Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
FHWA/TX-05/0-4768-1		

4. Title and Subtitle	5. Report Date
FEASIBILITY OF ETHERNET AS A CENTER TO FIELD NETWORK FOR ITS FIELD DATA COMMUNICATIONS	October 2004
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Leonard G. Ruback, Edward Brackin, and David Rickerson	Report 0-4768-1

9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
Texas Transportation Institute The Texas A&M University System College Station, Texas 77843-3135	11. Contract or Grant No. Project 0-4768

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
Texas Department of Transportation Research and Technology Implementation Office P. O. Box 5080 Austin, Texas 78763-5080	Technical Report: September 2003 – August 2004
	14. Sponsoring Agency Code

15. Supplementary Notes

Project performed in cooperation with the Texas Department of Transportation and the Federal Highway Administration.

Project Title: Investigation of Ethernet Technologies as a Transport Mechanism for ITS Field Data

URL: http://tti.tamu.edu/documents/0-4768-1.pdf

16. Abstract

The objective of this research is to explore the benefits of an Intelligent Transportation System (ITS) field cabinet communication design utilizing local area network concepts and Ethernet technologies. The primary goal of the research is the definition of field cabinet network architectures and a concept of operation that can operate effectively given a relatively slow (narrowband) communication link or a broadband connection to a data collection center/point allowing the technology to be used in a variety of deployments. Vendors were encouraged to participate in this research by providing Ethernet-ready devices and network equipment for a technology demonstration. An experimental network was constructed and operated to demonstrate the capability of the architectures utilizing live data from field cabinets moved to a transportation management center. The researchers found that an Ethernet field cabinet and network architecture can supply the needs for today's ITS field devices and provide a significant amount of capability to support the devices of tomorrow.

17. Key Words	18. Distribution Statement
Ethernet, Intelligent Transportation Systems, Communications, Network	No restrictions. This document is available to the public through NTIS: National Technical Information Service Springfield, Virginia 22161 http://www.ntis.gov

19. Security Classif.(of this report)	20. Security Classif.(of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	97	

FEASIBILITY OF ETHERNET AS A CENTER TO FIELD NETWORK FOR ITS FIELD DATA COMMUNICATIONS

by

Leonard G. Ruback
Research Scientist
Texas Transportation Institute

Edward Brackin
Assistant Research Specialist
Texas Transportation Institute

and

David Rickerson
Student Technician
Texas Transportation Institute

Report 0-4768-1
Project 0-4768
Project Title: Investigation of Ethernet Technologies
as a Transport Mechanism for ITS Field Data

Performed in cooperation with the
Texas Department of Transportation
and the
Federal Highway Administration

October 2004

TEXAS TRANSPORTATION INSTITUTE
The Texas A&M University System
College Station, Texas 77843-3135

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official view or policies of the Federal Highway Administration (FHWA) or the Texas Department of Transportation (TxDOT). This report does not constitute a standard, specification, or regulation. The researcher in charge was Leonard Ruback.

None of the product selections identified in this report should be construed as product endorsements. Products were selected or provided based on their stated capability and should not necessarily be considered to reflect a "best of breed" choice.

ACKNOWLEDGMENTS

The authors would like to thank the Texas Department of Transportation (TxDOT) who sponsored the research and the following individuals who provided guidance and expertise during various phases of the research: Charlie Brindell of TxDOT who served as the project director; Al Kosik of TxDOT who served as the project coordinator; Steve Barnett, Nelson Wellspeak, Tony Parlamas all of TxDOT, and Robert Bacon, formerly of TxDOT, who participated as members of the Project Monitoring Committee. Special thanks to Robert Castelli of GarrettCom, Inc. and Deb Smith of Digi International for graciously providing demonstration equipment for this project.

TABLE OF CONTENTS

Page

LIST OF FIGURES

Page

LIST OF TABLES

Page

CHAPTER 1:
INTRODUCTION

Current TxDOT intelligent transportation systems (ITS) field cabinets typically use a combination of fiber optic, twisted pair copper for limited distance modems (LDMs), and leased telephone service for communications from the traffic management center (TMC) to field equipment. Baseband video is modulated onto fiber and transported to either the TMC or a satellite hub building. Fiber is also used to create a link from a field cabinet or a group of cabinets. Add/drop multiplexers can be daisy chained along a corridor and only the needed RS-232 channels are pulled off in each cabinet. The RS-232 channels are the link into field devices such as the local control unit (LCU) and a dynamic message sign (DMS). The LCU is used for monitoring traffic sensor (loop data), controlling lane control signals, and ramp meters. For areas where no TxDOT fiber (or partner fiber) exists, integrated services digital network (ISDN) telephone lines and regular public switch telephone network (PSTN) lines, also known as "plain old telephone system" or POTS, are used to provide connectivity to items such as encoded video.

In Texas, a popular ITS network architecture is based on a backbone service being supplied by an asynchronous transfer mode (ATM) network and narrowband service supplied by T1 (also known as a DS1 line) over fiber as shown in Figure 1.

This design is well founded to provide communications for legacy traffic sensors but can not capitalize on software services that are available on newer and emerging sensor devices, and it is limited to the T1 signaling speed of 1.544 megabits per second (Mbps). The T1 architecture can be relatively easily migrated over to a scheme using Ethernet to provide the legacy services but also enable broadband digital access in each ITS cabinet. This new feature enables a greater range of equipment to be considered for installation at the site. In addition, most TxDOT ITS devices (and ITS devices in general) do not need dedicated bandwidth. Video can be considered an exception to this statement although the recent advances in Ethernet quality of service (QoS) may solve this issue.

