) with no impact upon the user. The second system would serve as the secondary DNS (138.87.1.11) and would require the phaseout of the secondary DNS (138.87.1.2). DHCP would coexist on both systems. Relay agents running on building routers and edge switches would evenly distribute address requests to these DHCP servers for address pool management. In the event of a failure, the other server would be defined as a backup through failover relay agents.

To render user mobility and provide for the convergence of underutilized IP address spaces, a combination of DHCP, LAM (Local Area Mobility), and policy-based vLAN (Virtual LAN) architecture must be introduced. Using this architecture, vLANs will be directly mapped to select IP subnets to allow these networks to be announced throughout ***ISUnet***. To provide mobility, this process allows select users to connect to the network using the same IP address regardless of their proximity. To provide convergence, multiple statically assigned subnets defined to a single department distributed across campus can be converged upon a single logical subnet. This process better optimizes IP address spaces on campus, reduces the number of defined subnets, and provides for a greater number of available addresses as the network continues to grow.

5.4 – Residence Hall Networks

Within a residence hall network, each student will be provided a dedicated network jack. This configuration allows multiple students sharing a room to have access to the network through separate ports. Using switched ethernet for connectivity (ie: dedicated 10Mb/s for each desktop), students can bring a variety of different hardware and software products to campus and obtain connectivity to the campus network for minimal cost at relatively high performance. This architecture also enhances network security since unauthorized packet sniffing will not be able to decode secure information residing within unicast frames on a switched network.

Residence hall networks require large concentrations of high-density switches to provide desktop connectivity. These switches attach to a common edge switch using aggregated fast ethernet uplinks. Because of this aggregation, a high performance connection is required into the ATM core using OC-12 or better.

5.5 - Building Network Integration

Once portions of the new core network infrastructure have been introduced, building networks can be rehomed to select edge switches to provide enhanced connectivity. In the beginning, this new connectivity will be in support of Internet based services. As such, the building routers connecting these networks to the existing core will remain in service to support other applications. Over time, these applications will migrate towards Internet protocols (ex: WINS for Windows 95 and NT workgrouping and TN3270E for mainframe emulation). Once this convergence has been completed, the original core will be phased out and all connectivity will be provided through the new infrastructure.

Most legacy networks on ***ISUnet*** will be connected to the new core using high performance uplinks. High-density token ring environments will be connected to designated edge switches from token ring building switches using proprietary technologies that optimize edge switch integration. Low-density token ring networks will be rehomed to select sites using FDX (Full DupleX) token ring. Existing ethernet environments will be connected to designated edge switches using aggregated fast ethernet uplinks from ethernet building switches.

One of the major objectives of the new core network is to provide connectivity to legacy network environments while concurrently supporting newer architectures. On ***ISUnet***, these new architectures include switched 10BaseT ethernet and switched fast ethernet. These environments will be attached to the new core using aggregated fast ethernet and gigabit ethernet uplinks.

5.6 - Internet Integration

As mentioned in the beginning of this document, it is impossible to modify one component of ***ISUnet*** without impacting another. This is especially true for Internet connectivity. If the existing Internet connections based upon T1 circuits remain in place, the new core would quickly saturate these circuits. As a result, a new means for providing Internet connectivity needs to be introduced.

Internet connectivity projections for the fall of 1998 suggest that as much as 12Mb/s will be needed to provide basic connectivity for the ever-growing number of users consuming Internet resources from Illinois State University. Many unknowns such as the full impact of the residence hall networks and the new general education curriculum make this projection extremely uncertain.

Currently, the University pays approximately $3,000 a month for each Internet T1 circuit that is maintained with an ISP. With four Internet circuits in production, current costs are approximately $12,000 monthly or $144,000 annually. If these costs are projected for eight T1 circuits for the fall of 1998, this expense rises to $24,000 monthly or $288,000 annually for additional T1 circuits.

The obvious solution to this dilemma is to evolve Internet connectivity beyond T1 (1.54Mb/s) connections and implement T3 (45Mb/s) circuits. Unfortunately, leased line T3 connectivity is extremely expensive (ex: MCI estimate of $40,000 monthly from Willow Springs NAP). It is also important to remember the requirements for I2 (Internet 2) connectivity (ie: ATM). As a result, the most optimal solution would be the introduction of ATM over DS3 or better. In this model, standard Internet connectivity as well as I2 services could be provided using a common physical circuit.

