3490 Piedmont Rd.
Atlanta, GA 30305
Phone (404) 262-1555
Fax (404) 262-2055

Convergent Media Systems

Inter-Networking Strategic Plan

Submitted by: Philip Elwood

Inter-networking Strategic Plan

Meeting today's networking needs while preparing for tomorrow's demands.

Purpose

Provide the bandwidth necessary for the projected increase in network usage created by the implementation of Microsoft Exchange, a Client-server based accounting system, video, voice and large data file transfers. Create an environment capable of supporting a corporate based Intranet providing secure internal and external access to confidential data. Provide increased redundancy for the voice and data traffic, preventing down time and limiting loss of work. Install the tools necessary for the monitoring of the system and for the conducting of preventative maintenance reducing the need for performing corrective measures. Create a hardware infrastructure that will provide maximum expandability and flexibility to meet current needs while preparing for future demands.

Executive Summary

This proposal is meant to provide a solution to the problems currently existing on the network. This will provide Convergent users with a network infrastructure capable of meeting existing needs as well as future requirements. The proposal is broken into three levels: Wide Area Network, Primary Site Network (100 or more devices) and Secondary Site Network. The three areas have different requirements with only minimal interaction between each area. Estimation of the costs in this proposal is based upon personal experience and the consideration of keeping budgeted costs to realistic figures. Actual implementation may cause possible fluctuations of these figures within ten percent of listed costs.

Wide Area Network

The Wide Area Network encompasses the interconnection between the different Convergent sites. The Wide Area Network ends at each site with the connection to the local network.

The proposal consists of a change to a frame relay network from the existing direct connection. This will provide the best growth plan for each site connection. It also allows the best current means of increasing the number of Convergent sites. Each site will be provided with Internet connectivity through this network.

The Wide Area Network will run predominantly TCP/IP protocols between sites. Each network device needs to have the ability to be remotely monitored and managed. All network equipment will have redundancy built in and a maintenance agreement. Security will be provided through a firewall at each site providing access to the Internet.

The implementation will be done in six phases for a total cost of $145,000.

Primary Site Network

Primary network sites are the Convergent locations with the greatest number of machines and demands to the network. Typically, sites with over a hundred network devices. Currently two sites meet these criteria: Atlanta and Littleton.

The proposal consists of a change to a switched Ethernet environment with each server and closet having a dedicated connection to a 100-megabit Ethernet connection. Each closet will be connected to the central switch with a fiber connection.

The backbone of the network will be a chassis based unit with modules for multiple 100mbit Ethernet connections. Closets will each contain at least one 10mbit switch that will connect to the backbone switch with a fiber connection. Clients will connect to each closet through 10BaseT hubs that are connected to ports on the closet switches.

The Network will run the TCP/IP protocol predominantly. Each network device needs to have the ability to be remotely monitored and managed. Security will be provided through a firewall at each site providing access to the Internet. Remote access security will be provided at each site providing remote access to the network. Management will be through an Enterprise management server based in Atlanta. All network equipment will have redundancy built in and a maintenance agreement.

The implementation of the Atlanta site will be done in four phases for a total cost of $196,632.

Secondary Site Network

A secondary site network has fewer requirements than the primary sites. The equipment is typically what is based in a primary site closet.

The proposal consists of a change to a switched Ethernet environment with each server and closet having a dedicated connection to a 10BaseT Ethernet connection. Clients will connect to each closet through 10BaseT hubs that are connected to ports on the central switch.

The Network will run the TCP/IP protocol predominantly. Each network device needs to have the ability to be remotely monitored and managed. Remote access will be provided through the primary sites. Security will also be provided through the primary sites for both Internet and remote access security. Management will be through an Enterprise management server based in Atlanta. All network equipment will have redundancy built in and a maintenance agreement.

The implementation will be done in three phases for a cost of $50,200.

Current Configuration

Convergent's current Local Area Network utilizes technology that is over ten years old. This technology, while suitable for the network needs at that time, has difficulty with the network needs of today.

All the closets and the computer room are interconnected with a thickwire Ethernet backbone cable capable of handling Ethernet traffic up to 10 megabytes per second. Each closet has a single tap and the primary servers have individual taps into this cable.

Each closet utilizes repeaters to broadcast packets to the client machines. The Ethernet tap extends to a Digital Ethernet Local Network Interconnect (DELNI) at the top end of the closet network. This DELNI then extends the network to multiple Digital Ethernet Multiport Repeaters (DEMPR). Each DEMPR then distributes the packets to client machines through Thinwire Ethernet connections. Each Thinwire connection has between one and six devices connecting to the chain of cable.

Problems

Increased network usage – We have moved from a terminal (text) based environment to a PC (graphics) based environment. This is causing greater congestion of the existing network due to the increased bandwidth requirements for graphic transmissions. The network utilization will continue to increase as further applications and Internet/Intranet usage increases. Network congestion causes slowdowns due to packet retransmissions as multiple devices try to send over the limited amount of bandwidth. This is like trying to cram a banana into a coke bottle, not all of it makes it inside.

Collisions – Collisions are a major concern on our existing Thinwire network. They occur for two reasons, packets are transmitted at the same time or there is a problem with a hardware device. The more common of the two possibilities is the transmittal of packets at the same time. This problem is occurring regularly and is showing a definite impact to the speed of our network. The less common but more dangerous of the two possibilities is the hardware device problem. This problem has shown itself multiple times in our existing wiring, causing network storms during each occurrence. A network storm occurs when packet collisions happen at a rate where all activity on the network ceases. This has been predominantly caused by a faulty cable connecting one of the client machines to the network.

Management – The only way a problem is located on the existing network is to isolate it down to the specific cable or device by going between closets and experimenting with the connections. This is a very slow and tedious process. A recent problem took an arduous two weeks to isolate and an additional six agonizing hours to correct. This problem intermittently affected five users over a two-week period.

Wide Area Network

Goal

Provide a Wide Area Networking infrastructure capable of handling voice, video and heavy data traffic. The environment needs to be able to handle the addition of multiple sites and easy growth of bandwidth requirements for existing sites. Redundancy needs to be built in to limit the loss of communication between sites. Tools for monitoring and analyzing data traffic between sites that provide reports and alarms are required.

Configuration

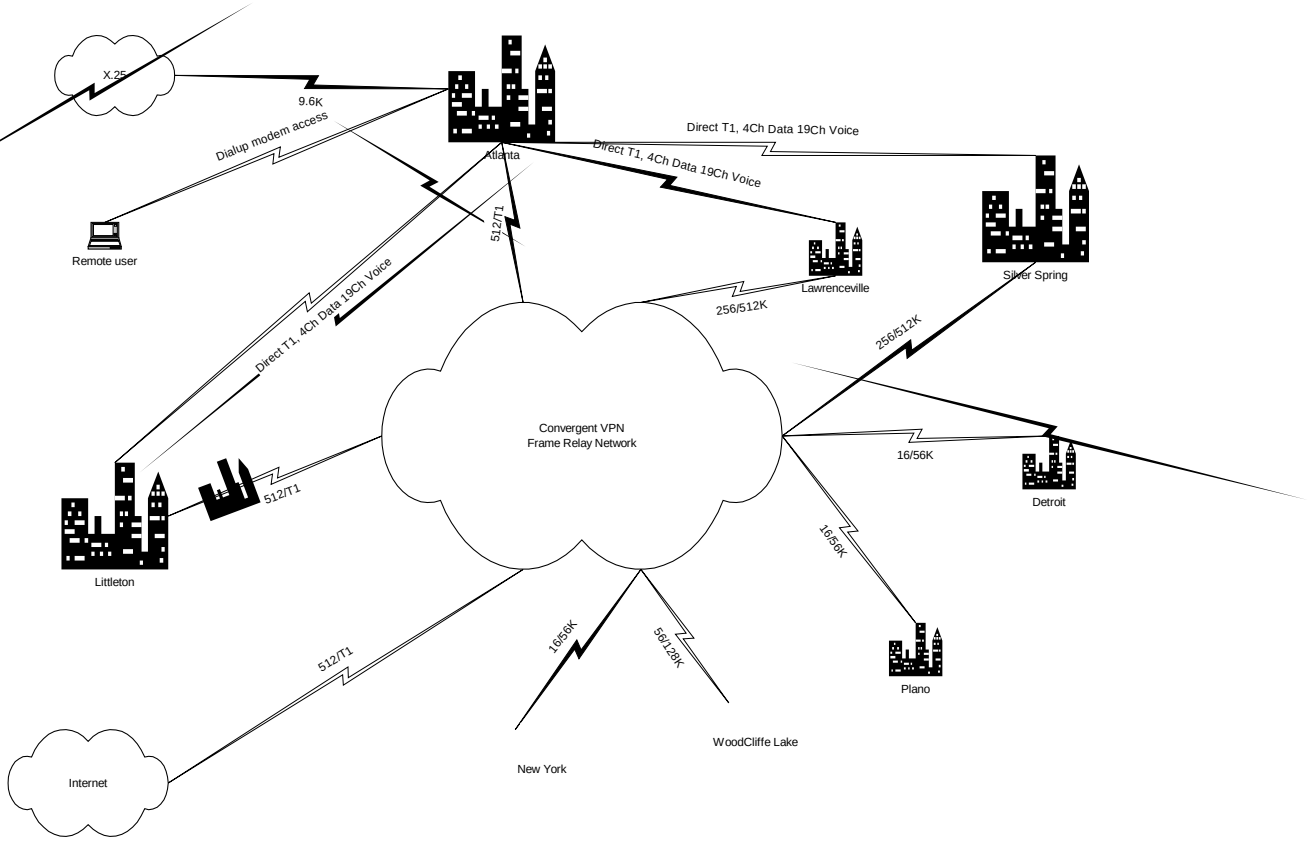

Connections

The main portion of this proposed Wide Area Network proposal is the conversion of the data access between sites from dedicated links to a frame relay cloud. This will create a virtual private network (VPN) allowing for single logical connections to exist between sites. All sites will be given committed information rates (CIR) and access rates needed for normal business operations. The primary and secondary sites will require the highest amount of bandwidth.

A primary control site for the frame relay network would be established in Atlanta. A secondary control site would be established in Littleton for redundancy purposes. In the event of a failure to the primary control site, the network would be able to function through the

secondary control site. The frame relay provider will require a 30-day notice to replace either control site.

The existing dedicated connections between Atlanta, Lawrenceville, Littleton and Silver Spring need to remain in place to handle the existing voice traffic. Hardware is currently being developed to handle voice and data across the same frame relay network. This option is still in the early stages of development and the Quality of Service (QoS) currently available for our business needs is poor. Once the QoS is standardized and reaches an acceptable QoS level, these direct lines may be eliminated.

The Internet connection will be placed in the frame relay network with access granted through the two control sites. This provides redundancy as well as allowing for greater Internet throughput.

Protocols

The network will be structured as a routed network. TCP/IP will be the only protocol routed between sites. This will better utilize the existing resources and eliminate unnecessary transfer of packets between sites. The Windows NT networks will use WINS, DHCP and DNS to communicate over the Wide Area Network.

Data

The majority of the current use of our Wide Area Network links is small text transfers with only occasional large file transfers. The use of a frame relay topology allows us to easily utilize our normal information transfer rate while allowing for the occasional large transfer. This is done through the purchase of a CIR at a lower level than the access rate.

Management

Management of the Wide Area Network will be done with SNMP and RMON Management. The SNMP management platform will allow us to monitor real time traffic statistics. RMON management will provide us with the ability to make quick device changes. All major network devices need to be SNMP compliant. RMON ability is preferred.

The network devices will be managed through a central interface that will monitor each device notifying the manager of a device failure or problem. The management console also needs to have the ability to track network usage between devices and provide reports showing network trends.

Maintenance

All primary Wide Area Network equipment is to be on a 24/7 four-hour response maintenance agreement. Spares of primary items need to be available for quick replacement for major sites.

Redundancy

Having a secondary control site at a different location allows us the ability to still utilize the Wide Area Network if the primary control site fails. The frame relay network will reroute through this secondary site if the primary cannot be found. This is the same case with the Internet connection.

Having a secondary path into the building of the primary control site by a different carrier will limit the possibility of a primary site failure.

Maintaining a data link through the existing T1 connections to the primary sites will also provide a redundant link in the event of a line failure.

All major pieces of network equipment should have redundancy built in. A secondary power supply and extra ports will limit the effect of a hardware failure. Hot swappable units will limit or eliminate the amount of downtime required to replace a failed part.