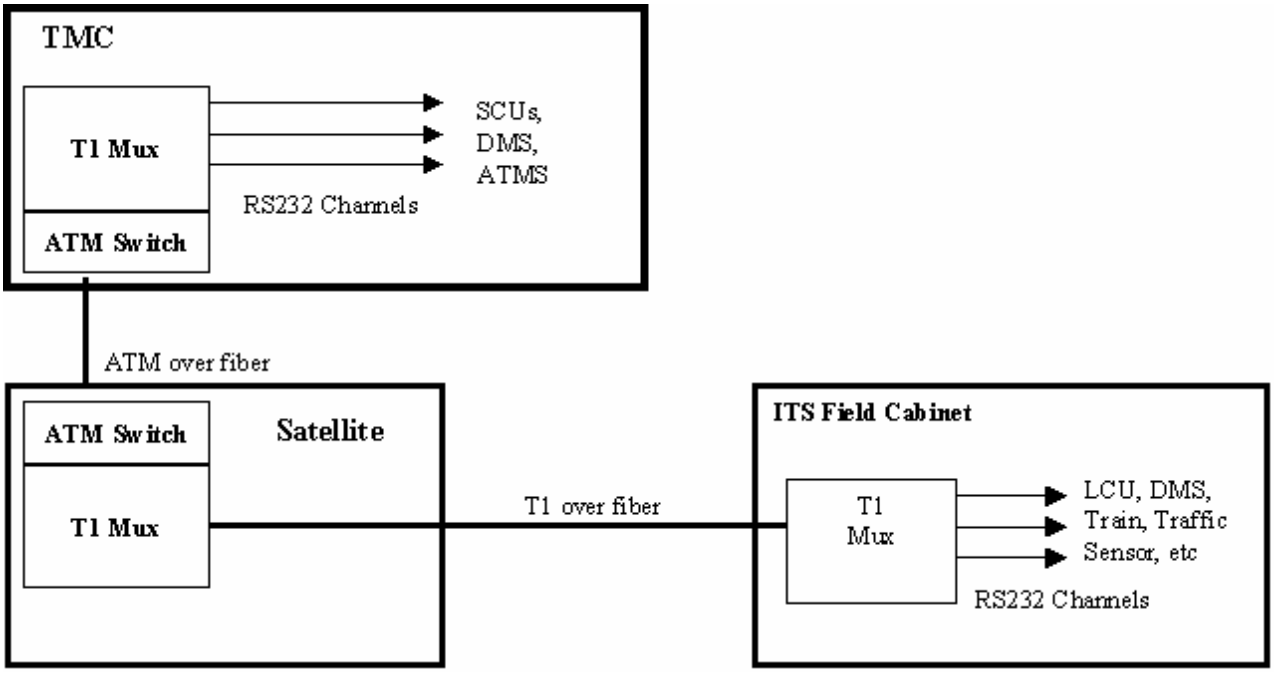

Figure 1. TxDOT ITS Communication Network Architecture Example.

It is important to note that little extra equipment is required to convert this architecture to Ethernet. The Ethernet edge switch may already be in place as part of the ATM deployment. The media converter is only required if the Ethernet edge switch is outfitted with twisted pair copper interfaces (100BaseT). An Ethernet switch could be selected that essentially has the media converter onboard. As for the field cabinets, hardened Ethernet manageable switches are available with a 100 megabit uplink port. The conversion to Ethernet will make available at minimum 100 Mbps service to field cabinets as opposed to the 1.544 megabit service from the T1 solution.

ETHERNET AS A REPLACEMENT

The attractions of using Ethernet in traffic and ITS applications are many. Ethernet enjoys widespread acceptance, is nonproprietary, and supports the universally popular transmission control protocol/Internet protocol (TCP/IP) suite of protocols. These Internet protocols support numerous utilities, such as telnet and simple network management protocol (SNMP), which can be used to remotely set up and/or configure equipment. From the user's point of view, the most obvious benefit is that Ethernet is based on an open standard. Since

Ethernet is an open standard, it guarantees more support on future technical advances compared with some proprietary control networks, increasing its flexibility. Because of the widespread acceptance of the technology, there are continual efforts to increase its capability and affordability. Bandwidth capacity has increased by a factor of 10,000 percent in recent years and is now approaching 10 gigabits per second (Gbps).

Ethernet has become the norm in large and small office networking. Large office systems prefer Ethernet technologies (switches, routers, hub, bridges, etc.) because of their ability to deliver exceptional performance for the dollar investment. The network must provide scalable, high reliability performance to maintain and increase the productivity of the company. On the performance side, Ethernet standards support communication speeds from 10 Mbps up to 10 Gbps. This speed easily rivals any other technology in the market today including ATM. The technology, at the same time, is also tremendously cost effective. Ethernet's use in small offices and even in homes is prime evidence of the overall value of an Ethernet solution. The benefit of large demand and high volume production is that common pieces of the Ethernet networking environment can be very competitively priced. An example of this would be the cost of the physical interface (10BaseT, 100BaseT) and a simple network "fabric" device such as a hub.