Recently, GTE announced a public frame relay service based upon DS3 that would allow local Internet consumers (such as Illinois State University) to attach to a major POP locally. This offering includes fractional DS3 with a CIR (Committed Information Rates) as high as 10.5Mb/s and a burst rate as high as 21Mb/s. With this model, higher performance Internet connectivity could be obtained at a lower cost per Mb than conventional T1. From the local POP, Internet traffic would be backhauled to a Chicago POP with direct access to major NAPs in the region. As a result, the Office of Telecommunications has entered into a contract with GTE for the introduction of a 21Mb/s DS-3 frame relay circuit that could enter production as early as August of 1998. In the future, a second DS3 frame relay circuit could be introduced with NAP connectivity out of the St Louis region. This connectivity is also expected to migrate towards ATM over the course of the next year or so.

In addition to DS3 frame relay, an opportunity has emerged with the State of Illinois to obtain connectivity to a shared Internet T3 circuit that would be located at Illinois State as early as the fall of 1998. Currently, ISBE (Illinois State Board of Education) is developing the infrastructure for a K-12 Internet network throughout the state of Illinois. As a result, ISBE has approached the University about constructing an Internet POP (Point of Presence) that would provide basic Internet connectivity for as many as 28 local school districts using T1 circuits. In exchange for providing space and environmentals for the equipment and circuits needed to provide this connectivity, Illinois State would be able to consume a portion of the bandwidth on the ISBE circuit that connects to an ISP in Chicago. Once this new T3 circuit was operational, the ISBE connection would serve as a means for surviving major outages on the University's primary Internet circuits.

To research I2 connectivity, Illinois State University become an affiliate I2 member in the spring of 1998. This relationship allows the University to closely monitor the evolution of this new Internet infrastructure initiative. The University can then become a full I2 member at a future date for $25,000 annually plus the cost of connectivity (estimated at $400-500,000 annually) in preparation for bringing these services closer to reality.

5.7 – Private WAN Connectivity

As more and more users demand enhanced remote network connectivity, new means will need to be found and implemented beyond existing ISDN BRI (128Kb/s) solutions. In the fall of 1997, evaluation HDSL (High-speed Digital Subscriber Line) products were introduced using a standard four-wire circuit. These products offer speeds between E1 (2.0Mb/s) and 384Kb/s. By the fall of 1998, an ADSL (Asymmetric Digital Subscriber Line) service offering from GTE will become available which can also be used for WAN connectivity as it is very similar to frame relay.

5.8 – Remote Access Connectivity

The current remote access system for ***ISUnet*** is based upon X2 based (56K) analog dialup technologies. With the introduction of these technologies in February of 1998 came a restructuring of the campus dialup pool. A second pool containing 48 X2 based channels was introduced with a 30-minute connection limit. This pool will serve as the means for servicing users that have short duration connection requirements. The primary pool will continue to contain 282 X2 based channels with a 6-hour connection limit. By the fall of 1998, an additional 144 channels will be added to the campus dialup inventory.

The long-term futures of remote access services are still unclear. Currently, technologies such as ADSL (Asymmetrical Digital Subscriber Line) and bi-directional cable modems are being researched as potential architectures for a next generation remote access system for ***ISUnet***. An ADSL pilot has been tentatively scheduled for the fall of 1998.

5.9 – Network Management

There are many elements to network management including consoles, topology maps, statistical databases, and resource managers. Within ***ISUnet***, several major changes will occur within the arena of network management. These changes will be triggered by the core upgrade project.

The first major change is the introduction of a new management console. This console will integrate network topology maps and subsystems that manage specific network elements and devices using SNMP (Simple Network Management Protocol). A statistical database engine for collecting and reporting on RMON (Remote MONitoring) data will also be included.

Another major change is the introduction of a JAVA based front end to HP OpenView. This application allows the central network management console to be viewed using a common web browser. This enables support staff within Telecommunications to have immediate access to the management console from anywhere a browser is installed. Security within the application prevents information managed by this system to be viewable by non-authorized users.