Security

A firewall at each of the sites allowing access to the Internet will provide a secure connection to the Internet and limit access into the network. Setting up proxy servers at these sites will also limit the machine addresses sent to the Internet. All dial-up access into the network needs to be through a secure connection that is Radius compliant.

Equipment Required

Description	Quantity	Unit Cost	Extended Cost	Existing units	Total Cost
Large-office Router	3	$25000	$75000	0	$75000
Multi-port Router	1	15000	15000	1	0
Small-office Router	6	5000	30000	3	15000
56K CSU/DSU	3	500	1500	4	0
T1 CSU/DSU	13	1000	13000	6	7000
Firewalls	2	15000	30000	0	30000
WAN Analysis	2	5000	10000	0	10000
Training	6	2000	12000	0	12000
Travel	6	500	3000	0	3000
Total			**$189,500**		**$152,000**

This model replaces just a few items. The central router in the Atlanta office will be replaced with a Large-office router. The bridges that connect Atlanta to Lawrenceville and New York will be replaced with Small-office routers at each of those sites.

The other two Large-office routers are to provide secure Internet access, in conjunction with the Firewalls, to the Convergent network.

Options Considered

Routers

Bay Networks

Type	Model	Cost
Large-office router	Backbone Link Node	$23450
Multi-port router	Access Stack Node	13500
Small office router	Advanced Remote Node	4295

Strengths

- Regarded as one of the two best router lines
- Excellent redundancy/hot-standby options and fault-tolerant design
- Best throughput performance
- Most complete and extensive WAN interface support
- VLAN (Virtual LAN) capable

Weaknesses

- Minimal and inconsistent statistics
- Minimal security features

Cisco Systems

Type	Model	Cost
Large-office router	7200	$23800
Multi-port router	4500	14500
Small office router	2524	4200

Strengths

- Regarded as one of the two best router lines
- High-density Ethernet cards
- Rich features, including Hot Standby Routing Protocol, for configuring alternate, redundant and backup links
- Extensive features for traffic prioritization, WAN link compression, Virtual LANs and security
- Web server within router supports a browser-based version of the command-line interface
- Modular design allows for insertion of additional modules when required. Modules on the 2524 include a CSU/DSU module

Weaknesses

- No ATM or ISDN support yet for router
- Some module compatibility issues between models
- Most arduous, nonintuitive command-line console interface
- Inconsistent set of management applications provides disparate management capabilities and very different views of router performance and statistics

DSU/CSU

Adtran

Type	Model	Cost
56K CSU/DSU	DSU IV	$500
T1 CSU/DSU	TSU100	1000

Strengths

- Considered a leader in the industry
- SMNP management interface includes ability to report on line failures

Weaknesses

- Stand-alone units
- Not easily rack-mountable

Other

No other standalone DSUs were considered due to the quantity of Adtran units currently owned by Convergent.

DSU/CSU units internal to routers are the only other option considered. The strengths and weaknesses reflect that of internal units.

Strengths

- Direct bus interface to router.
- All-inclusive unit with router prevents downtime due to single unit power sources.
- All-inclusive units easy to place in remote locations.

Weaknesses

- Only available on lower end routers.
- A swap of DSU unit requires router restart.

Firewalls

Cisco Systems

Type	Model	Cost
PIX Firewall	PIX	$3600
256 user upgrade	SWPIX-256-3.0	11000
10/100 Ethernet interface	PIX-1FE	400
Total		**$15,000**

Strengths

- Outside network access based on industry-standard authentication protocols such as TACACS+ and RADIUS
- Conceals internal network architecture
- Quick easy configuration with web interface
- Audit information through Syslog MIB support
- No downtime required for installation
- No day-to-day management required
- Standalone machine, no extra machine required
- Full outbound Internet access from unregistered internal hosts
- Interoperable with Cisco routers
- Internal client addresses hidden from external network
- Can support more than 16000 simultaneous connections
- Proprietary operating system difficult for hackers

Weaknesses

- Proprietary operating system
- No third party tests and comparisons have been done
- Limited functionality reviews by some users

CyberGuard Firewall

Type	Model	Cost
Firewall	3.0	$10000
Machine	Pentium 200, 80meg memory	4000
Total		**$14,000**

Strengths

- Industry leader
- Excellent reviews
- Proxy filter and limits to user applications (ex. Realaudio)
- Remote user Secure ID support
- Xwindows user interface for easy management
- Remotely manageable

Weaknesses

- Unix operating system
- Does not come with a machine

Checkpoint Firewall

Type	Model	Cost
Firewall	12.1 for NT	$10000
Machine	Pentium 200, 80meg memory	4000
Training		2000
Total		**$16,000**

Strengths

- Industry leader
- Excellent reviews
- NT or Unix based package
- Remote users managed with Radius and/or SecureID
- Rules based management
- Intuitive and easily managed
- Virus scan
- Administrated remotely with Windows 95, NT and/or Unix

Weaknesses

- Does not come with a machine

Recommendation

Router - Cisco

Both families of routers will more than adequately support our environment. Both Bay Networks and Cisco are considered leaders in the industry. The Bay Networks routers can handle more data throughput and are easier to manage while the Cisco routers are very feature rich.

Due to our current investment in the Cisco router family, Cisco is the logical choice for our router family. There are no overwhelming benefits to choosing another family.

CSU/DSU

Where the option is available, the internal CSU/DSU is preferred. This is typically only available in the low-end router families. All external routers will be Adtran units capable of SMNP monitoring.

Firewalls

All the Firewalls appear to meet our needs of securing the network. CheckPoint for NT has been chosen as our Firewall because of its excellent status as an industry leader, NT based platform, Radius support and easy manageability. The Firewall will be placed at the point of contact to the Internet, between an external router and an internal router.

Stages for deployment

Cost

Stage	Cost excluding WAN	Cost
Frame Analysis	$5000	$5000
Atlanta, New York and WoodCliffe Lake	33800	33800
Plano and Detroit	10400	29200
Littleton	1000	1000
Silver Spring and Lawrenceville	6200	6200
Internet	88600	88600
Total	**$145,000**	**$163,800**

Analysis of current Frame Relay requirements

A third party who has the means of testing the current use of our existing direct connections will be brought in. An analysis of the links between Atlanta and all other sites will take place. The analysis will determine current normal and peak bandwidth utilization. It will also determine the types and percentage of each protocol being broadcast.

Type	Number	Unit Cost	Total Cost
Frame Relay Analysis	1	$5000	$5000
Total			**$5,000**

Atlanta (control site), New York, WoodCliffe Lake,

Goal

Move the two existing 56K lines to a Frame Relay network. This will provide us with a robust environment in which to grow our Wide Area Network. The Atlanta location will be the hub of the Frame Relay network with a router that allows for growth of this network. New York will be provided with a 16/56K connection due to its low utilization. WoodCliffe Lake may require higher bandwidth utilization. This will be determined by the analysis.

Implementation steps

1. Training

   Prior to the purchase of the routing equipment, training on the equipment needs to occur. Training prior to configuration will allow for a more effective configuration of the routers. This training should encompass both the regular and advanced router configuration courses for the network administrator and the regular course for a backup administrator. The network administrator may also require additional coursework to properly handle Internetworking.

Type	Number	Unit Cost	Total Cost
Router training	2	$2000	$4000
Advanced training	1	2000	2000
Network training	1	2000	2000
Total			**$8,000**

2. Purchase of equipment

   A large router will be purchased for the Atlanta site. The small router currently located in Atlanta connecting to the WoodCliffe Lake connection will be moved to the New York facility. The router currently located at WoodCliffe Lake will be left in place. The existing CSU/DSU units currently existing at each site will be left in place. The two 56K CSU units in Atlanta that currently connect those sites will be removed and replaced with a T1 CSU/DSU for Atlanta's Frame Relay connection.

Type	Number	Unit Cost	Total Cost
Large-office Router	1	$23800	$23800
Small-office Router	0	4200	0
T1 CSU/DSU	1	1000	1000
Total			**$24,800**

3. Installation and configuration

   The configuration and testing of the equipment, with the exclusion of the existing WoodCliffe Lake router, will be done at the Atlanta location prior to it being installed at the remote location. An administrator will need to be at both the Atlanta location and the remote location during the installation and testing of the equipment. The WoodCliffe Lake router will be configured remotely with an administrator onsite in case of an emergency.

Type	Number	Unit Cost	Total Cost
New York travel	1	$500	$500
WoodCliffe Lake travel	1	500	500
Total			**$1,000**

Plano and Detroit.

Goal

Move these sites to a Frame Relay network. Each of these sites currently uses slow individual and expensive dial-up lines to connect into the Convergent network. This connection will lower the cost and increase the speed of these connections. Both are small sites and will be provided with a 16/56K connection due to initial low utilization.

Implementation steps

1. Site survey of Plano and Detroit

   These are two new sites that currently have minimal or no established network. This will be done in conjunction with a LAN analysis of the locations. The survey will determine location of wiring closet, number of networked machines, needs of the users and the wiring needs of the location.

Type	Number	Unit Cost	Total Cost
Plano travel	1	$500	$500
Detroit travel	1	500	500
Total			**$1,000**

2. Purchase of equipment

   Two Small-office routers will be purchased to cover these two sites. The two available 56K CSU units will be used to connect the sites to the Frame Relay network.

Type	Number	Unit Cost	Total Cost
Small-office Router	2	$4200	$8400
56K CSU/DSU	0	1000	0
Network*	2	5000	10000
Cabling*	2	4900	7800
UPS*	2	500	1000
Total			**$27,200**

** Local Network equipment purchases and cabling will be done at this time.*

3. Installation and configuration

   The configuration and testing of the equipment will be done at the Atlanta location prior to it being installed at the remote location. An administrator will need to be at both the Atlanta location and the remote location during the installation and testing of the equipment.

Type	Number	Unit Cost	Total Cost
Plano travel	1	$500	$500
Detroit travel	1	500	500
Total			**$1,000**

Littleton (secondary control site)

Goal

Provide the Littleton site with the bandwidth needed for attachment between Littleton and the other Convergent sites. Provide a flexible, quick and easy growth path when more bandwidth is required. This site will also serve as a redundant Frame Relay host site (secondary control site). This site will not handle host functionality unless the primary host site in Atlanta has failed.

Implementation steps

1. Purchase of equipment

   A T1 CSU will be purchased for the connection of this site. It will connect into a port on the existing router.

Type	Number	Unit Cost	Total Cost
Multi-port Router	0	$14500	$0
T1 CSU/DSU	1	1000	1000
Total			**$1,000**

2. Installation and configuration

   Configuration will be done remotely with the assistance of the local administrator. Installation and testing will be done by the local administrator in conjunction with the network administrator.

   The existing connection will be left in place with the channels being made available for voice as needed.

Silver Spring & Lawrenceville

Goal

Provide the Silver Spring and Lawrenceville sites with the bandwidth required for attachment between the either site and the other Convergent sites. Provide a flexible, quick and easy growth path when more bandwidth is needed.

Implementation steps

1. Purchase of equipment

   A router will be purchased for the Lawrenceville site. The existing router in Silver Spring will be left in place. Two T1 CSU units will be purchased for the connection of these sites. The Silver Spring CSU will connect into a port on the existing router.

Type	Number	Unit Cost	Total Cost
Small-office Router	1	$4200	$4200
T1 CSU/DSU	2	1000	2000
Total			**$6,200**

2. Installation and configuration

   Configuration of the Silver Spring equipment will be done remotely with the assistance of the local administrator. Installation and testing will be done by the local administrator in conjunction with the network administrator.

   Configuration of the Lawrenceville equipment will be done in Atlanta. The installation and testing will be done with the use of the Atlanta staff in conjunction with the network administrator.

   The existing T1 connections will be left in place with the channels being made available for voice as needed.

Internet connection

Goal

Provide a connection to the Internet closest to the hub of the Convergent network. To have a secondary Internet site at a different location with a different Internet provider for redundancy.

Implementation steps

1. Multi-site plan for the Internet

   An Internet expert will be given the task of providing Convergent with a plan that will meet the organization's Internet demands.

Type	Number	Unit Cost	Total Cost
Internet Analysis	1	$5000	$5000
Total			**$5,000**

2. Training on Firewall

   Prior to the installation of the Firewalls, it is important for the network administrator and a backup administrator to be trained on the software.

Type	Number	Unit Cost	Total Cost
Firewall training	2	$2000	$4000
Total			**$4,000**

3. Purchase of equipment

   Two routers will be purchased to handle the dual connections to the Internet. These routers will be placed outside the two firewalls, which will protect the inner network routers.