REPORT ORGANIZATION AND SCOPE

The purpose of this project is to investigate the ability of current field-ready (hardened) Ethernet technology equipment to perform the job of center to field data transport. The project will review Ethernet equipment that is designed to operate over traditional single mode fiber as well as methods of extending the Ethernet over leased or purchased copper and wireless. The effort will focus on the use of Ethernet as a replacement for TxDOT's traditional T1 over fiber field service and extensions beyond the fiber plant. Given these criteria, attention will not be directed toward a wide area network (WAN) architecture but rather a local area network (LAN) architecture. A chapter is devoted to network architectures that provide the required field services both on fiber as well as purchased media. Chapter 2 addresses the concept of operation for each of the architectures, and Chapter 3 presents a discussion of device integration. Chapter 4 is devoted to system management for an Ethernet network using standard SNMP protocols and tools. To fully investigate the capability of Ethernet, a demonstration network was constructed for testing and analysis and its construction and operation is documented in the final

chapter. A glossary of terms and appendices dedicated to additional information on various network topics are included.

CHAPTER 2:
ETHERNET CABINET ARCHITECTURES

To create a system architecture to replace the current T1 based communication system, several areas must be investigated and defined. A device needs assessment was conducted to determine the characteristics of the field equipment that will need to be supported by the new system. The devices need to be reviewed for physical interface requirements, bandwidth and throughput needs, and their latency sensitivity. An infrastructure review defines the physical media and service options (fiber optic plant, leased copper services, etc.) available for moving data from the field to the center. Finally, several architectures will be presented to utilize the different physical media available but still using standard, hardened industrial Ethernet equipment.

DEVICE NEEDS ASSESSMENT

A field equipment survey reveals that there is not a vast array of controllable devices currently deployed at the roadside. The following list identifies the most common devices that require communication at the roadside:

- camera pan, tilt, zoom (PTZ) controller;
- DMS controller;
- inductive loops and lane control signals (LCSs) via an LCU;
- traffic signal controller; and
- environmental sensor (weather) station.

Each of these devices typically provides an RS-232 or RS-422 physical interface and will require the native data to be encapsulated into Ethernet packets for transport over the communication infrastructure. Hardened terminal servers are available to perform this task and are available in both an RS-232 and RS-422 physical interface on the equipment side. All the sensor devices are relatively low bandwidth, typically 9600 bits per second, and can easily be supported by an Ethernet architecture system. Although they are low bandwidth, all but the dynamic message sign and weather station have some latency requirements.

Latency is a significant issue for some traffic management devices. Camera control, traffic signal controllers, and TxDOT LCUs are latency sensitive. Camera movement commands

must traverse the network with a minimal amount of delay in order for TMC operators to have a near instantaneous response for precise motion control. Latency is not an issue when moving to a preset position, but it becomes noticeable when trying to "freehand" control the camera with the only feedback being the received video from the camera. Excessive delay requires operators to anticipate when to issue commands resulting in inaccurate movement. TxDOT LCUs are deployed along the roadway to extract traffic speed, volume, and count data from inductive loops buried in the pavement, as well as control LCS units. Each LCU repeatedly exchanges data frames with a system control unit (SCU). Up to eight LCUs can be connected to one multi-drop communication channel on an SCU. The SCU orchestrates communication on the channel by issuing a request and waiting for a response from the addressed LCU. Each LCU on the channel receives three distinct requests for data transfers during a reporting interval. Data frames must traverse the network in a very timely manner to allow the SCU to manage all LCUs in the specified reporting interval, typically 20 seconds. Traffic signal controllers operate in a similar manner as the SCU/LCU. A more detailed discussion of the integration of LCUs is included later in this report.

Other sensor types and technologies can easily be visualized for the roadway of the future. The use of video detection systems is gaining more popularity with a push to make the raw video from the field camera available to the TMC or the responsible traffic engineering staff. There are numerous, relatively low cost devices in the marketplace to encode video and transport it over a TCP/IP network, including networks without tremendous bandwidth. Integration of simple video encoders on an Ethernet system is as simple as plugging in a cable and assigning the device a network address. Simple viewing can be done on a network connected computer (either in the field, the satellite, the TMC, or an external facility) instead of a traditional video display wall.

Roadway sensors currently gather more information and more accurate information than they can effectively convey to the LCU, which was designed to manage very basic contact closure sensors such as inductive loops. The information available on these new breed sensors will require their own serial communication link to a management application residing on the network, possibly in the TMC itself. Newer and upgraded sensors are now beginning to offer direct interfaces to Ethernet. Sensors are beginning to gather, manage, and share more information than legacy copper multi-drop networks can support.

In the future, all field sensors are expected to support common, standardized interfaces specified by the National Transportation Communications for ITS Protocol (NTCIP). The protocol suite has included Ethernet as a way to communicate with equipment (NTCIP 2104, Ethernet Subnetwork Profile). NTCIP utilizes many of the concepts of standard TCP/IP networking to manage devices, and an adoption of an Ethernet field cabinet will make the integration of NTCIP compliant devices even easier. NTCIP compliant devices are managed using SNMP, which is the same protocol used by the network infrastructure (switches, routers, etc.). Today, dynamic message signs, environmental sensor stations, and traffic signal controllers are available that support NTCIP.