At the same time, the existing resource database (currently known as PCLN) will be integrated with the telephone billing and resource databases to provide a single system that manages these resources collectively. This data will be viewable through a web browser and will be more current than the systems of today.

5.10 – IP Convergence

File and print sharing services based upon legacy protocols must migrate to IP (Internet Protocol) in order to provide universal access for all operating system and application platforms. For Windows 95 and NT workgroups and domains, this can be accomplished by implementing WINS (Windows Internet Name Server). WINS is a service that resolves NetBIOS hostnames to IP addresses that allows Windows clients to attach to Windows resources using native IP. In the next few years, this architecture will be augmented by an industry wide directory service based upon Active Directory. This service requires a central directory model using an architecture called LDAP (Lightweight Directory Access Protocol). This architecture will complete the application migration towards IP.

To provide mainframe emulation over IP, TN3270E can be implemented. TN3270E is an enhanced specification of traditional TN3270 that includes supports for printer session emulation. The lack of printer emulation support has been a major stumbling block to the elimination of SNA on campus. The new release of TCP/IP for MVS includes support for TN3270E. By the fall of 1998, Administrative Computing has indicated that this new release available.

6 - Translating the Model to Reality

6.1 - Research

Often, new application requirements appear with little or no warning. As such, the network architecture must be flexible so that it can evolve to support these emerging requirements. In this manner, services can be quickly offered once needs have been defined. This goal may only be accomplished when bi-directional communication exists between network consumers and providers. Without meaningful dialogue, responding to emerging needs may not occur in a timely manner.

Engineering a network infrastructure requires a structured planning process to develop both a functional design and an implementation scheme. Issues such as integration, scalability, and manageability are key to the success of evolving a network infrastructure. The result of this process is a clear and defined roadmap of technological evolution that matches current and anticipated needs of the University.

Once the model has been developed, the next step in the process is to match products with the developed architecture. In the case of ***ISUnet***, these products must support the requirements of the new core architecture defined in this document. As mentioned in the previous sections, these products are organized into three groups: 1) ATM core switch products, 2) edge (distribution) switch products, and 3) access switch products. Research for this project was based upon these three categories.

Another issue to be addressed through this research is whether these products should be “best of breed” or a vendor based suite of products. Considering the level of interaction with the vendor, the maturity of many of the technologies needed, integration issues, and training, the decision was made to implement products from a single vendor.

6.2 - Presentations to Vendors

The first step in the process of selecting products was to bring vendors to campus to review the requirements as defined in this document. These first meetings were purely informative. Each vendor was given a clear and concise overview of the requirements and objectives defined earlier in this document. The following vendors were invited to campus between February and August of 1997 for this phase of the project:

3Com
Madge
IBM
Accacia
Newbridge
Fore Systems
Xylan
Cisco Systems
Bay Networks
Cabletron

During these meetings, the required features and components of the electronics in each group of products were defined to each vendor.

6.3 - Presentations in Response

Once the vendors had been exposed to the requirements of this project, each of them were asked to return to campus and present their product line and address how these products could be implemented to meet the objectives defined to them in the previous meeting. These meetings focused predominantly on architecture with a limited introduction of products. The following vendors returned to campus to present information regarding their product suites:

3Com
IBM
Cabletron
Cisco Systems
Xylan
Fore Systems
Bay Networks

6.4 - Planning and Engineering

At this stage, information obtained from each vendor was compiled to develop models of the new ***ISUnet*** core using vendor specific products. These models defined the overall architecture along with supporting services and electronics.

After reviewing the materials presented during the previous sessions, it was possible to begin narrowing the list of vendors who had products that were viable for the new core network. Some product lines did not fit into the model. Some vendors had no legacy products of their own and some had no token ring product lines at all.

After narrowing the list, the remaining vendors were invited back to campus to further refine the model using specific products in great detail. The following vendors were asked to return:

3Com
Xylan
Cisco Systems
Bay Networks
Fore Systems
Cabletron

During these meetings, vendors were asked about products and services relating to such topics as integration, redundancy, scalability, and management. Neither Bay Networks nor Cabletron had products that could properly integrate token ring, ethernet, and ATM technologies. As such, these vendors were excluded from subsequent discussions.