Type	Number	Unit Cost	Total Cost
Large-office routers	2	$23800	$47600
T1 CSU/DSU	2	1000	2000
Firewall	2	15000	30000
Total			**$79,600**

4. Installation and configuration

   The configuration of the equipment will be in the Atlanta office. It is very important for both Firewalls to be configured the same. The installation will be performed at each site by a local administrator in conjunction with the network administrator.

Primary Site Network (100 or more devices)

Goal

Provide a high-speed network with a centralized redundant backbone to meet current network needs while providing for a robust environment that will grow to accommodate future network requirements. Provide an environment that minimizes the amount of downtime through redundancy and monitoring.

Provide clients with the fastest reliable connection to network resources required. Create an environment allowing clients to run voice, video as well as data with minimal impact to the network.

Provide servers with individual network connections to limit the extra amount of data the server network cards have to process. Segment the client connections to a maximum of 50 users per segment.

Configuration

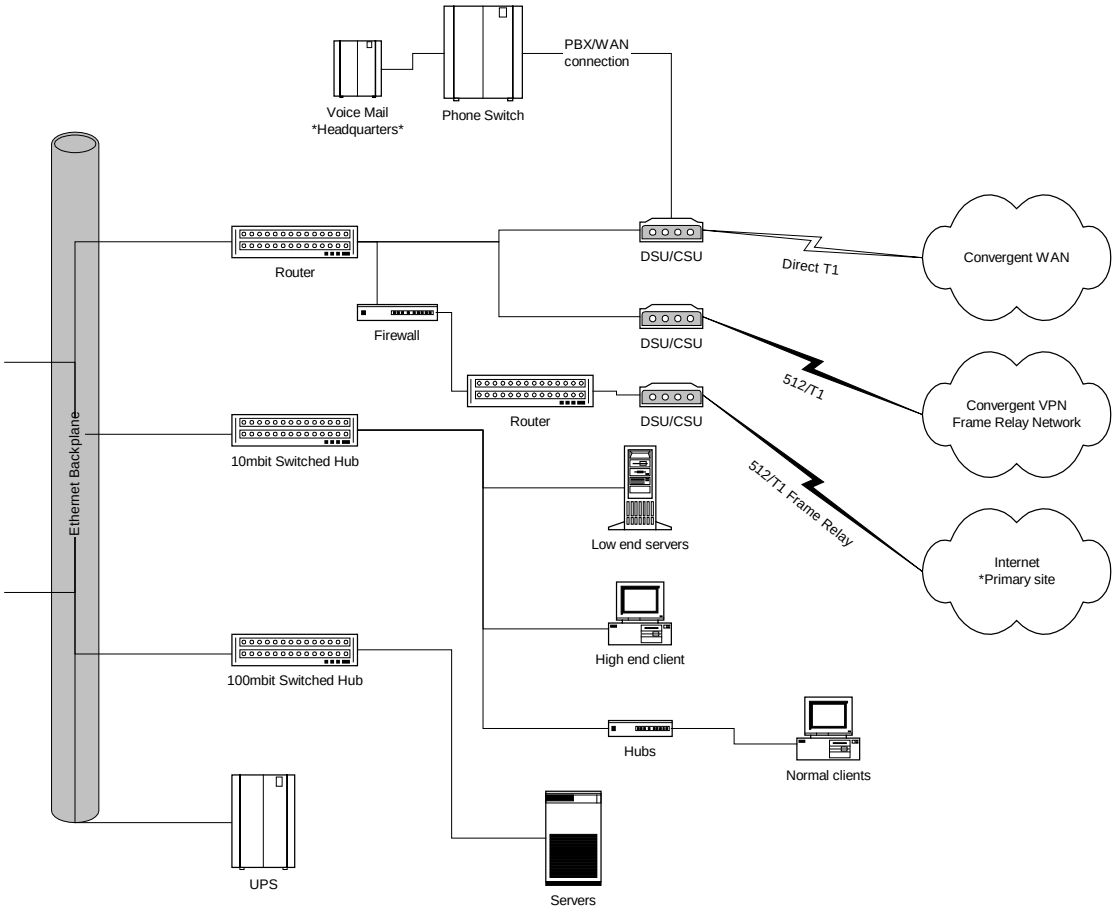

Backplane

The current need is for a chassis based system with dedicated 100BaseT switched ports for server connections and two 100Mbit (Fiber or Cat5 cabling) switched ports for each closet segment. In this configuration, each major server would be given its own network segment thereby limiting the network traffic that each server sees to only the packets dedicated to that server. By segmenting the closets, we will provide faster and more reliable links between the closets and the servers to which they need access. This will limit the packet flow on the network to the machines based off that segment.

Server connections

The proposed configuration is for a 10 port chassis based system that provides for expandability and flexibility. The chassis needs to be able to utilize 10BaseT, switched 10BaseT, 100BaseT, switched 100BaseT, Fiber Ethernet, Fiber 100, ATM and Fiber ATM. Dual power supplies and extra ports are needed for redundancy.

Closet/client connections

Each closet should have two 100Mbit Fiber or Cat5 switched Ethernet segments for redundancy and performance. Fiber is the preferred connection due to cleaner signal flow and no interference from electrical devices. A redundant connection should be made to each closet to minimize problems caused by cable breaks or connection device failures. A redundant connection that has the ability to be utilized as well as having it for redundancy purposes is preferred. A ten port switched hub will be at the base connection of each closet, accepting the 100Mbit connections and passing the packets to secondary hubs and high end clients. A normal client in this configuration will have a 10BaseT connection to a closet hub. Growth should allow each closet to segment itself into 100BaseT and 10BaseT segments depending upon client needs. Closet client connections should all be to patch panels from which patch cords connect them to the hubs. Each client connection should be to a level 5 jack in either a wall plate or a desk box. A accurate network map of each closet needs be maintained.

Protocols

Data transfers are to be through TCP/IP. 100BaseT and 10BaseT will connect client and server devices. 100Mbit fiber or Cat5 Ethernet will connect closet and backbone devices. Cat5 Ethernet connections will connect every server and client to a hub or switch.

Network devices

Switches

Switches need to have a minimum of 10 ports, all of which should all be SNMP configurable for monitoring and remote management. They should also be chassis based for redundancy and the ease of adding devices in the computer room. Closet switches can be stand-alone or stackable units. For ease of management, all switches need to be from the same manufacturer.

Hubs

Hubs should all be SNMP and RMON manageable and either chassis based or stackable. A minimum of 16 ports per unit is needed. For ease of management, all hubs need to be from the same manufacturer.

Remote Access

Remote connectivity should be as fast and reliable as is currently available. The remote access server should allow for ISDN access and allow for future technologies being made available by the manufacturers and carriers (ex: ADSL and 50K modems). Connecting should be easy for the client with the loss of data over the telephone link being minimized as much as possible. Built in modem connections that are easy to upgrade and hot swappable are preferred.

Security

Firewall

A firewall at each of the sites allowing access to the Internet will provide a secure connection to the Internet and limit access into the network. Setting up proxy servers at these sites will also limit the machine addresses sent to the Internet. The firewall should limit FTP, Telnet and Web access to specific machines. All external access should be enabled. A separate router will provide the external connection to further secure the internal network.

Remote Access Security

A remote access device provides a backdoor into the network and should have similar security to an Internet firewall. The main concern is former employees that have used the connection in the past. Usernames and passwords are required to gain access into any server on the network. However, guest access is provided to some drives that contain company-privileged documents. Any access provided through a remote access device will provide free Internet access and may be used for break-in attempts to network servers. Security on the remote-access devices should be username and password protected with an easily maintainable database. The security should conform to industry standards. Being able to use an already existing security database is highly desirable since it would limit the number of passwords a client would have to remember as well as the databases needing to be maintained. The security system should maintain logs of who dials in, how often, for how long, when and how much traffic is generated. Reports from these logs should be easy to generate.

Management

Management of the network will be done with SNMP and RMON management. The SNMP management platform will allow us to monitor real time traffic statistics. RMON management provides us with the ability to make quick device changes. All major network devices need to be SNMP compliant. RMON ability is preferred.

The network devices will be managed through a central interface that will monitor each device notifying the manager of a device failure or problem. The management console also needs to have the ability to track network usage between devices and provide reports showing trends.

A protocol analyzer is needed to provide preventative and detective measures. This will allow us to monitor the transfer of data and quickly isolate problems. The analyzer needs to allow us to monitor multiple segments at multiple locations.

A cable-testing tool is necessary in troubleshooting a cabling problem. It will ensure that the cable can support the data running through it reliably.

Redundancy

Since the chassis will be the backbone of the network, redundancy and solutions need to be provided to eliminate any foreseen as well as unforeseen failure points. Since the replacement of equipment may take anywhere from 4 hours to 3 days, an immediate solution needs to be available. The most likely failure points are the power supply and network card. To limit the possibility of a power supply problem, a secondary power supply will be purchased. To limit the impact of a card failure, extra ports will be left blank on each of the cards and a 16 port 10BaseT hub will be held in backup.

All devices need to be on a manageable UPS that can isolate power ports and shutdown systems prior to a UPS shutdown. A UPS will prevent power fluctuations and brief power outages from affecting the network. A manageable UPS will also shutdown connections if a long outage occurs, prior to a shutdown of itself. This may be accomplished by either having a central unit with power ports in each closet or by having a small UPS in each closet and a larger UPS at the central location.

Support

All major network devices need to be on a 24/7 service contract. All other network devices need to have a standard maintenance contract with next day replacement. A spare for each network device should be kept at each site.

Equipment Required

Total

Device	Quantity	Unit Cost	Extended Cost	Existing units	Total Cost
Backplane switch	2	$35000	$70000	0	$70000
Closet switch	7	5000	35000	0	35000
Hub	13	2000	26000	0	26000
Remote Access	2	10000	20000	0	20000
Storage rack	7	800	5600	1	4800
Patch panel (1 per 32 stations)	12	300	3600	1	3300
Patch cable (Copper)	280	6	1680	0	1680
Patch cable (Fiber)	18	110	1980	0	1980
Desk cable	280	10	2800	0	2800
Patch panel Fiber (1 per 12 stations)	7	350	2450	0	2450
Cabling (connectors)	75	40	3000	0	3000
Wiring (Copper)	320	140	44800	75	34300
Wiring (Fiber)	9	550	4950	0	4950
Enterprise Management server	1	35000	35000	0	35000
Enterprise Management client	2	5000	10000	0	10000
Hub Management station	1	6000	6000	0	6000
Hub Management client	2	1000	2000	0	2000
Network analyzer	2	15000	30000	0	30000
Cable tester	1	5000	5000	0	5000
UPS Central	2	5000	10000	2	0
UPS Closet	5	500	2500	3	1000
Training	6	2000	12000	0	12000
Total			**$334,360**		**$311,260**

Atlanta

Device	Quantity	Unit Cost	Extended Cost	Existing units	Total Cost
Backplane switch	1	$35000	$35000	0	$35000
Closet switch	4	5000	20000	0	20000
Hub	9	2000	18000	0	18000
Remote Access	1	10000	10000	0	10000
Storage rack	4	925	3700	1	2775
Patch panel (1 per 32 stations)	9	300	2700	1	2400
Patch cable (Copper)	190	6	1140	0	1140
Patch cable (Fiber)	16	110	1760	0	1760
Desk cable	190	10	1900	0	1900
Patch panel Fiber (1 per 12 stations)	5	350	1750	0	1750
Cabling (connectors)	75	40	3000	0	3000
Wiring (Copper)	220	140	30800	75	20300
Wiring (Fiber)	8	550	4400	0	4400
Enterprise Management server	1	35000	35000	0	35000
Enterprise Management client	1	5000	5000	0	5000
Hub Management station	1	6000	6000	0	6000
Hub Management client	1	1000	1000	0	1000
Network analyzer	1	15000	15000	0	15000
Cable tester	1	5000	5000	0	5000
UPS Central	1	5000	5000	1	0
UPS Closet	4	500	2000	3	500
Training	4	2000	8000	0	8000
Total			**$216,150**		**$197,925**

Littleton

Device	Quantity	Unit Cost	Extended Cost	Existing units	Total Cost
Backplane switch	1	$35000	$35000	0	$35000
Closet switch	3	5000	15000	0	15000
Hub	4	2000	8000	0	8000
Remote Access	1	10000	10000	0	10000
Storage rack	3	800	2400	0	2400
Patch panel (1 per 32 stations)	3	300	900	0	900
Patch cable (Copper)	90	6	540	0	540
Patch cable (Fiber)	2	110	220	0	220
Desk cable	90	10	900	0	900
Patch panel Fiber (1 per 12 stations)	2	350	700	0	700
Cabling (connectors)	0	40	0	0	0
Wiring (Copper)	100	140	14000	0	14000
Wiring (Fiber)	1	550	550	0	550
Enterprise Management server	0	35000	0	0	0
Enterprise Management client	1	5000	5000	0	5000
Hub Management station	0	6000	0	0	0
Hub Management client	1	1000	1000	0	1000
Network analyzer	1	15000	15000	0	15000
Cable tester	0	5000	0	0	0
UPS Central	1	5000	5000	1	0
UPS Closet	1	500	500	0	500
Training	2	2000	4000	0	4000
Total			**$118,710**		**$113,710**

The following items: Enterprise Management server, Hub Management station and Cable tester only require one item for the company. The Atlanta site will manage these items.