INFRASTRUCTURE OPTIONS

There are multiple ways to link a TMC with field cabinets and equipment using Ethernet technology. The ultimate solution would be to construct and operate totally private networks. This solution is valid for large scale deployments where supporting fiber optic resources are available, and it is common in urban areas. The private network is logically and physically an extension of the TMC network. For areas without an in-place fiber physical plant, options still exist. Private broadband wireless products are making their way in the ITS marketplace and could be deployed to provide communications to isolated locations. Communications services can also be purchased from private sector providers. Vendors include:

- local cable television providers,
- wireline telephone providers,
- cellular wireless providers, and
- wireless Internet service providers (WISPs).

Purchased services can be organized into leased private and Internet access categories. Leased private circuits are available to provide point to point connectivity between facilities and use similar networking equipment to that found on the TMC network. Example services and bandwidths are:

- plain old telephone system at 56 kilobytes per second (kbps),
- ISDN at 128 kbps,
- fractional to full T1 (DS-1) at 64 kbps to 1.544 Mbps, and
- T3 (DS-3) at 44.736 Mbps.

Costs on the circuits cover the entire spectrum. Dial-up lines are low cost but are limited in bandwidth and may only be reasonable for simple sensors. Higher bandwidth comes at a cost with speeds over a megabit incurring hundreds of dollars per month per circuit. The bandwidth on these services is specified by the technology or contract and should not fluctuate. Being point to point, the services are not shared with other users; therefore, security is ensured by the provider. Copper services, POTS, and ISDN, may be the only options available in rural areas.

Internet access is now available through numerous entities. In urban areas, cable television and telephone companies can deliver broadband access to the Internet at the field cabinet. It may be possible to negotiate a service through them that does not touch the Internet, thus reducing security threats. The service provider would create a private network within his own infrastructure and include links to field sites and to the TMC. Typical Internet access is delivered via a cable or digital subscriber line (DSL) modem, and a router/switch will be required to effectively use and secure the service. Data rates vary but a common target is to provide close to T1 or fractional T1 level service.

Wide area wireless is becoming more prevalent in today's market, and it is expected to increase in the future. Cellular carriers are currently offering code division multiple access (CDMA) data services with a bandwidth of 56 kbps to 100 kbps. The major carriers have announced full third generation (3G) 500 K to 1 megabit service roll outs beginning this year with wide area coverage available in the next 2 years. Wireless Internet service providers are offering service within their coverage area at rates that are very competitive with similar leased copper services. Typical bandwidth is 1.5 megabit, which is offered to compete with T1 pricing. Internet access tends to be less expensive than private leased bandwidth, but it does not offer guaranteed, consistent bandwidth. Internet bandwidth from these providers is a shared resource and performance can be impacted by the number of subscribers in a local area. Promoted bandwidth levels are averages or peaks and thus may have difficulty supporting latency sensitive applications.

PRIVATE NETWORK ARCHITECTURE

Although this project does not address the communication network backbone, a brief discussion is pertinent to understand the entire network. Typically, a communication backbone is deployed to move very high bandwidth data among a TMC and remote stations called

satellites. A common architecture for a communication backbone is a ring with the TMC and satellites as nodes on the ring, as depicted in Figure 2. The technology used on the ring may utilize Synchronous Optical Networks (SONET), Asynchronous Transport Modes (ATM), or gigabit Ethernet over fiber optic cables. The satellites are geographically dispersed around the metropolitan area and act as bandwidth concentrators, where multiple lower bandwidth field links (for instance, T1 lines) are multiplexed onto the backbone. Communication links spread out from the satellite and touch numerous equipment cabinets in the region or neighborhood of the satellite. The field links act as data collectors from field cabinets that host roadside sensor data connections. For districts using ATM, it is common to deploy an Ethernet edge switch in the satellite. The edge switch provides a mechanism to bring standard Ethernet traffic onto the ATM backbone. With the edge switch in place, Ethernet is effectively extended from the satellite to the TMC over the backbone.

The link from the satellite to the field cabinets is the focus of this project. This link is normally configured in a daisy chained fashion among neighboring cabinets. In other words, the link moves from the satellite to the first field cabinet, then to the second, then to the third, and so forth threading its way through a group of cabinets. Each field cabinet may support multiple sensors but normally does not require very high bandwidth since the sensors themselves do not require it. It is advisable to create a ring out of the field cabinet "thread" if possible. The ring, when properly operated, provides redundancy (more than one path to a node) and, therefore, increases the overall reliability of the network.

Industrial Ethernet has the capability to supply the communication needs for the satellite to field component of the ITS data network. As previously pointed out, the satellite edge Ethernet switch provides the point of introduction for field Ethernet to hop onto the backbone. To begin an architecture study, a review of required network services that the network must provide is in order. Based on discussions with TxDOT project staff and within the guidelines of this research project, the Ethernet replacement system must provide at minimum the services that the T1 over fiber system does today.