6.5 - RFI (Request for Information)

After meeting with the remaining vendors, an RFI (Request for Information) was developed in order to provide information on product features from each vendor. The RFI was broken into two components. The first component dealt with edge switch requirements while the second component dealt with ATM core switch requirements. Since these were two separate groups of products, it made sense to separate this information.

After reviewing the information provided at previous meetings with these vendors, it was determined that only four of them had products that came close to meeting the requirements defined in this document. The RFI was sent to the following vendors:

Xylan
Fore Systems

Cisco Systems
3Com

After reviewing the responses to the RFI, both 3Com and Fore admitted that a third party set of products would be needed to complete the architecture if they were selected. As a result, 3Com and Fore were excluded from subsequent phases of this process.

6.6 - Product Evaluations

After much research, only Xylan and Cisco remained as candidates for product evaluations. In November of 1997, a request was submitted to both Xylan and Cisco for evaluation equipment. This gear was to be used to simulate the model that had been defined throughout this process. The goal of this phase was to gain the experience necessary to select the best group of products from a single vendor.

During the evaluation process, a plan was developed to test services such as LANE, policy based vLAN membership, and integrated routing with cut through techniques. More extensive testing would occur in the pilot phase once the products had been selected.

6.7 - Consultant Review

After the architecture and implementation of the new core had been clearly defined, several consultants were solicited for an objective review of this plan. Unfortunately, few fealt that they had the level of experience to effectively evaluate the plan for this project. Eventually, a vendor (Kent Datacom) who had no vested interest in this project was found to perform this review. This review took place during the months of December 1997 and January 1998 and was based partly upon this document and multiple interviews that were conducted with members of the Networking Systems staff.

6.8 - Proof of Concept Visits

During the evaluation process, both Xylan and Cisco were asked if the University could review other sites that were using their products in an environment that was similar to the proposed model. Xylan suggested a site visit to another university. This visit never occurred because the vendor was unable to locate an environment with similar requirements. Instead, Xylan was asked to provide references.

Cisco offered to have a meeting with product managers and engineers at their corporate office in RTP (Research Triangle Park), NC to assist in implementing the model that had been developed. Cisco was also asked to provide references.

6.9 – Selecting the Vendor

The most important criteria in selecting products for the ***ISUnet*** core upgrade project are the technologies that enable services to support the needs of the University. However, these requirements are only one aspect of this process. Other issues include the following:

- Can the vendor produce what it claims?
- Are these products developed within time frames that meet the needs of the University?
- Can the University escalate problems within the support organization?
- Is the vendor responsive to the needs of the University?
- Does the vendor provide adequate resources for resolving problems?
- Does the vendor devote adequate resources for researching new technologies and services?
- Does a user support group exist for these products?

- Can the vendor offer solutions for future requirements?
- Does the vendor offer educational pricing?
- Are volume discounts available?

6.10 – Why Cisco?

After completing the product evaluations and finalizing the selection process, the current and emerging products from Cisco Systems were found to be best suited for this project. These products are both flexible and feature rich. Cisco has a strong history of routing, switching, and ATM technologies, in the LAN, CAN, WAN, and remote access venues of networking giving them a high level of maturity. Cisco also has the largest educational marketshare of these products. These products met the requirements defined in this document better than the products from the other vendors reviewed within this process.

In addition to having a strong product line, Cisco has an outstanding post-sales support process. Unlike other vendors, Cisco has a local presence in Bloomington-Normal with a responsive account team. From this office, Cisco has brought resources from across the country to Illinois State University to meet and discuss topics and issues of interest.

Since 1993, the University has purchased products such as routers, terminal servers, and ethernet switches from Cisco totaling near $750,000. Based upon research, the ***ISUnet*** core upgrade project is estimated at approximately $4 million over the course of three years. Nearly $2 million of this estimate will be used for the purchase of electronics. As such, the University is making a large, long-term commitment towards a single vendor. Cisco has responded to this commitment with an aggressive discount rate.