Options Considered

Network

Quotes for configuration and cost have been received from the following vendors: Bay Networks, Cabletron and Cisco. These companies are considered to be three of the top four networking companies. The fourth company, 3com, was not considered because of their lack of a strong backbone option. Each vendor was asked to provide their quotation broken into two categories called backbone and closet. Although the backbone and the closet categories may be thought of separately, to provide an easily maintained network both the backbone and closet network devices should be purchased from the same vendor.

Bay Networks

Backbone Option 1

Type	Model	Cost
Switch Chassis	5000N	$2500
Power supply	5001	2000
Supervisory Module	5110	1000
Network Management Module	5310SA	5400
10/100 Switch (16 port)	CH2004001	9600
Hub Management	636-03	3500
Support		4320
Total		**$28,320**

Strengths

- Dedicated switched 100BaseT ports
- Easy to switch out ports
- Hub management provides easy control over switch

Weaknesses

- Redundant power supply is external to unit
- Fiber modules are not cost effective
- The Backplane must be switched to move to a different environment

Backbone Option 2

Type	Model	Cost
Enterprise Switch	28115R	$14000
Hub Management software	636-03	3500
Support		3150
Total		**$20,650**

Strengths

- Dedicated switched 100BaseT ports
- Low cost 100BaseT switch
- Hub management provides easy control over switch

Weaknesses

- Does not come in a chassis model
- Only expandable with the addition of a new unit
- Does not allow for other interconnections besides 100BaseT
- Redundant power supply is external to unit
- Fiber modules are not available

Closet

Type	Model	Cost
Ethernet Switch	CW2001001	$4000
Support		720
Total		**$4,720**

Strengths

- Low cost switch
- Hub management provides easy control over hubs

Weaknesses

- Redundant connection not available
- Redundant power supply is external to unit
- Fiber modules are not available

Cabletron Systems

Backbone

Type	Model	Cost
Switch Chassis	9C106	2500
Power supply	9C206-01	2500
System monitor module	9C306	700
10/100 Switch (11 RJ45 ports)	9H422-12	10800
10/100 Switch (12 Fiber ports)	9H421-12	15750
VLAN Manager Server	SFVLAN-01NT	5000
VLAN Manager Client	SFVLAN-02NT	1000
Support		6005
Total		**$44,255**

Strengths

- Able to provide dual active connections to devices
- Thresholds can be set on dual active connections to overflow to other connection
- Two second change-over for redundant connections
- Independent Backplane
- Highest Backplane throughput at 5Gbps
- Easy monitoring and management
- VLAN management is icon based and easy to manage
- Reasonably priced Fiber module
- Hot swappable ports
- Internal power supply

Weaknesses

- Noisy unit
- High price

Closet

Type	Model	Cost
10BaseT Switch (24 port)	2E42-27	$4600
100BaseTX Switch interface	FE-100FX	350
Switch Support		425
Total (switch)		**$5375**
10BaseT Hub (24 port)	SEHi-24	1400
Hub Support		240
Total (hub)		**$1640**

Strengths

- Able to provide dual active connections to devices
- Thresholds can be set on dual active connections to overflow to other connection
- Two second change-over for redundant connections
- Broadcast storm protection
- Can handle up to four 100BaseT connections
- Easy monitoring and management
- VLAN management is icon based and easy to manage
- Reasonably priced Fiber modules on switch
- Internal power supply
- Moderately priced manageable hubs

Weaknesses

- Can handle up to four 100BaseT connections
- High price for switch

Cisco Systems

Backbone

Type	Model	Cost
Switch Chassis	WS-C5000	$3000
Power Supply	WS-C5008A	4000
10/100BaseTX	SW-X5213A	20000
Management software	CWPC-2.1-WIN	2500
Support		3500
Total		**$33,000**

Strengths

- Provides standby option for ports
- High speed Backplane at 1.2Gbps
- 100BaseT and ATM options available
- Dual internal power supply
- VLAN available with the use of a Cisco Router

Weaknesses

- Standby port can only be used if primary port fails
- VLAN available only with the use of a Cisco Router
- FDDI and Fiber Ethernet switching not currently available

Closet

Type	Model	Cost
10BaseT Switch	WS-C1900	$4000
Support		450
Total		**$4,450**

Strengths

- Redundant standby uplink port
- Smallest rack footprint
- VLAN available with the use of a Cisco Router

Weaknesses

- Standby port can only be used if primary port fails
- Only switches are available, no hubs
- FDDI and Fiber Ethernet switching not currently available
- VLAN available only with the use of a Cisco Router

Enterprise Management

Cabletron Systems

Type	Model	Cost
SpectroServer	SA-CSI1000	$8500
SpectroGraph Client	SA-CSI1001	4250
Enterprise Configuration manager	SA-CSI1008	8500
Cisco Router module	SM-CIS1001	2550
9H421,9H422 module	SM-CSI1066	1000
9E428 Module	SM-CSI1071	1000
8H02 Module	SM-CSI1068	1000
SEHI Module	SM-CSI1020	850
Support		4800
Training		1650
Total		**$34,100**

Strengths

- Interoperates on both Windows NT and Unix
- Excellent autodiscovery
- Discovers to the port level of devices
- Identifies nodes that have nodes based off of them (routers, bridges, etc.) and will only notify you of the highest point of failure
- Ability to assign importance to thresholds
- Policy based alarm filtering and forwarding
- Server to server distribution
- Has enterprise configuration manager option for alarms and verification of changes
- Future release will support Microsoft's SMS

Weaknesses

- Requires a separate client unit
- Does not currently support SNMP-v2

Hewett Packard

Type	Model	Cost
Network Node manager		$10000
Distributed management		10000
Extendable SNMP agent		250
Cisco Management		10000
Cabletron Management		10000
Support		7200
Training		2000
Total		**$49,450**

Strengths

- Industry leader
- Large number of third party software support
- Defines the length of time an event must exist before it is valid

Weaknesses

- Time consuming network creation process
- Each filter must be created separately and cannot be combined with other filters
- A failure on a device with devices under it will cause alarms in all the failed devices
- NT version does not have the features of the Unix version

Remote Access

3com/USRobotics

Type	Model	Cost
EdgeServer	001098	$7500
Dual PRI Card		2500
Total		**$10,000**

Strengths

- Industry leader
- Built in NT 4.0 server on a card
- Fully based on NT RAS
- Can host web pages local to server for Intranet support
- Flexible solution for adding T1, ISDN and Phone lines
- Excellent management on server
- Modems may be upgraded through firmware

Weaknesses

- Adequate speed
- Minimal management of multiple devices

Shiva

Type	Model	Cost
LanRover AccessSwitch		$13000
Dual PRI		$10000
Total		**$23,000**

Strengths

- Industry leader
- One of the fastest IP transfer rates
- Special speed enhancing client caching software
- Good management interface with Web management capabilities

Weaknesses

- Management did not work well with multiple clients
- Expensive per port solution
- Modems must be exchanged in an upgrade

ISP provided

Type	Model	Cost
T1 CSU/DSU	T120	$1000
Total		**$1,000**

Strengths

- No large hardware costs associated with service
- Helpdesk function performed by provider
- Management is the responsibility of the provider
- Lower cost per call with remote users utilizing local connections
- Technology upgrades are the responsibility of the provider

Weaknesses

- No management control over remote network
- Possible security hole

Network Analyzer

Network General

Type	Model	Cost
10/100 Sniffer server	SS-6045-1MA	$10000
Sniffmaster Windows console	SMW-0SX-1	4000
Reporter for Windows	NMA-WIN-1001	5000
Total		**$19,000**

Strengths

- Expert system for quick analysis and correction
- Distributed site monitoring with centralized management
- Minimal network management traffic
- Accurately measures utilization to 98% of maximum

Weaknesses

- High price

Other

Network General is regarded as being the best in the industry. No other options were considered. No other product comes close to the Network General's in features, management or analysis.

Cable tester

Fluke

Type	Model	Cost
Cable tester	DSP-100	$3800
Fiber option	DSP-100F	1000
Total		**$4,800**

Strengths

- Digital tester
- Provides location of problem in the wire
- Detects crosstalk problems
- Analysis in a quarter the time of an analog tester
- Graphical display
- Downloads to a PC
- Fiber option

Weaknesses

- Expensive solution

Microtest

Type	Model	Cost
Cable tester	COMPAS	$3000
Network addition	TCP/IP	1200
Total		**$4,200**

Strengths

- Provides traffic monitoring and analysis
- Easy to use
- Information downloadable to PC
- Ping test capability
- Can perform full LAN monitoring functions between a client and a port
- Graphical display

Weaknesses

- Less functionality with 100BaseT networks
- No fiber option available

UPS

American Power Conversion

Type	Model	Cost
UPS Main	Matrix 5000	$4400
UPS Closet	Smart-UPS 700	$475

Strengths

- Hot-swappable batteries
- Amplify and reduce voltage without switching to battery power
- Plug-in accessory slot for communications
- Environmental measuring device option
- SNMP manageability
- Easy to use Windows interface
- E-mail and pager support

Weaknesses

- Poor UPS MIB support for enterprise management packages

Liebert

Type	Model	Cost
UPS Main	UPS 5000	$4850
UPS Closet	UPStation D600	$475

Strengths

- Good SNMP support
- Good graphical tie in to enterprise management packages

Weaknesses

- DOS based management tool
- Supports only power boost not power reductions

Recommendation

Network

Each vendor's solutions will be adequate to meet our current needs while providing plenty of growth potential to meet our future demands. The solutions are mostly similar with some variation as to hardware configurations and networking solutions.

The vendor selected to meet Convergent's networking needs is Cabletron. They were selected because of the robustness of their solution for our network. They provide the best redundant connections through ports that provide concurrent dual paths to each connection, which in essence doubles the throughput capability of each connection. These ports could also have thresholds established to provide for quicker connections. All secondary power supplies are internal to the units. Their VLAN management package was the easiest to work with. They were also the only vendor to provide us with a cost effective Fiber option.

Enterprise Management

Both of these products would easily meet our needs. The Cabletron system however, provides extra functionality to more quickly locate and respond to problems. It also has superior reporting features and is more flexible to the addition of multiple sites.

Network Analyzer

Network General has been selected as our network analysis platform. They have consistently provided capabilities that the other vendors have yet to implement. Their systems include an intelligent monitor that determines many network problem symptoms and provide you with suggestions on solutions.

Cable Testing

The Fluke DLC-600 has been selected as our cable tester. It provided testing capabilities, that the other testers did not have to quickly determine and locate problems. It also provided a monitoring module for fiber.

Remote Access

This has yet to be determined. An ISP run virtual LAN appears to be the most flexible and cost effective means of supporting our remote client base. There are concerns that have not been addressed as of yet. In addition, the actual cost benefits have not been dealt with.

Both of the RAS vendors have excellent products and excellent reputations in the market. US Robotics provides a less expensive and more flexible solution that overshadows the speed benefits that the Shiva solution has to offer.

UPS

APC has been selected as our standard UPS. Their units provide the greatest power fluctuation support, hot-swappable batteries and a multiple monitoring options.

Stages for deployment

Atlanta network

Cost

Stage	Option 1	Option 2
Stage 1	$53865	$68365
Stage 2	51265	51265
Stage 3	67000	67000
Stage 4	10000	10000
Total	**$182,131**	**$196,632**

Stage 1 - Closet

Goal

Provide each desktop/client connection with a tested 10BaseT connection to the network. Provide a closet configuration that is easy to manage using patch panels. Segment the closet network with a switch on the top end of the network. Clients would connect to a port on a hub that connects to a port on the switch. High-end clients requiring more bandwidth connect directly to a port on the switch. Patch cords will connect the port connections on the patch panel with the ports on the hubs/switch.

Implementation steps

1. Rewiring of closets.

   There are two options to rewiring the closets to handle 10BaseT network connections.