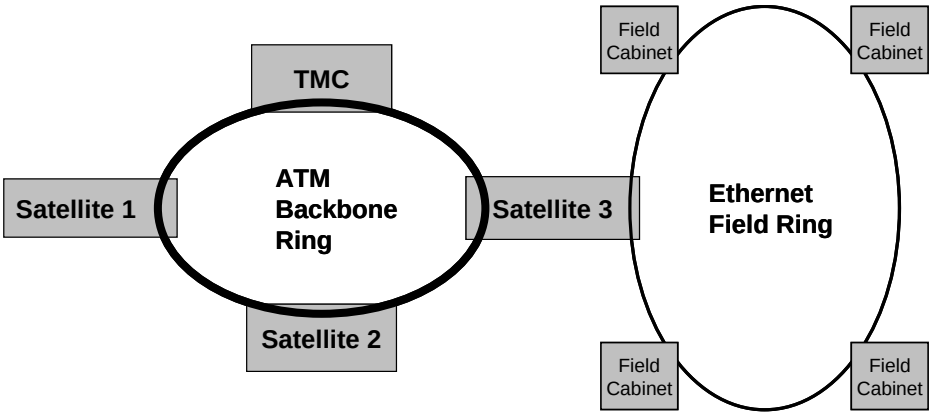

Figure 2. ITS Backbone and Field Network.

The T1 satellite to field system may contain 10 to 20 nodes, and it operates at a bandwidth of 1.544 Mbps. The point of origination of the T1 system will be the satellite or possibly the TMC itself. One plan may be to deploy nodes at regular intervals along a freeway and use limited distance modems or private wireless to move data from the installed sensor location to the field cabinet/node. The T1 system would work its way through all the field cabinet/nodes and may have an independent path back to the satellite if fiber resources are available. Return paths that lie in the same cable conduit are not considered independent. Multiplexers are installed in field cabinets and provide a network node. The multiplexer chassis is outfitted with modules or channel cards that break out low bandwidth RS-232 channels from the T1 stream. The channels are then assigned to individual sensors. The channels can be configured to operate in point to point or multi-drop fashion.

An Ethernet replacement solution for this type of network is very similar in physical architecture and is shown in Figure 3. A hardened, managed Ethernet switch would be installed in field cabinets replacing the T1 multiplexer. Industrial grade Ethernet switches are available from several manufacturers and operate at a speed of 100 Mbps. The Ethernet switch can be specified with a variable number of ports, typically multiples of four. Port modules are available that support either fiber or copper physical interfaces. Fiber interface ports support the link to and from other network nodes (cabinets). A stand-alone or integrated (into the switch) terminal server provides the same functionality as the channel card in the T1 system. The terminal server can support point to point and point to multi-point data transport. Additionally, the terminal server can also operate in a modem emulation fashion. The mode can be especially useful for

direct integration of systems that are currently operating with a dial-up modem connection where a network replaces the telephone line.

If the physical media (fiber) permits, the field network should be configured in a ring architecture. The Ethernet switch installed in the satellite hub provides the anchor point for both ends of the field cabinet fiber ring and will require two ports to support fiber connections. The design provides redundancy by creating a secondary path if the ring is broken in one place. Depending on the complexity and logical design of the backbone, a router may be deployed at the satellite hub to logically isolate the field cabinet network from the backbone network while providing an intelligent, manageable uplink to the backbone.

The architecture of a switch feeding providing multiple Ethernet ports for devices and a capability to uplink to another service is a basic building block in all the architecture models discussed. In the following cases, the only difference is in the equipment needed to interface to the wide area or long haul media. In this case, the switch itself supports a pair of fiber uplink ports to the broadband fiber. In subsequent cases, a secondary piece of equipment performs the task.

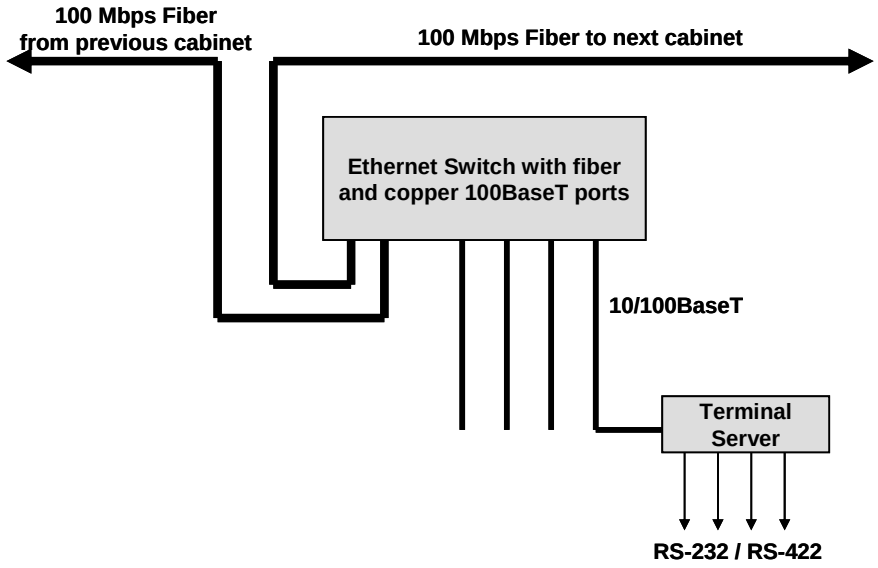

Figure 3. Private Network Cabinet Architecture.

LEASED SERVICE ARCHITECTURE

Leased services from private providers can be utilized to build a link to an isolated cabinet or node. One of the commonly used services is ISDN-Basic Rate Interface (ISDN-BRI).