6.11 – The Product Line

Cisco has a large variety of product suites to serve as ATM core switches, edge switches, and access switches. The Cisco Catalyst 8500 series switch will serve multiple roles. Initially, the 8500 with cell based hardware will be used as the ATM core switch within the new ***ISUnet*** core. This switch has 13 slots (one of which is used by the switch processor), a 20Gb/s cell based backplane, support for quad-port OC-12 (622Mb/s) modules, and single-port OC-48 (2.48Gb/s) modules. This switch also accepts PAMs (Port Access Module) from the previous generation of ATM switches called the Lightstream 1010. This includes modules such as quad-port OC-3, quad-port DS-3 (for WAN connectivity), and single-port OC-12.

The Cisco Catalyst 5500 series switches will be used as both edge and access switches within the new core network. This switch can be deployed as a 5, 9, or 13 slot chassis. A supervisor module controls clocking, network management, and related functions within the switch. A hardware-based daughtercard that resides on the supervisor module provides flow caching (called NetFlow) which enables cut-through routing services which enhances the scalability of the switch. The supervisor module also contains a pair of FEC (Fast Ether-Channel) uplinks. In addition, this switch supports services such as IP multicast pruning, PIM (Protocol Independent Multicast), VTP (Virtual Trunking Protocol) for announcing vLANs to other switches, VMPS (vLAN Management Policy Server) for implementing centrally administered vLAN policies, and aggressive network management using SNMPv2 and RMON.

As an edge switch, the Catalyst 5500 will be populated with a virtual routing engine called the RSM (Route Switch Module), one or more 12 port FEC (Fast Ether-Channel) modules for building switch integration, and a dual port MPOA OC-12 module for ATM core connectivity. The Julian Hall edge switch will also be populated with a pair of 16 port token ring modules to provide connectivity to the original core network.

By the end of 1998, the Catalyst 8500 using frame-based hardware will be used to displace the Catalyst 5500 series switch as an edge product on ***ISUnet***. At 15Gb/s, this new switch has greater backplane capacity than the 5500. The 8500 also includes an integrated, silicon based routing architecture that is highly scalable. Once this new product ships, the original edge switches will begin being redeployed as access switches. Modules from these edge switches will then be introduced into the new frame switches.

As an access switch, the Catalyst 5500 product can be populated with a variety of media cards including switched ethernet, 10/100 ethernet, and token ring. In densities of 24 ports or less, the Catalyst 1900 will be used. This switch supports all of the same features as the Catalyst 5500 but with a fixed number of ports.

To provide connectivity to legacy networks, the FEC (Fast Ether Channel) module can provide 100Mb/s connectivity from the building network to the edge using an architecture called ISL (InterSwitch Link). This technology allows the extension of vLANs through a logical frame multiplexing scheme. This prevents the need from deploying token ring modules within all edge switches to provide connectivity to legacy token ring networks. To make this process work, many existing token ring switches will be redeployed into lower density environments and displaced with the Cisco Catalyst 3900 series switch.

7 - Implementation Schedule

7.1 - Schedule Outline

The core upgrade project timetable is largely driven by the residence hall network project. This new environment will be provided connectivity to ***ISUnet*** through the new core. The table below outlines the general rollout schedule:

Table 7.1.1 – Core Upgrade Project Implementation Schedule

Phase	***Start Date***	***End Date***	***Description***
Pilot	2/1998	6/1998	finalize and test model
I	6/1998	12/1998	deploy ATM core, co-located edge switches
II	7/1998	1/1999	deploy residence hall switches in south, west
III	11/1998	12/2000	deploy remaining edge switches
IV	1/1999	12/2000	enhance matrix

7.2 - Pilot

The pilot project is designed to take the basic model reviewed during the evaluation phase and refine it through more extensive testing. Once developed, actual production applications will begin consuming resources on this network on a limited scale through a select group of users.

During this phase, the Cisco Lightstream 1010 will be used as the ATM core switch. This switch has 5 slots (one of which is used for the switch processor), a 5Gb/s cell based backplane, and support for single port density OC-12. At the same time, dual-port OC-12 MPOA (Multi-Protocol Over ATM) will be introduced into pilot edge switches. At that point, testing will include features as circuit aggregation and cut through routing using NHRP (Next Hop Routing Protocol).

7.3 - Phase I

At the start of the first phase, the Cisco Lightstream 1010 switch and related modules will be purchased and used to construct the new ATM core. As part of the product evaluation agreement with Cisco, these switches and related modules will be replaced with the 8500 and associated (higher density) modules once the products are available.