   - Use existing terminal wiring for 10BaseT connections. Each station requires the replacement of the terminal jack with a 10BaseT jack. The closets require the replacement of the terminal punch-down blocks with 10BaseT patch panels. Each connection will be tested to insure that it can handle 10BaseT signals. Any connection not meeting this standard will be rewired.

Note: This wiring will not handle signals greater than 10Mbit and any station requiring higher bandwidth will require new cabling at $140 per connection.

Type	Number	Unit Cost	Total Cost
Cabling	0	$140	$0
Cabling (connectors)	220	40	8800
Desk cable	190	10	1900
Patch cable	190	6	1140
Patch panel	8	300	2400
Wall rack	3	800	2400
Rack shelf	3	75	225
Rack install	3	50	150
Total			**$17,015**

- Add new wiring for every station that does not currently have Cat5 wiring. Each station requires the replacement of the terminal jack with a 10BaseT jack. The closets require the replacement of the terminal punch-down blocks with 10BaseT patch panels.

Type	Number	Unit Cost	Total Cost
Cabling	145	$140	$20300
Cabling (connectors)	75	40	3000
Desk cable	190	10	1900
Patch cable	190	6	1140
Patch panel	8	300	2400
Wall rack	3	800	2400
Rack shelf	3	75	225
Rack install	3	50	150
Total			**$31,515**

The second option is the preferred option. It is more expensive but will prepare us to handle higher levels of data throughput without the replacement of wires. If the first option is selected, all the pieces with the exclusion of the wire may be reused when the connection is upgraded. The existing cabling has been tested supporting a network signal and will provide an acceptable low cost alternative.

Note: The table in Closet B with two older unused printers will be removed at this time. Closet C will have connections for New York, DFX modem and DSN modem moved to the computer room.

2. Purchase of equipment.

   The equipment will be set up and configured with one switch per closet and enough hubs to handle all the closet's connections. The switch and hubs will be configured and tested prior to the connection of client stations.

Type	Number	Unit Cost	Total Cost
Closet Switch	4	$5375	$21500
Hub	9	1650	14850
UPS Closet	1	500	500
Total			**$36,850**

   This equipment will replace the existing 14 Thinwire repeaters and 6 Thickwire repeaters found in the 4 closets.

3. Transfer of Network connections.

   Connection of the client stations will be done in an unobtrusive manner either on a weekend or on a gradual basis. The client workstations will need to be reconfigured to connect to the network via a 10BaseT connection.

Stage 2 - Network Backbone

Goal

Replacement of network backbone with a chassis based switched backbone having 100BaseT single segments running to each server and 100BaseFX multi-NIC segments running to each closet. Dual pair fiber cables will connect the backbone to each closet. This will provide the highest quality signal as well as data security to each closet. Two connections will provide a redundant link in the case of a wire failure.

Implementation steps

1. Wiring from computer room to closets. (This will be done during Phase 1 wiring.) All wire will be tested to ensure that it can handle 100mb traffic.

Type	Number	Unit Cost	Total Cost
Cabling (FX)	400ft	$2.25	$900
Cable installation	20hrs	$50	1000
Connector installation	8x4=32	$50	1600
Patch cable (FX)	16	110	1760
Patch panel (12 port FX)	5	350	1750
Total			**$7,010**

   FX stands for Fiber Optics.

2. Purchase and configuration of Backplane switch.

Type	Number	Unit Cost	Total Cost
Backplane Switch	1	$37755	$37755
Management software	1	6500	6500
UPS Central	0	5000	0
Total			**$44,255**

3. Testing of switch connected to current network backbone.

4. Placing and testing of servers on switched ports. This will be done during off-hours.

5. Placing and testing of closets on switched ports. This will be done during off-hours.

6. Management software configuration, testing and monitoring.

Stage 3 – Network Management

Goal

Create an environment where all network devices and servers are monitored and centrally managed. Provide the tools to quickly notify and correct problems that occur.

Implementation steps

1. An Enterprise management station to monitor the network devices. This will be configured to provide usage information and alarms for device failures for all the sites. Only one server is required for an organization.

Type	Number	Unit Cost	Total Cost
Management server	1	$35000	$32000
Management client	1	7000	7000
Training	2	2000	4000
Total			**$43,000**

2. Tools to correct network cable and protocol problems that occur.

Type	Number	Unit Cost	Total Cost
Network sniffer	1	$15000	$15000
Cable tester	1	5000	5000
Training	2	2000	4000
Total			**$24,000**

Stage 4 – Remote Access

Goal

Provide a fast, reliable and secure means of access to the corporate network. A chassis based system with 12 initial modem connections will meet the current need.

Implementation

1. Purchase and installation of new unit. This will obsolete the existing DEC700 and Hayes modem rack that is currently meeting this demand.

Type	Number	Unit Cost	Total Cost
Remote Access	1	$10000	$10000
Total			**$10,000**

2. Configure and test the server against the Radius server.

3. Replace existing unit.

Secondary Site Network

Goal

Provide a high-speed network with a centralized redundant backbone meeting current network needs while providing for a robust environment that will grow with future requirements. Provide an environment that minimizes the amount of downtime through redundancy and monitoring.

Provide clients with the fastest reliable connection to network resources required. Create an environment allowing for clients to run voice and video with little impact to the network.

Provide servers with individual network connections limiting the extra amount of data the server network cards have to process. Segment the client connections to a maximum of 50 users per segment.

Configuration

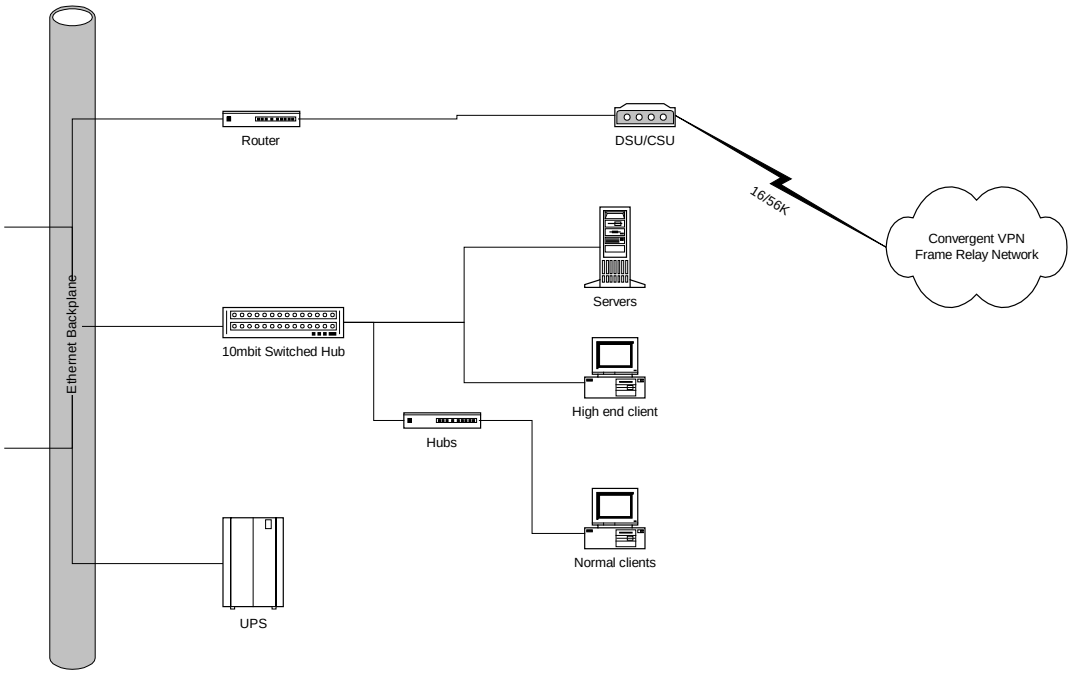

Backplane

The current need is for a 10mbit switched hub with dedicated 10BaseT switched ports for server and high end client connections. These sites will be configured initially with a regular hub instead of a switched hub. In this configuration, each major server would be given its own network segment, limiting the network traffic that each server handles to the packets designated for that server. Smaller sites, less than 10 persons, will typically not require a local server.

Client/Server connections

A 10 port switched hub will be at the base connection of the network, accepting the 10Mbit connections and passing the packets to secondary hubs and high end clients. A normal client in this configuration will have a 10BaseT connection to a closet hub. Growth should allow the network to easily segment itself into 100BaseT and 10BaseT segments depending upon client needs. Client connections should all be to patch panels from which patch cords connect them to the hubs. Each client connection should be to a level 5 jack in either a marked wall plate or desk box. A network map of each closet needs be maintained.

Protocols

Data transfers are to be through TCP/IP. 10BaseT will connect client and server devices.

Network devices

Switches

Switches need to have a minimum of 10 ports. Each switch should all be SNMP and RMON manageable for monitoring and remote management. Switches can be stand-alone or stackable units. For ease of management, all switches need to be from the same manufacturer.

Hubs

Hubs should all be SNMP and RMON manageable and either chassis based or stackable. A minimum of 16 ports per unit is needed. For ease of management, all hubs need to be from the same manufacturer.

Remote Access

Remote access will be provided through one of the primary sites.

Security

Security will be managed through the primary sites providing Internet and remote access to the site.

Management

Management of the network will be done through the use of SNMP and RMON Management. The SNMP management platform will allow us to monitor real time traffic statistics. RMON management provides us with the ability to make quick device changes. All major network devices need to be SNMP compliant. RMON ability is preferred.

The network devices will be managed through a central interface at one of the primary sites that will monitor each device notifying the manager of a device failure or problem.

Redundancy

Since the replacement of equipment may take anywhere from 4 hours to 3 days, an immediate solution needs to be available. The most likely failure points are the power supply and network card. To limit the possibility of a power supply problem a secondary power supply will be purchased. To limit the impact of a card failure, extra ports will be left blank on each of the cards and a 16 port 10BaseT hub will be held in backup. All devices need to be on a manageable UPS that can isolate power ports and shutdown systems prior to a UPS shutdown.

Support

All major network devices need to be on a 24/7 service contract. All other network devices need to have a standard maintenance contract with next day replacement. A spare for each network device should be kept at each site.

Equipment Required

Device	Quantity	Unit Cost	Extended Cost	Existing units	Total Cost
Backplane switch	6	$5000	$30000	0	$30000
Hub	8	2000	16000	0	16000
Storage rack	6	800	4800	3	2400
Patch panel (1 per 32 stations)	6	300	1800	4	600
Wiring	150	140	21000	125	3500
UPS	6	1000	6000	2	4000
Travel	2	500	1000	0	1000
Total			**$80,600**		**$57,500**

Options Considered

The closet switch and hubs selected for the Primary sites will be used for the Secondary sites.

Recommendation

Each site will start with one 10mbit switch. The servers and router will each get a dedicated port. Hubs will be added as need arises. A spare hub will be kept on site for emergencies.

Stages for deployment

Cost

Stage	Cost excluding WAN	Cost
Plano and Detroit	$18800	$29200
New York and WoodCliffe Lake	15600	15600
Silver Spring and Lawrenceville	15800	15800
Total	**$50200.00**	**$60,600**

Plano and Detroit

Goal

Provide a network to effectively handle the local needs of the site. This would be done in conjunction with the Wide Area Network configuration for these sites.

Implementation steps

1. Site survey of Plano and Detroit

   These are two new sites that have minimal or no currently established network. This will be done in conjunction with a LAN analysis of the locations. The survey will determine location of wiring closet, number of networked machines, needs of the users and the wiring needs of the location.

Type	Number	Unit Cost	Total Cost
Plano travel	1	$500	$500
Detroit travel	1	500	500
Total			**$1,000**

Travel is considered as part of the Wide Area Network expense.

2. Purchase of equipment

Type	Number	Unit Cost	Total Cost
Small-office Router*	2	$4200	$8400
56K CSU/DSU*	0	1000	0
Switch	2	4600	9200
Network cards	8	100	800
Storage Rack	2	800	1600
Patch Panel	2	300	600
Cabling	40	140	5600
UPS	2	500	1000
Total			**$27,200**

** Wide Area Network equipment purchases will be done at this time.*

3. Installation and configuration

   The configuration and testing of the equipment will be done at the Atlanta location prior to it being installed at the remote location. An administrator will need to be at both the Atlanta location and the remote location during the installation and testing of the equipment.

Type	Number	Unit Cost	Total Cost
Plano travel	1	$500	$500
Detroit travel	1	500	500
Total			**$1,000**

Travel is considered as part of the Wide Area Network expense.