ISDN-BRI provides data speeds up to 128,000 bits per second and is available in many places where other higher speed copper services are not. ISDN can be used to extend an Ethernet network to a field cabinet by using a combination of modems and network bridges or routers. Modems are used to create a data connection on the ISDN line and to convert data coming in on the modem's RS-232 interface to ISDN level signals. An asynchronous bridge or router is connected to one of the Ethernet switch ports and is used to convert packets from the normal 10/100BaseT interface to asynchronous RS-232 frames for transport over the ISDN channel. An identical setup on the far end converts the ISDN level signals to Ethernet frames.

A bridge is a simple device that spans two networks and transfers all packets available at one bridge to the other. It is very important to understand this concept and to understand the traffic that is expected to run over the link. If a significant amount of broadcast LAN traffic is present on the broadband side of the bridge, the transfer of the traffic onto the relatively low bandwidth ISDN connect will swamp the link. Response times from devices on the low bandwidth side will rise and packets may be lost in buffering. Proper switch port configuration can help alleviate the situation but not necessarily remove the issue. Routers can be used instead of bridges to limit the ISDN side from some undesirable traffic.

The ISDN connected cabinet architecture is very similar to that of the private network. A managed Ethernet switch is used to share the wide area connection with multiple cabinet devices. The Ethernet switch uplinks to an asynchronous bridge or router that is coupled with an ISDN modem. A terminal server is a likely cabinet device and installed to support RS-232 interface sensors. A simple video encoder and/or traffic sensors with a direct Ethernet interface may also be present. Bandwidth is limited to a maximum of 128 kbps and typical asynchronous modems connect at 115 kbps. Bandwidth management is critical in order to maintain acceptable performance from latency sensitive equipment. Figure 4 illustrates a simple leased line connected field cabinet using an ISDN line.

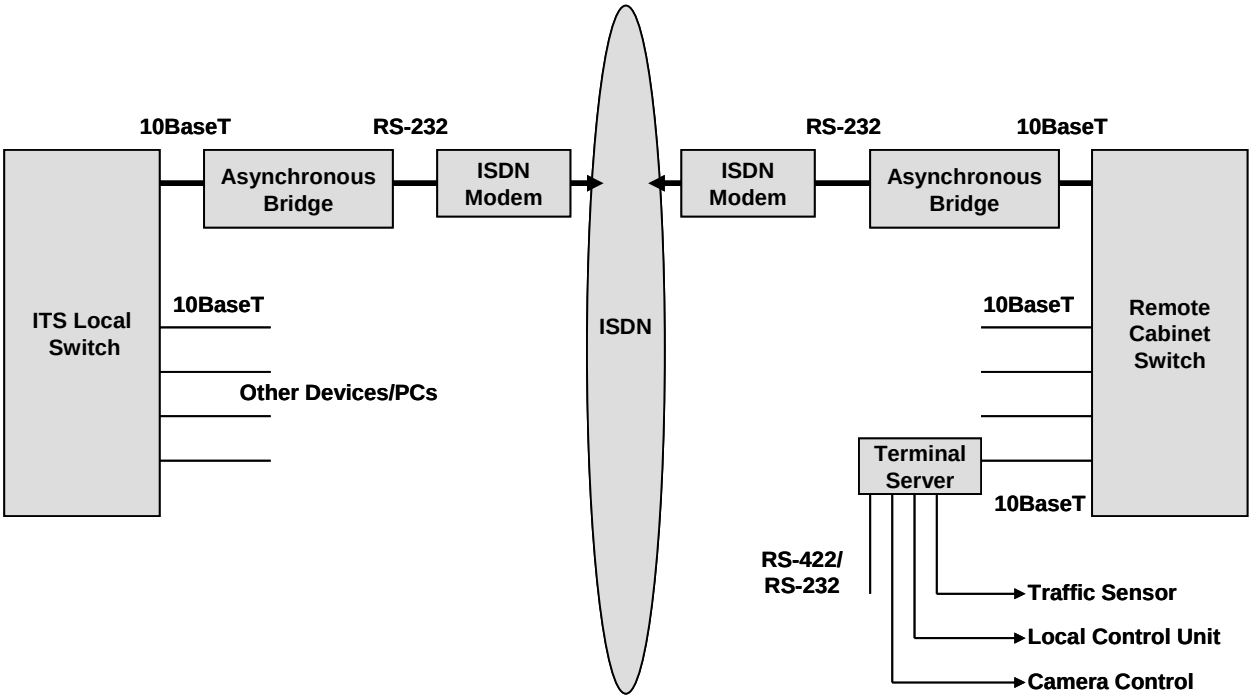

Figure 4. Purchased Services Cabinet Architecture.

PUBLIC INTERNET ARCHITECTURE

Access to the public Internet has become commonplace and available nearly anywhere. Internet access can be purchased from numerous Internet service providers (ISPs) and can be delivered over various media from simple dial-up modems to ISDN, DSL, cable, and wireless. Using the Internet to move traffic sensor data is possible in areas where other options are not attractive, but there are pitfalls. Using a public access network forces the topic of security to the forefront. Network security for systems that touch the Internet is typically handled through the use of a router, firewall, and a virtual private network (VPN). As before, a switch is used to share the protected network connection with devices in the cabinet.