In this phase, each ATM core switch will be populated with multiple OC-12 modules for core and edge switch integration and dual ASP (ATM Switch Processor) modules for redundancy. These switches will be introduced in four locations: 1) the Julian Hall data center, 2) the Cook Hall switching center, 3) the Stevenson Hall switching center, and 3) the Redbird Arena switching center. All four locations have available space, appropriate environmentals, and fiber densities in place. Although SMF (Single Mode Fiber) has already been installed between all but 3 locations as described in this project, all strands are currently dark and unterminated. Many of these strands will need to be terminated by before Phase I can begin.

Once deployed, a redundant and distributed architecture for LES (LAN Emulation Server), LECS (LAN Emulation Configuration Server), and BUS (Broadcast Unknown Server) will be introduced using SSRP. At the same time these core switches are deployed, co-located edge switches will also be introduced. Each edge switch will attach to a primary ATM core switch and to a secondary ATM core switch using OC-12. Each ATM core switch will also attach to the other three switches at OC-12 rates.

During the summer of 1998, the core router servicing the dialup and all dialup related equipment will move to the expanded Stevenson Hall switching center. Once placed, a high performance link will be installed in this core router to provide ***ISUnet*** with connectivity to the new ISBE (Illinois State Board of Education) K-12 POP (Point of Presence). Intemet connectivity is provided to this POP by a Sprint DS-3 (45Mb/s) circuit that terminates in a POP in Chicago. In exchange for space providing space for the ISBE POP, the University can consume a percentage of this circuit for Internet connectivity. This project enhances the survivability of ***ISUnet*** and provides high performance connectivity into a statewide Internet backbone infrastructure. In the long-term, it is likely that the ISBE network may evolve into the Illinois Century Network and would provide even more opportunities for the University.

To provide support for MPOA (Multi-Protocol Over ATM), an ATM port adapter will be installed within the Julian Hall and Stevenson Hall core routers. These routers will be connected respectively to the Julian Hall and Stevenson Hall core switches through OC-3. This configuration will allow these routers to act as route servers using MPS (Multi-Protocol Server) for the MPC (Multi-Protocol Client) residing within the edge switches. This architecture provides for cut-through routing within the ATM core. A mechanism will be developed to distribute this process between the two routers so that MPS will not fail in the event a core router fails.

Within Julian Hall, the edge switch will also be used to provide connectivity to the original core network. This activity will be done in parallel with the existing core products. All token ring trunks will be attached to a pair of 16 port 16Mb/s token ring modules within this switch. This configuration will provide basic connectivity between the resources that reside on the old and new core networks.

7.4 - Phase II

In the second phase, the electronics for the residence hall networks will be deployed. This includes an estimated 17 high-density access switches within select residence hall complexes and edge switches that will provide these environments connectivity to the new core network. Intelligent, rack mount UPS (Uninterruptible Power Supply) will be placed in each location to provide line conditioning and limited battery backup for this equipment.

The south campus residence hall network is based upon approximately 1,500 ports of switched ethernet. This connectivity will be provided by 8 distributed access switches that are located evenly throughout the 4 towers of this complex. Each access switch will connect to the south campus edge switch that services this environment using 400Mb/s FDX ethernet (a potential aggregation of 3.2Gb/s) and an architecture called FEC (Fast Ether-Channel). This edge switch will then be attached to the new ATM core switch located in Cook Hall using OC-12 (622Mb/s).

The west campus residence hall network is based upon approximately 1,200 ports of switched ethernet. This connectivity will be provided by 9 distributed access switches that are located evenly throughout the 3 towers of this complex. Each access switch connects to the west campus edge switch that services this environment using FEC (400Mb/s FDX) ethernet (a potential aggregation of 3.6Gb/s). The edge switch will then be attached to the new core ATM switch located in Redbird Arena using OC-12 (622Mb/s).

At the same time rollouts are occurring for the south and west campus, the Smith House network will be optimized. Although the equipment located in Smith has already been upgraded to the same type used in the south and west campus, the connectivity for this gear has not been changed since the network went online in 1995. This network environment will be rehomed to the new edge switch deployed in Stevenson Hall using FEC (400Mb/s FDX). This change will allow users in Smith House to enjoy the same performance levels as those users in south and west campus residence hall networks.