New York and WoodCliffe Lake,

Goal

These sites have currently existing local networks that sufficiently meet the current need. The addition of the standard equipment will allow the site to be monitored more effectively. It will also provide the site with the same capabilities as the other sites.

Implementation steps

1. Purchase of equipment

Type	Number	Unit Cost	Total Cost
Backplane switch	2	$4600	$9200
Hub	2	1650	3300
Storage rack	1	800	800
Patch panel	1	300	300
UPS	2	500	1000
Total			**$14,600**

This equipment will obsolete the two existing hubs located at these sites. These hubs will be kept in reserve in the case of an equipment failure.

2. Installation and configuration

   The configuration and testing of the equipment will be done at the Atlanta location prior to it being installed at the remote location. An administrator will need to be at the remote location during the installation and testing of the equipment.

Type	Number	Unit Cost	Total Cost
New York travel	1	$500	$500
WoodCliffe Lake travel	1	500	500
Total			**$1,000**

Silver Spring & Lawrenceville

Goal

These sites have currently existing local networks that sufficiently meet the current need. The addition of the standard equipment will allow the site to be monitored more effectively. It will also provide the site with the same capabilities as the other sites.

Implementation steps

1. Purchase of equipment

Type	Number	Unit Cost	Total Cost
Backplane switch	2	$4600	$9200
Hub	4	1650	6600
Total			**$15,800**

2. Installation and configuration

   Configuration of the equipment will be done in Atlanta. The installation and testing of the Lawrenceville equipment will be done with the use of an Atlanta administrator visiting the site. The installation and testing of the Silver Spring equipment will be done by the local administrator.

Appendix: Concepts and Definitions

Definitions: Network Components

Segment - The smallest piece of a network on which stations can exchange data without intervention from another intelligent device.
Extended Segment - A number of segments joined together by bridges. Any broadcast or multicast made on the extended segment should be seen by all stations on an extended segment.
Network - The term itself has come to be rather ambiguous, referring to a segment, extended segment or internetwork. We often call any of these "The Network."
Internetwork - A set of segments or extended segments joined together by a router.
Unicast Packet - A data packet addressed to a single station. An example might be data from a client to a server.
Multicast Packet - A data packet addressed to a group of stations. The destination address is formed in such a way that stations realize that the packet may be destined for many other stations.
Broadcast Packet - A data packet addressed to any and all stations on the local segment. Broadcasts are often used by stations who have just joined the network - broadcasts are made to find out information about the segment that has just been joined.
Repeater - A device that facilitates connecting stations onto the segment. It does not understand network addresses - it merely copies data bit by bit from and to the physical media to which it is attached. On Token-Ring segments, this device is often called a Media Access Unit or MAU. A repeater is not considered an intelligent device.
Bridge - A bridge is used to connect two or more similar segments together (for example, Token-Ring to Token-Ring or Ethernet to Ethernet). A bridge has two purposes. The first is to extend the length and number of stations that a segment can support. Secondly, the bridge reduces overall traffic flow by only passing data packets that are not destined for a hardware address on a local segment. All broadcast and multicast traffic must cross a bridge - since no true destination can be known. In recent years, bridging technology has been used between dissimilar media (for example, Ethernet to FDDI), this sometimes may cause problems as we will see later. A bridge is considered an intelligent device. (See also bridging.)
Router - Sometimes called a gateway, it is used to connect two or more (potentially extended) segments. The segments may be similar or dissimilar. Routing information beyond the hardware address must be contained within the data packet. Virtually no broadcasts or multicasts are ever propagated across a router since no exact destination information is typically contained within these packets. Hardware addresses have only local significance to a router - higher level routing information is globally significant. (See also routing.)

Definitions: Glossary of Terms

10BaseF - Ethernet standard for running over fiber optic cable.

10BaseT - A cabling standard that allows you to transmit Ethernet data over unshielded twisted pair wire. It derives its name from a baseband network that transmits at 10 megabits per second (Mb/s) over twisted pair cable.

10Base2 - Also called Thinnet, it is a cabling standard that allows you to transfer Ethernet data across coax cables. It also allows you to place multiple connections on a link.

100BaseX - Fast Ethernet standard. Allows you to transmit Ethernet data over unshielded twisted pair wire. It derives its name from a baseband network that transmits at 100 megabits per second (Mb/s) over twisted pair cable.

Backbone - Also called backplane. The main segment of a network, to which all other segments are connected.

Backbone network - The major transmission path for network interconnection. It interconnects all nodes of the LAN network.

Backplane - See Backbone.

Bridge - A device that links two local area networks of the same type, passing only traffic that originates on one network and is intended for the other network.

Client – The user's computer or a peripheral networked device. High end clients require faster access into the network for data access.

DSU/CSU - Allows a network to traverse a phone connection to a remote site with the addition of a bridge or router. Works the same way as a MODEM on a dedicated phone line.

Ethernet - A network communications standard allowing you to transfer data between networked devices.

Fiber - A means of transferring data from point to point using light as opposed to the electrical signal that is transferred across standard copper cables.

Firewall - A routing device that limits external access to a network from the internet.

Frame Relay - Type of inter-site connection allowing you to create a Virtual Private Network (VPN) across phone lines. Creates a switched network across the WAN. Allows multiple sites to be interconnected with a single link between them. Allows you to buy a lower normal data transfer rate with the possibility of occasional bursts of data to a higher level.

Hub - Allows you to attach devices to the network. Also called a repeater or concentrator.

Internet - A group of interconnected organizations allowing for data transmissions between them.

Local Area Network (LAN) - A local collection, usually within a single building or several buildings, of personal computers and other devices connected by cabling to a common transmission medium, allowing users to share resources and exchange files.

Patch panel - A system of connecting blocks, patch cords, and backboards that facilitates administration of cross-connect fields for network moves and rearrangements.

Repeater - See Hub.

Router - A protocol-specific network layer device that connects two networks and may direct traffic based on network resources availability. Forwards packets between subnetworks on the basis of network-layer destination addresses.

Server – A machine which provides file, print and application access to client machines. Primary servers or High End servers run major business applications or provide connectivity to primary areas. A Primary server is critical to the business and has much data flow through it. Secondary servers or Low end servers have specific purposes that require minimal data flow and user access.

Switched Hub - Allows you to attach devices to the network, dedicating an entire network segment to that link.

T1 - A dedicated phone connection that allows you to transfer data and voice at a speed $1/10^{th}$ that of a normal network.

Thinnet – See 10Base2.

Twisted Pair - A form of wiring used to transmit baseband electrical signals. It consists of two insulated wires loosely spun to help cancel out any induced noise in balanced circuits. An example is a phone line.

WAN - A Wide Area Network is any network that extends beyond the single site.

X.25 - A standard for transmitting protocols through public data networks.

Concepts: Shared-medium Networks

Sharing data on a network means multiplexing it on the basis of either frequency or time, and arranging for some sharing, or contention, scheme.

Multiplexing

Frequency Division Multiplexing (broadband)

Some of the earliest data networks devised used Frequency Division Multiplexing or FDM. These networks are also known as broadband networks. Just as we divide up the radio spectrum into channels, so can we divide up the spectrum over a cable. A possible example would be to use, say, 200 MHz of bandwidth and divide it into 20 channels of 10 MHz each. Each 10-MHz channel could, theoretically, be used to transmit up to 10 Mbps of data.
Apart from the obvious advantage of providing more bandwidth, the analog techniques used can also permit signals to be transmitted over fairly long distances on standard cabling - sometimes up to ten times further or more than straight digital signals on the same cable. So broadband systems can be used to provide very high data rates over fairly long distances.
The advantages of broadband seem significant. However, it is seldom used for local area networks today. In fact, other than digital systems that piggy-backed on private CATV systems, broadband is almost strictly the domain of telephone carriers. There are good reasons for this. First, analog broadband networks are expensive to build and maintain. They must be tuned and retuned as the network is extended. The equipment required to do this is expensive and the expertise is rare. Equipment that is attached to a broadband network is also expensive as it must have a digital-to-analog modem and transmitter - much like any radio system would have.
The real nail in the coffin for broadband systems were the time of their introduction. Indeed, a good broadband network could accommodate 500 Mbps or more data traffic and do it over a few square miles. The problem was that in the mid-70's to early 80's no one had computers that use 500 Mbps of bandwidth and we had barely begun to build local area networks. Few had even thought about extending the network over areas like a few square miles. About the only commercial example of broadband being used for LANs was IBM's PC-LAN products. It used only two fairly narrow channels of a broadband network and hence only could transmit and receive data at about 1 Mbps.

Time Division Multiplexing (baseband)

If sharing a network by dividing up the frequency spectrum doesn't fly - then sharing by dividing up time is the only other alternative. The idea here is to use baseband signaling - essentially putting digital signals right on the wire - and sharing it by devising mechanisms for

computers to take turns accessing the bandwidth. Through the 70's, a few companies devised ways to build baseband networks. The four most popular systems where IBM's Token-Ring, Xerox's Ethernet, Advanced Interlink's ARCnet and Apple's LocalTalk. These all use a few basic techniques for arbitrating bandwidth.

Packet Contention

CSMA/CD Networks

When Xerox began building high-end printers that would produce copy directly from workstations, they needed a mechanism to get images from the moderately priced workstations to the fairly expensive printers. In the 70's at Xerox's Palo Alto research center, work was under way to develop a shared data network that could do the job.

The goal was to find a simple algorithm - one that could be implemented in the fairly basic silicon available at the time. Ethernet was the result of the efforts and the method for sharing the wire was called Carrier Sense Multiple Access with Collision Detection or CSMA/CD.

The idea was a fairly simple one. Listen before you talk - that's the multiple access part and stop talking if you hear some one else, that's the collision detection part.

Essentially, if a station found no traffic on the wire, it could start putting data onto the network. If, while it was putting data on the network, it sensed a collision, the station would stop immediately and wait a random amount of time before transmitting again. The specifics of Ethernet were designed so that once a station got a few bytes into its transmission, all stations on the segment should be able to detect the signal and remain silent until the transmitting station finished. So, on a properly implemented Ethernet, the collision rate should only be a few percent of the packets even when the network was 60 to 70 percent busy or more.

With its 10-Mbps data rate, Ethernet was easily able to handle the transmissions of many PDP-11 or VAX 750 machines (the prominent candidates for networking at the time) without difficulty. Networking a computer at the introduction of Ethernet was not a trivial decision. While today it is rare to find an Ethernet adapter with a price tag above a few hundred dollars, a network interface from the late 70's was likely to cost $5,000 or so.

Today, however, Ethernet has been rendered a single-chip solution and provided by many vendors right on the motherboard of $1,200 PCs. In these configurations, Ethernet now adds no more than $25 to the price of a PC.

Token Passing Rings

Token-Ring, and later FDDI, which employs token-ring techniques, is a newer and more complex method for sharing network bandwidth. Before we get into the specifics of Token-Ring, lets talk a bit about why anyone would find it necessary to build a more complex technology than Ethernet.

In the early 80's, shortly after the introduction of the IBM PC, it was observed that on a network with fast minicomputers and comparatively slow PCs, the PCs could be starved on the network. Further, due to the variety and quality of Ethernet implementations available, many found that Ethernet was unusable when network utilization reached only into the 20 or 30 percent range or so.

Because Ethernet employed random back-offs and was subject to network hogs, it was thought to be unsuitable for mission critical networking. The term bandied about was non-deterministic. In other words, there was no way to mathematically assure that a given station could transmit a given amount of data within any particular time frame. In fact, the folks who thought up Ethernet could show that they could make assurances within high probabilities - but that wasn't good enough.

Token Ring's approach was to arrange stations into a logical ring. Once the ring was formed, a token was generated and passed between the stations on the ring. If a station had data to transmit, it removed the token from the network, transmitted its data and then passed the token along to the next station. Each station could transmit data up to some maximum time or until it

was out of data to send - whichever was shorter. In this way, every station is assured access to the network regardless of the station's speed or network interface design.
This sounds really good, doesn't it? There, of course, are trade-offs. First and foremost, the algorithm behind Token-Ring is an order of magnitude more complex than Ethernet. When the ring is running normally, Token-Ring seems pretty simple. However, consider the added complexity of loosing a token, finding two tokens on the network, unexpectedly loosing a station from the ring or bringing a new station into the ring.
Each of these complex cases must be handled correctly. It's just a lot harder to implement than Ethernet's listen-then-talk model. Add to this the fact that IBM picked 4 Mbps as the base data rate for Token-Ring and you can see how simpler, faster Ethernet might be more attractive. Token-Ring also had the problem that if it became deterministic (that is, every station held the token as long as it could), it, too, became painfully slow to use.
If ASIC technology were as advanced then as it is today, Token-Ring might have Token-Ring networks everywhere. It wasn't very advanced, Token-Ring cards were and are expensive when compared to Ethernet and it is hard to make an argument for using IBM's style of Token-Ring.