The functions identified above can be contained in a single network appliance in the cabinet. The router function provides a method of segmenting the intra-cabinet network from the open Internet. The firewall enforces a set of rules for access to the intra-cabinet network from the outside world and can hide the workings of the inner cabinet network from the Internet. The

"wide open Internet" is restricted to a few specific, defined ports of entry, thus eliminating unwanted predator traffic from entering the intra-cabinet network through unexpected ports. A VPN provides a means to authenticate a user (field cabinet) and to encrypt data. Authentication keeps unintended users out of the system and ensures only the specific field cabinet is allowed to log onto the TMC data network. A VPN also encrypts all the data that moves over the Internet to the TMC, thus masking content as well as key information that hackers may use to attack either the field or TMC network. On the TMC side, a similar configuration is assembled. A router with embedded VPN can be used to allow trusted access onto the private network. A block diagram is shown in Figure 5.

The architecture for the devices that sit inside the protected network is the same as that for a private network. The VPN and router provide the security between the trusted network and the Internet. The intra-cabinet network can utilize the basic building block from the private network model. Planners can adopt a standard design for components that sit on the trusted (intra-cabinet) network and reuse it at other field locations. The router feeds an Ethernet switch that in turn provides multiple Ethernet ports for cabinet devices such as a terminal server or a video encoder.

As highlighted before, a method to reduce the security risk for purchased public services is to negotiate a private network arrangement with a provider of public services. Being that the bandwidth will never reach the Internet, the provider will not have the cost of his access to the Internet to offset. The data will flow totally within the provider's network. The arrangement may be particularly attractive in the WISP arena. Similar agreements and services have been provided by wireless cellular data vendors for large deployments of cellular digital packet data (CDPD) or CDMA services for public sector customers.

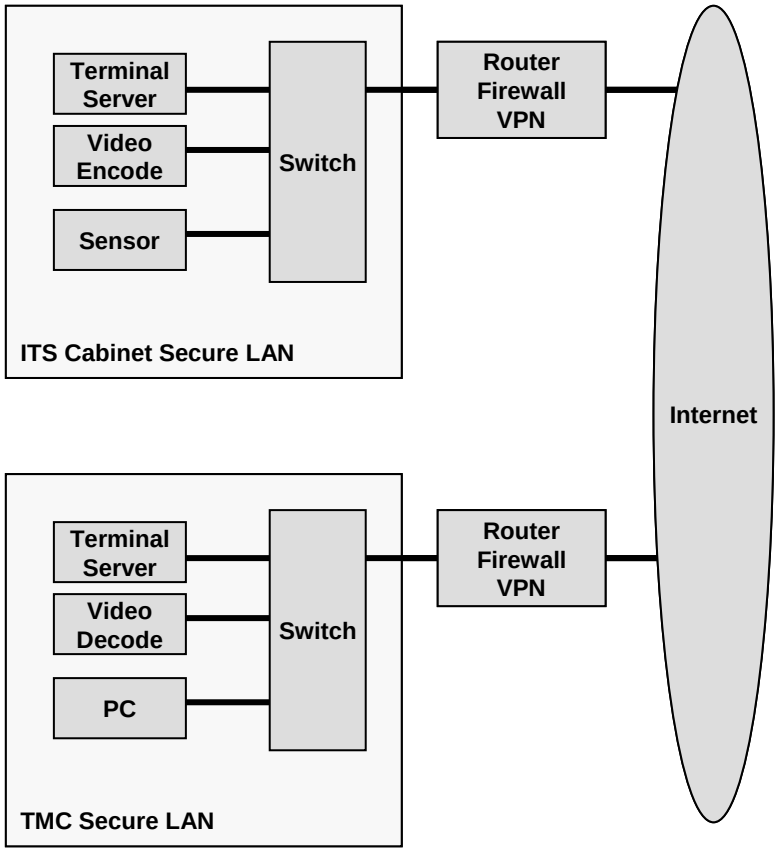

Figure 5. Internet Cabinet Architecture.

CHAPTER 3:
CONCEPT OF OPERATIONS

The concept of operations for each of the previously identified cabinet architectures is very similar. The intra-cabinet design in each architecture is that of a simple LAN. Multiple cabinets can be linked together through the cabinet Ethernet switch to form a larger LAN, and they can be logically managed as if they were located in one building. This concept is valid for the private network case and bridging over leased lines. For architectures that use a router to provide segmentation and security, the cabinet appears as a single isolated LAN.

The private network architecture simply creates a geographically extended local area network among the daisy chained switches in field cabinets. A router/gateway may be used to provide connectivity to the backbone, if the field sensor network does not terminate in the TMC. The router segments the field cabinet LAN from the backbone, therefore reducing the amount of housekeeping traffic introduced onto the backbone. Another approach for segmentation without the use of a dedicated router is to utilize virtual LAN (VLAN) technology. VLAN will be discussed in greater detail below.

Each device attached to the cabinet LAN should link directly to its associated switch. The one switch port to one device assignment will greatly improve the manageability of the network. Each device requires configuration prior to being added onto the networks, and the process is simple. Three basic network parameters are required to bring a device online: address, network mask, and gateway. The network mask and gateway are the same for all devices on the LAN, therefore only one "new" parameter is required per device. One of the most valuable assets of an Ethernet communication system is its ease and simplicity in setup. An Ethernet network requires no reconfiguration as new devices are added, thus effectively delivering a "configure it once" system. Typically, a system administrator configures two tiers of settings in a network switch.

Ethernet switches require little configuration to perform the typical network functions. Some very basic parameters such as device name, location, responsible party, and device IP address are configured in nearly all network devices, including switches. A second tier of setup requires configuring the spanning tree protocol (STP), VLAN, multicasting, and SNMP. Each of these configuration groups provides specific network capabilities.