Once these new networks are online, only Internet based resources will be available to users from within the residence hall networks. Since there are projects in the works that address access to NT networks and the mainframe using IP, this architecture should not impede student’s usage of this environment.

During the summer of 1998, the new DNS/DDNS/DHCP servers based upon Sun Ultra-1 servers will be brought into production. The existing DNS servers will be reconfigured to run other services for the University. These new servers will provide DNS and DHCP services for the entire campus community.

7.5 - Phase III

In the third phase, additional edge switches will be deployed within the new core to complete the preliminary matrix. The following table lists the remaining edge switch locations and the time frame for deployment of each. All edge switches will be connected to a respective core switch using OC-12 (622Mb/s) running MPOA.

Table 7.5.1 – Edge Switch Deployments in Phase III

Edge Switch	***Core Switch***	***Time Frame***
Old Union	Julian Hall	12/1998
Center for Visual Arts	Stevenson Hall	12/1998
Green Food Service	Redbird Arena	12/1998
Science Laboratory Building	Stevenson Hall	1/1999

Milner Library	Cook Hall	2/1999
Student Services Building	Cook Hall	3/1999
Fairchild Hall	Cook Hall	4/1999

Enough fiber is already in place to support OC-12 connections from each edge switch to a designated ATM core switch with the exception of the Milner Library, Fairchild Hall, and Student Services Building locations. New SMF (Single Mode Fiber) will need to be installed befor these environments can come online.

During this process, building network integration within the new core infrastructure will begin to occur. All token ring building switches will be connected to a respective edge switch using FTC (Fast Token Channel) with transfer rates up to 400Mb/s FDX using 802.5 framing. All ethernet building switches will be connected to a respective edge switch using FEC (Fast Ether-Channel) with transfer rates up to 400Mb/s FDX using 802.3 framing.

In order for building network integration to occur, the existing MMF (Multi-Mode Fiber) plant will need to be reconfigured to support uplinks from each building network to the designated edge switch locations. In most cases, this will require connections to pass through existing patch panels. These pathways will need to be determined and prepared prior to this phase of the project.

All building networks will be implemented during this phase by the schedule below. It is likely that the actual implementation process will vary from this timeline.

Table 7.5.2 – Building Network Integration Timetable

Network	***Type***	***Time Frame***
Julian Hall	token ring	12/1998
Julian Hall	ethernet	12/1998
Stevenson Hall	token ring	12/1998
Stevenson Hall	ethernet	12/1998
Old Union	token ring	1/1999
Old Union	ethernet	1/1999
Williams Hall	token ring	1/1999
Williams Hall	ethernet	1/1999
CVA	ethernet	2/1999
CECW	ethernet	2/1999
Cook Hall	ethernet	2/1999
Science Building	token ring	2/1999
Science Building	ethernet	2/1999
ORL	token ring	3/1999
Milner Library	token ring	3/1999
Milner Library	ethernet	3/1999
Student Services	token ring	3/1999
Student Services	ethernet	3/1999
Fell Hall	token ring	4/1999
DeGarmo	token ring	4/1999
DeGarmo	ethernet	4/1999
Felmley	ethernet	4/1999
Felmley Annex	token ring	4/1999
Moulton Hall	token ring	5/1999
Moulton Hall	ethernet	5/1999
Schroeder Hall	token ring	5/1999
Hovey Hall	token ring	5/1999
Hovey Hall	ethernet	6/1999

University High	token ring	6/1999
Turner Hall	token ring	6/1999
Turner Hall	ethernet	6/1999
Physical Plant	token ring	7/1999
Bone Student Center	token ring	7/1999
Fairchild Hall	token ring	7/1999
Metcalf	token ring	7/1999
Student Accounts	token ring	7/1999
Instructional Tech	ethernet	8/1999
Vidette	ethernet	8/1999
McCormick	token ring	8/1999
Redbird Arena	token ring	8/1999
Nelson Smith	token ring	9/1999
Nelson Smith	ethernet	9/1999
Ropp	token ring	9/1999
Huddleson	ethernet	9/1999
Green Food Service	token ring	10/1999
Edwards	token ring	10/1999
Conferencing Unit	token ring	10/1999
Conferencing Unit	ethernet	10/1999

There are many other environments that will be addressed by this project that are not listed in this schedule. These networks will be targeted once the environments listed above have been completed.