Concepts: Switched Networks

In either case, whether your technology is Token-Ring, Ethernet or some other shared medium technology, the success of the shared network depends on each station having comparatively little data to transmit or on having relatively few stations on each segment of the network. The more a station can saturate a network, the better it is to have fewer stations on the network. Switching is basically a technology that is meant to facilitate reducing the number of stations per segment. The term switching is taken from the telecommunications industry where the devices that routed telephone calls were originally called mechanical switches. Switching has come to imply an architecture where any inbound traffic can be redirected to any outbound port with relatively little concern for traffic loss or congestion.
The way in which a switch decides how to direct traffic could be by almost any mechanism. It could use bridging or routing techniques, or it could use some other mechanism to pre-determine the path that subsequent data transmissions will take.

Packet-based

To realize the goal of very few stations per switched port, the switch market aims to provide the bandwidth advantages of bridges and routers at a price closer to that of a repeater. The only way to achieve this sort of price point is to rely heavily on ASICs and other custom silicon which has only recently become reasonably cheap to produce.
These chips analyze packets and determine how to direct the packet. Virtually every major networking vendor either has developed such chips or is working on them. As new generations of chips emerge, the chip count on switches go down and so do the prices. In fact, in the current generations of Ethernet switches the high-speed uplink ports are the most expensive pieces of the switch. In some cases, an FDDI or ATM port can cost as much as 10 to 12 Ethernet ports.

Cell-based

As good as packet-based switching is, there are certain types of traffic for which they are not ideally suited. Further, the complexity associated with handling variable length packets, each containing their own detailed addressing, makes packet-based switching an expensive proposition.
Cell-based switching is a solution aimed at handling non-data traffic (for example, voice and video) along with data. One problem with router and bridge-based systems is the latency that they introduce to the network. Routers and Bridges usually fully capture and then forward a data packet. If the router or bridge can do this instantly, up to 1.4 milliseconds of delay is

introduced to Ethernet packets traveling through them. Routers, in particular, are likely to introduce more delay because they often must process the packet with a single central CPU. For data networks, these delays are usually not serious - in fact, they usually go unnoticed. However, these delays are significant for video and audio traffic. Cell-based networks in general, and ATM in particular, are architected to handle general digitized data including voice, video and computer-originated data.

The idea behind cell-based networks is to chop standard data packets into much smaller fixed-length cells. In ATM's case, these cells are 48 bytes long with another five bytes for addressing and control.

One fact that should immediately become obvious is that a five-byte address field is too small to hold even a single six-byte physical address. Obviously, there must be something else going on. Indeed, ATM requires that a route be determined before data starts flowing.

The five-byte addresses are only relevant from an end station to a switch or between switches. Each switch then builds a table that includes the translation of incoming addresses with out-going addresses.

By predetermining the flow of data using that predetermined path throughout a data exchange, ATM assures that cells will arrive in the proper order at the end station. In fact, ATM includes no mechanism for retransmission of cells. Higher order protocols must take care of any data lost in the ATM network. However, when cells do arrive in the correct order and an in a timely fashion, it can be a simple matter (that is - cost effective) to retrieve the data, voice or video information contained within the cells.

We will touch on ATM only lightly here. A much more in-depth discussion will be included later.

Concepts: Network Addressing

Since all traffic on a network must be seen by each station on the network, there must be some way to designate which data packets are destined for which stations. In other words, each station must have an address that is unique to its hardware.

It seems clear that if two stations are on two completely separate networks, they really don't need to have different hardware addresses. After all, they'll never see each other's traffic. Up to this point, we haven't talked about what defines a network, so we must further define some terms. These terms are used to describe the pieces of hardware that tie a network together.

It should be obvious that hardware addresses and their uniqueness is most important on segments and extended segments. To that extent, the hardware address of any station could be set by the local administrator of any particular segment.

This is how ARCnet works. Up to 255 addresses may be configured for ARCnet stations. These addresses are usually set by configuring jumpers on the network card. Apple's LocalTalk network takes a slightly different tact. Rather than worry about setting addresses, each station just picks one and then broadcasts to see if any other station is using it. For small segments, both of these techniques work well. However, the bigger the network, the less comfortable this scheme becomes.

Ethernet, Token-Ring and FDDI have employed a different technique. Their hardware addresses are considerably longer - six bytes long rather than just one or two. The upper three bytes are assigned to hardware manufacturers who then assign the lower three bytes themselves. The scheme allows for 16 million manufacturers, each of whom can then assign 16 million addresses to their products.

This scheme was originally administered by Xerox for Ethernet until the standards for Ethernet were turned over to the IEEE, which now administers hardware address assignments.

Each packet that is sent on a network must contain a source and destination hardware address. Some topologies have allowed for different length addresses as well as local assignment of addresses. However, in almost all cases, the globally administered six-byte addresses are used.

Concepts: Bridging

How it works

We've described the basic functioning of bridges. They essentially build a list of known physical addresses and note which port those addresses reside. These addresses are valid only for a certain length of time, after which, if no traffic has been seen from the address, it is removed from the table. Any packet that has a destination address unknown to the bridge is retransmitted on all ports of the bridge except the originating port.

If the bridge is a little smarter, it will determine if the address is known on a different port and only transmit the packet to the port that contains the known address. This is essentially the functioning of a very basic switch.

Keep in mind that broadcast packets have no known destination and therefore must be sent through all ports of the bridge. This can lead to problems on large networks.

Other problems can occur when media are mixed on a bridge. The most significant problem here occurs when one media has a different maximum allowable packet size than the other (known as MTU or Maximum Transmission Unit).

Some protocols provide for a mechanism called MTU discovery. This is fine as long as the stations are using some connection-oriented protocol and it makes sense to store the discovered MTU. However, if they are using a connectionless protocol, it makes no sense to rediscover the MTU with each transmission.

When to use it

In general, the solution to the MTU problem is to make bridges that are at least smart enough about higher-layer protocols to participate in MTU discovery if it is used. In the case of TCP/IP specifically, the bridge must be capable of fragmenting packets.

Packet fragmentation is normally performed by routers and can be a taxing task for some of them. IP fragmentation is a process of taking large packets and breaking them down into packets as small or smaller than the MTU of the destination media. Bridges (or switches for that matter) that can perform IP fragmentation that are generally able to handle any protocol that might be thrown at them.

The bridge's requirement to pass on all broadcasts can cause problems as well. On large networks, usually ones with multiple bridges and hundreds of stations, the propagation of broadcasts through the network can result in other stations creating broadcasts as well. This is known as a broadcast storm. They can last a while and consume as much network bandwidth as is available.

A more common problem occurs when a significant number of broadcasts occur on a fast backbone and have to be propagated to slower media. If broadcasts consume 5 percent of the bandwidth on 100-Mbps media, it probably isn't a problem. However, those same broadcasts would saturate a 4-Mbps Token-Ring segment or take 50 percent of the available bandwidth on an Ethernet segment. That is a significant problem.

Most bridges provide mechanisms for filtering broadcasts and in some cases, this may provide an adequate solution. However, on larger networks at least some routers should be used.

Concepts: Routing

How it works

TCP/IP, IPX/SPX, AppleTalk and other protocols all operate at the network layer. That is, they employ at least two levels of addressing where bridged systems have a flat, universal addressing scheme. The bridging technique of forwarding packets with unknown destination addresses doesn't scale to global proportions, indeed, it doesn't scale well past a few hundred nodes.

By dividing addresses into a network field and a node field, it is possible to more accurately direct packets. In fact, just this two-level hierarchy is enough to build a global network.
If a router's job were just to steer packets around an internetwork, we'd probably have much cheaper routers than we do. The fact of the matter is that routers usually do much more. They also store and rebroadcast information about the internetwork, keep protocol dependent tables, enforce administrative rules on network traffic and provide redirection for special purpose broadcasts. All of this is fairly CPU intensive, and routers, consequently, tend to be bottlenecks in networks.

When to use it

Keeping this in mind, routers do have their uses and should be considered for certain reasons. There is no better way to erect a wall between two different parts of an organization than the use of a router (ex. separating marketing and engineering). Routers are also the only game in town when it comes to connecting your private network to a public network (like the Internet). Furthermore, routers are the best means of connecting networks via comparatively slow wide area networks. If you're paying for wide area bandwidth, you'll want all the control possible over the data that flows across the network.
These are the instances where there is no substitute for routing. However, in the local area network, routing is not the best way to increase the overall bandwidth within your network. That is best done with switches that have some routing smarts.

Concepts: Switching

What's the difference?

Switching has matured beyond simple multiport bridging. There are a number of important features that not only make switching the most economical way to get more bandwidth in your network, they also make a switched network much easier to administer.
In terms of bandwidth, switches provide high speed, low latency bandwidth. Latency is usually much lower for a switch than for a router as there is usually less processing occurring in a switch and multiple processors (most often ASICs). In instances where traffic is flowing between like media (say Token-Ring to Token-Ring), switches can begin retransmitting the packet before having completely received the packet. This is called cut-through bridging (as opposed to store and forward) and can reduce latency more.
On the administrative side, virtual LAN (VLAN) is now a feature commonly found on switches. VLAN technology addresses some of the flaws in bridging without necessarily introducing the complexities of routing.
The idea of VLANs is to take some group of ports on the switch and treat them together as a LAN segment. The net effect of this is to create broadcast domains since all traffic is directed only at the port for which it is destined.
Traffic flowing between VLANs must be routed. However, VLANs can usually encompass many more segments than a regular bridged network might have. This reduces the number of router ports needed and often results low levels of traffic between VLANs (often just mail). Some switch vendors have built routing functions into their switches and others have not. While some route IPX, IP, AppleTalk and DECnet, most only handle IPX and IP - bridging all other protocols. Depending on the configuration of your network and the ease with which you can reconfigure your network addresses, routing may be worth its additional cost.

When to use it

Switching, particularly as a means to accessing an ATM backbone, will likely be the preferred mechanism for building high bandwidth networks over the next three to five years. Virtually any network that has outgrown a single segment design can benefit from switching. Probably the bigger issue is converting networks that currently employ routers. Reworking network

addresses can be a challenge and in some environments, it can be almost impossible. (See also switched networks.)

Today's networks

Ethernet

Physical characteristics

Ethernet has been essentially described in four specifications from the IEEE. These build upon the work done initially by Xerox and later by Xerox, Intel and Digital Equipment Corp. These specifications involve various types of cables, connection rules and other hardware considerations. They all employ the general CSMA/CD algorithms discussed earlier.
Note that in order for CSMA/CD to work properly, there must be a minimum packet size on the network. That minimum size has been set at 64 Bytes and the length of the various network segments where more than two transceivers can exist has been determined based upon the propagation speed of data over the media.

Topologies

10BASE-5 or Thick Ethernet

10BASE-5 is the original Ethernet system. It employs a quarter of an inch diameter, 50 ohm coax cable (with a minimum bend radius of 10 inches). 10BASE-5 segments can run in length up to 500 meters with as many as 100 transceiver connections spaced at least 2.75 yards apart. 10BASE-5 transceivers access the media by piercing the thick coaxial cable. These transceiver taps are known as vampire taps. Since they don't actually require breaking the physical cable, the electrical signals over the cable are typically fairly clean.
10BASE-5 systems were originally envisioned to be cheap and fairly easy to build. The large cable needed simply to be run by rooms where computing equipment would be located. Taps would be made into the cable by using external transceivers. As it turned out, the requirement of an external transceiver and the thick cable, which was expensive and difficult to work with, limited to use of 10BASE-5.

10BASE-2 (A.K.A Thin Ethernet and CheaperNet)

Thin Ethernet was a fairly popular specification and is still used in many environments today. With a maximum segment length of 203.5 yards, it requires that the 50 ohm cable be only .2 inches thick (a bend radius of two inches). It also uses standard BNC connectors and "T's" to provide access to the media. Typically, T's are connected directly to the back of network interface cards, thus eliminating the need for an external transceiver .
Only 30 transceivers can be inserted onto a Thin Ethernet segment and they must be spaced at least 19.69 inches apart. 3Com was heavily involved in developing Thin Ethernet hardware, much as they are today. Their hardware was able to handle slightly longer segments, up to 220 yards in length. Unfortunately, mixing other vendors equipment into an environment where cable runs exceed 203.5 yards can cause problems. For this reason, keeping total lengths to 203.5 yards is a good idea.