7.6 - Phase IV

The basic foundation of the new core will have already been deployed by the time the fourth phase occurs. This phase will focus on evolving the existing matrix within the core, the introduction of server farms in select environments, and the introduction of fast ethernet connectivity for select desktop environments.

Within this phase, the matrix will be evolved using OC-48 (2.48Gb/s) between ATM core switches. This project will quadruple the overall bandwidth between ATM core switches and will prepare the new core for the introduction of native telephony and video services.

In previous phases of this project, 100BaseT will have been made available for many servers that reside on ***ISUnet***. These connections will have allowed for optimized connectivity of key services supporting large numbers of clients. In this phase, 100BaseT connectivity will also be made available to select desktop systems in prepared environments. All eligible desktop systems must implement a BUS architecture that is capable of speeds up to 100Mb/s (ie: PCI) and connectivity must be based upon UTP-5 cabling. 100BaseT cannot be implemented within environments based upon STP cabling.

With the emergence of the new ***ISUnet*** core will come the idea of a managed migration process from token ring technologies to switched 10BaseX and 100BaseX ethernet technologies. This process will occur over many years. Fiscal incentives will be put in place to encourage this migration at a steady yet manageable pace.

As the core network is enhanced and higher performance connectivity is made available at the desktop, server farms will begin to emerge. A high performance server farm will be constructed for central services in the Julian Hall data center. This will include the major Internet servers (electronic mail, network news, the web, authentication, etc), accounting system, debit card systems, and the mainframe. To accomplish this objective, a pair of high-density switches will be deployed to provide switched, fast ethernet connectivity to any central service within the Julian Hall data center. Each switch will be attached to the Julian Hall and Stevenson Hall ATM core switches through OC-12 to offer survivability in the event of an outage and scalability in the event of over consumption.

7.7 - Beyond

Once the last phase of the core upgrade project is complete, the evolution of ***ISUnet*** will become more managable. As density or bandwidth needs increase in select environments, additional edge switches will be deployed. These new deployments will also trigger the introduction of aggregated OC-12 links between edge and ATM core switches.

Additional residence hall networks in Central Campus (Dunn-Barton and Walker), East Campus (Hewett and Manchester), and Watterson Towers are also expected. These environments will be constructed using similar technologies to the environments deployed in the south and west campus residence hall networks. Requests for connectivity for the University owned apartment complexes of Shelbourne and Cardinal Court will also emerge. At the same time, proposals will be developed to extend ***ISUnet*** into select Greek houses.

By this time, gigabit ethernet (IEEE 802.3z) should be a full standard. It will first begin to appear as an uplink technology between building networks and edge switches. Gigabit ethernet will also become available for very high-performance servers within clusters. As gigabit technologies emerge and are deployed, the backplane capacity of the original edge switches of the new core will become inadequate. At this point, the original edge switches will then be displaced with higher capacity frame based products. These displaced switches will then be redeployed as access switches as the network continues to expand. This process could continue indefinitely.

Over time, newer applications, services, and technologies will be introduced into ***ISUnet*** in response to user requirements. These events will trigger further activities to evolve the network. In reality, this process never ends.

8 - Conclusion

8.1 - ISUnet is a Strategic Resource

The health and viability of ***ISUnet*** greatly impacts the activities that occur at Illinois State University. To meet the needs of these activities, this network must evolve. Since requirements continuously change, it is only natural that the components of this network will need to evolve to support these needs. As a result, evolution of the network is a continuous process.

8.2 – The Time is Now

This document defines a very clear process for addressing the networking needs of Illinois State University. A substantial amount of research has taken place in the effort to develop this plan. However, much more work lies ahead during the implementation phase of this project. Since implementation will take two years or more, now is the time to begin this work. As these projects progress, ***ISUnet*** will begin to emerge as an outstanding infrastructure for network based services and will allow the University to maximize information technology as it enters the new millennium.