10BASE-T

Neither of the coax-based Ethernet specifications lent themselves well to the structured wiring plants that telco workers had been building for decades. Using telco-style wiring was seen as necessary if networked computers were to populate most every desk in the corporate world. Various vendors realized this and began making Ethernet implementations that could run over standard category 3 twisted pair wiring. The same wiring that drives most every telephone in the world.

The standard eventually came down to supporting 110-yard segments of category-3 cable with a maximum of two transceivers per cable (the end station being one and the hub being the other). Standard RJ-45 phone jacks are used for host connections and transceivers are almost always built onto the network interface card, making the connecting hardware and card very economical.

10BASE-F

10BASE-F is an Ethernet over fiber-optics specification. Its main purpose is to provide long Ethernet runs and electrical isolation either up building risers or between buildings.
Like most other multimode fiber specifications, 10BASE-F segments can go as long as 1.24 miles and accommodate only two transceivers.

Token Ring

Physical characteristics

Token-Ring is heavily used in IBM mainframe environments. It's standardization has taken place in the IEEE 802.5 committee. Token Passing need not be a ring topology, IEEE 802.4 defines Token Bus. However, the ring topology is good since a station that put data on the ring can also take it off, therefore knowing whether the data made it all the way around uncorrupted.
Transmission speeds of 4 and 16 Mbps have been standardized. Data units are always at least 22 bytes long and their maximum length is determined by the Token Holding Time (THT), which usually allows for packets up to approximately 4,500 bytes.
One station on each Token-Ring segment will act as the monitor. This is usually the first station to enter the network, but each station must be capable of acting as the monitor. The monitor has a few very important responsibilities. It must create the original token, compensate for ring jitters, be able to store one whole token so that the token is occasionally fully removed the ring, remove unowned or mangled packets from the ring and finally establish the order of stations on the ring.
Each station must receive and retransmit each packet on the Token-Ring network, so the major concern in Token-Ring is the aggregate differences between the clocks on all of the Token Ring cards. This difference in clock rates - and the potential data loss - is known as 'jitter.' Almost all of the difficulties associated with multivendor Token-Ring networks center around jitter problems.

Topologies

Token-Ring can make use of a wide variety of topologies. The most common today is through active hubs with end-station runs using telephone grade wire. However different limitations exist for four different Token Ring topologies. For each 4-Mbps and 16-Mbps Token-Ring, there rules governing their use over both unshielded twisted pair (UTP) wiring as well as shielded twisted pair (STP).
Until as recently as late 1991, IBM was unwilling to admit that 16-Mbps Token-Ring could or should be run over UTP wiring, preferring STP wiring. Indeed, the Manchester II encoding used to put Token-Ring data onto a wire (it's the same encoding mechanism as used for Ethernet), requires a physical signaling rate of 32 MHz, and the FCC is quite careful about systems that run at these rates as they can interfere with a number of broadcast technologies. Companies such as Proteon and Synoptics (now Bay Networks) had shown an UTP 16-Mbps system commonly and IBM has since agreed to standardize the technology.
For any ring, 4 Mbps or 16 Mbps, the maximum number of stations has been set at 260 stations. The limit on the number of stations is due to total jitter present throughout the ring. Each station has a small buffer that can be used to compensate for differences in clock rates around the network. Only the monitor station, however, is responsible for correcting the ring's apparent jitter. This allows for the one-bit delay between stations. If more than 260 stations

are present on a ring, the monitor may not have enough "room" in its latency buffer to account for all the jitter present in the ring.
In reality, it is difficult to build a ring of 260 stations. Somewhat less than 100 is probably a more realistic number. Each adapter must maintain its own clock and each one on the network may meet the requirements for a 260-station ring. Few stations is more prudent.
Passive token ring MAUs or Multistation Access Units where the original means for building a ring as envisioned by IBM. MAU's were unpowered devices that simply allowed for star-shaped rings, thus permitting a structured wiring plant (like the telephone network.) Passive MAUs have given way to active devices that have a number of advantages. Active devices can act as repeaters, and thus elevate concern for signal degradation due to overall ring length. Whether active or passive, each MAU has a Ring-IN port and a Ring-OUT port. These ports are used to extend the ring beyond the MAU. As with Ethernet, Fiber can be used on these ports for distances of up to 1.24 miles, and up to 550 miles of IBM type 1 STP cable. Conversely, NICs are supposed to be able to drive signals on up to 770 miles of type 1 STP cable. Lobes from MAU to end station and back may not exceed 110 miles when using type 1 cable (since 110 + 550 + 110 would total to the 770 miles a single station can drive.)
If UTP is used for end station runs, then no more than 72 stations are permitted on the ring. Essentially all type 1 cable measurements can be used, however a conversion of 1.1 miles of type 1 cable to .495 miles of type 3 cable must be made.

FDDI

Physical Characteristics

Fiber Distributed Data Interface (FDDI) was the first standardized 100-Mbps technology. From day one, it was and is envisioned to be a backbone technology. Station management, redundant links and fairly flexible architecture give FDDI its backbone flavor. They also make it a fairly expensive technology - especially compared with other 100-Mbps technologies like Fast Ethernet and 100VG-anyLAN.
As the name implies, FDDI was intended to run on fiber. Standards have been written for data grade UTP (category 5) as well as STP wiring. FDDI raw baud rate is actually 125 Mbps as FDDI's minimum data units are expanded from four bits to five bits. The additional bit allows bit patterns to be chosen so that series of '0's are not permitted. FDDI uses a non-return-to-zero, invert on 1's encoding technique. By not allowing consecutive 0's, FDDI can maintain its clocking at a frequency of 125 MHz rather than the doubled frequencies required by Ethernet (although not Fast Ethernet) and Token-Ring.
FDDI uses token passing for its media access arbitration just as Token Ring does. However, rather than specify a flat Token Holding Time, FDDI uses a Target Token Rotation time. The Token Holding Time is then calculated by dividing the target rotation time by the number of active stations on the ring. This addresses one of the faults of Token Ring - that crowded rings get quite slow, and stations may not get a chance to send data for 50 milliseconds or more. On most FDDI rings, the TRT is usually 5 to 10 milliseconds, assuring that each station will have regular access to the ring.
The proper setting for the TRT is some matter for debate. Those who want to see lots of data on the ring and very little token time encourage a high TRT. Those who want to see very regular access and are less concerned about ring utilization push for low TRTs. The 5 milliseconds mentioned above should be viewed as a fairly low TRT, 10 milliseconds is moderate and anything above it is a high TRT. This number is usually configurable on a station by station basis. When new stations enter the ring, a new TRT will be determined, it will be the lowest TRT requested by all stations on the ring.
Jitter is less of a concern on FDDI rings than on Token-Rings. Each station on the FDDI ring has a buffer that is used to compensate for differences in clock rates - as opposed to Token-Ring where only one station responsible for managing jitter compensation. As a result, FDDI has maximum station count of 500 nodes. Each ring can be up to 62 miles in length and the distance between stations can be up to 2 km using multimode fiber. Single-mode fiber can be

used for distances of 12.4 miles, but get your wallet out as single-mode transceivers and fiber are extremely expensive.
All stations directly on the ring must be dual-attached stations. That is, four fibers will be used to build two distinct rings and each station must be attached to both rings. The secondary ring is normally not used for data transmission. It is there only to fix faults that may occur. Packets on the two rings flow in opposite directions and should there be a fault (broken cable or down station), stations adjacent to the fault wrap their transmit and receive lines from the two rings essentially forming one big ring.
This bigger ring may actually be up to 220 yards in length. The stations that have wrapped to form the new ring will constantly probe for a fix in the fault and will return the ring to its normal operation when the fault is no longer detected.
Single-attached stations including UTP-attached stations must go through concentrators to attach to the main ring. Concentrators are active devices that manage the insertion of station into the ring as well as provide for link integrity tests and other connection management functions. They also make the architecture flexible in that an end station with a dual-attach card can be "dual homed" to two different concentrators. If the primary connection should fail for any reason, the secondary connection can still be used to access the ring - a good fault-tolerant option.

ATM

We have described ATM briefly here as a cell-based technology. It is worth pointing out some of the differences between ATM and the technologies that we have talked about up to this point.
ATM is a point-to-point technology. There is no concept of sharing ATM's media. While this seems like a fairly odd choice, considering how expensive it can be to dedicate media and bandwidth to each and every station on the network, it is in fact a fairly logical choice. ATM was originally conceived as a wide area transport for use by telcos. In the telco world, the idea of many stations attaching to the same wire is as outdated as party lines.
This is not to say that two stations' data will never travel over the same wire, indeed this happens all the time, and must happen for the whole notion of a network to be reasonable. However, only two devices share each wire - one on each end. For that reason, the arbitration mechanisms that we've fairly carefully described for shared media networks are not appropriate for ATM. Rather, mechanisms need to be developed to arbitrate the bandwidth that will be available to two stations through the life of their data exchange.
The ATM specification for traffic flow control is called ABR or Available Bit Rate. It is a complex specification that must be implemented in silicon at the same level where packets are segmented into cells and cells are reassembled into packets. In the short term, this makes the economical promise of simple SAR chips (Segmentation and Re-assembly) a little hard to realize.
Another problem that faces ATM networks is their point to point nature. By definition, point-to-point, connection-oriented networks cannot support broadcasts and multicasts. However, we know that upper-layer protocols like TCP/IP and IPX require broadcasts to operate properly.
Finding a way to map the existing networking protocols onto ATM is a complex task. There are two approaches to attacking the problem. One is called LAN Emulation and it basically follows the same model as bridging. The other is called MPOA or MultiProtocol Over ATM.
LAN Emulation (LANE) provides mapping of six-byte LAN addresses into 20-byte ATM addresses as well as providing mechanisms for setting up virtual circuits between stations wishing to communicate, providing broadcast resolution mechanisms and handling for unknown packets. In this way, the ATM network looks like a bridge with various components exploded throughout the network. The devices that provide access from a shared LAN technology like Ethernet into the ATM network are called edge switches. Under LANE, the edge switches need to send broadcast packets to the device on the ATM network that can

handle them. However, once a virtual circuit is set up, the Edge switch sends to the intended end station without outside intervention.

The problem with LANE is that when stations need to communicate with other stations not on the same emulated LAN, they must go through a router. That router is a potential bottle neck. A better solution might be to have the ATM network emulate a router rather than a bridge. That is, provide mechanisms for resolving addresses at the network layer and making the edge switches smart enough to determine the route without a router.

That is essentially what MPOA does. It includes the mechanisms of LANE and adds a route server that builds routing tables and pushes them out to the edge switches. The edge switches then need only consult the table to determine the proper path to configure for the destination packet.

While this sounds simple, it isn't. Each protocol needs particular handling and may require more processing than simple routing. For example, if a packet originates on a Token-Ring node and is destined for an Ethernet node, the original data packet may be as large as 4,500 bytes - three times as is allowed by the Ethernet node. Each routed protocol handles problems like this in different ways and each must be accommodated.

Of course, the ideal way for a station to operate on an ATM network is to set up its own virtual circuits after consulting some central registry for an address (think of this as directory assistance or the White Pages). However, we have a couple decades of development and applications invested in our present applications and we can't just throw all that away - so LANE and MPOA are important to the success of ATM.

Topologies

A number of speeds have been suggested for ATM. Perhaps the first commonly implemented ATM system was the so-called ATM-TAXI system. TAXI is a chipset intended to implement FDDI's physical layer and therefore gave us 100-Mbps ATM. While this technology was instructional, it will not be commonly used in the end.

The telco industry has settled on 155 Mbps (also known as OC-3) as basic rate for ATM service. The next step up will be OC-12 or 622 Mbps. The step down from 155 Mbps is still a bit unclear. However, right now IBM's 25.6-Mbps ATM specification is winning favor as it uses many of the physical interface elements of 16-Mbps Token-Ring. In fact, the first switch port and end-station card combination to come to market with a price less than $1,000 was from IBM.

The problem for ATM25, as it is known, is that it may not represent enough of an advantage over switched Ethernet used in conjunction with LANE or MPOA. On the other hand, virtually any voice or video technology likely to run to personal computers during this century will likely work just fine over a 25-Mbps full duplex system. They may not all work so well over Ethernet. These three rates, therefore, are likely to be the ones to see significant volume over coming the months and probably years.

As might be guessed, ATM622 will require fiber. ATM155 will run over fiber or category 5 UTP cabling. ATM25 will work over category 3 or category 5 UTP wiring.