If field cabinets are connected in a ring arrangement, the network switches must employ a method to automatically identify the redundancy and to hold connections in reserve. An Ethernet network is designed to logically operate as a bus (linear) network. A ring physical topology can be constructed, but one of the hops in the ring must be deactivated by the associated network switches. Spanning tree protocol is the means for this to happen and it must be available and enabled on the switches if a ring arrangement is constructed. STP manages all network links within the LAN and automatically recognizes a link failure. Upon failure detection, it calculates and activates new network paths to overcome the physical blockage.

Switches now support the concept of virtual LANs. A virtual LAN segments a large flat LAN into smaller broadcast domains. A broadcast domain is characterized by the group of devices that are logically located on a single LAN. Devices on a broadcast domain hear all traffic from devices on that domain. In general, VLANs logically segment a large LAN into smaller LANs. One potential use of VLAN technology is to deploy an ITS network as one single large flat LAN and segment certain similar resources into smaller VLANs. A similar use may be to allow multiple agencies to use the physical network and virtually organize the portions of the network into agency specific broadcast domains (VLANs). Routers are still used to move packets between VLANs. The technology can be very useful, but the technique's value must be weighed against the increase in network configuration complexity.

Multicasting is a technique to share a data stream from a single source with multiple end users. Video distribution to multiple viewers over a network is a common use of multicast. A video stream from a single encoder is replicated by a switch or router for any connected decoder that desires to receive the stream. In essence, the network fabric creates a video stream copy and sends to the attached decoder. If multicast video is to move over the network, the switches must be properly configured. Typically the administrator must enable Internet group management protocol (IGMP) on each switch. Multicast traffic can be very damaging to an improperly configured network. One mode of multicast failure is to send, by default, all multicast traffic to all devices on the LAN and prune back only when the devices fail to respond to a "continue stream" request. The potential flood of packets hitting devices with slower speed interfaces can be devastating. It is also important to note that multicast sources continually transmit their data out onto the LAN containing the source no matter whether there is another device requesting its stream or not. In addition, the router on the LAN must handle the multicast traffic from each

multicast source on its LAN. The router must be powerful enough to handle this traffic as well as the normal unicast traffic.

All equipment deployed on any of the networks identified should be manageable. Remote manageability is highly advisable for all network fabric devices but also should extend to end equipment (e.g., terminal servers, video encoders, power management) as much as possible. Managed devices make available valuable information to diagnose equipment problems. SNMP is the standard that is used throughout the networking industry to manage equipment. System management with SNMP will be covered in more detail in the following chapter.

Security on private and leased networks is maintained by limiting physical access to the network itself. Security on public access networks is provided by firewalls and VPNs. Network scanning software packages are available to scan the network and look for new devices that are not "authorized" and generate alerts. The software acts as an ever vigilant network watchdog on the network searching for intrusions.

The private network and the bridged network can be operated as a single "flat" LAN or broadcast domain. Since there will not be a considerable number of devices being moved on and off the network, simple static addressing can be instituted. Dynamic addressing is popular on networks with many transitory users but would not be necessary on an ITS sensor network. Internet connected networks are operated as isolated LANs with a router/firewall at the head end.

The RS-232 interface devices in the field may operate in a point to point fashion (sensor direct to management program) or in a point to multi-point fashion. Examples of the latter include LCUs, camera controls, and traffic signal controllers. The terminal servers support numerous methods of operation but will typically be used in one of the following:

- point to point – transmission control protocol server/client
- point to multi-point – user datagram protocol (UDP)
- modem emulation

Point to point applications may have a software program talk directly to the sensor, where the sensor has its RS-232 encapsulated in Ethernet packets. Many traffic equipment vendors are making their management software "IP aware." Other software designed to work via hardware RS-232 ports only can be upgraded to network connectivity by installing software that creates

virtual communications ports. The virtual port software works with the computer operating system, and the network enables traditionally hardware-only applications.

Modem emulation is a convenient way to migrate software and devices that previously operated on dial-up lines onto an IP network. In general, a pair of terminal servers replaces the dial-up modems with the IP network replacing the public switched telephone network. The device management software issues a modified dial string to its terminal server that closely resembles the string it would have issued to the dial-up modem. The server at the device end receives the call and issues responses similar to those of a dial-up modem. In this fashion, older management software packages can continue to operate and the telephone lines can be deactivated.

SCU/LCU INTEGRATION

Multi-point systems such as the SCU/LCU design can utilize the UDP protocol to emulate multi-drop RS-232 lines. The challenge is to configure the terminal server equipment to properly frame data as it arrives from the sensor. Common methods include watching for a defined end of frame character or searching for a defined minimum inter-character time to determine the end of a data frame. The LCU/SCU system requires the latter to be used since their protocol does not specify a unique character as an end of frame marker. Many terminal servers support this method, and it will be required for those used with an SCU/LCU arrangement.

To better understand operations, a simple discussion of SCU/LCU communications is necessary. The SCU polls up to eight LCUs on a channel or line. The SCU sends out a command to an addressed LCU and starts a wait timer. When the SCU hears the LCU begin responding, the SCU's wait timer is reset. If the timer expires before a response begins to arrive, the SCU declares the communication has failed. This system works well in a multi-drop, non-packetized environment where each host on the multi-drop hears each bit of the communication. An Ethernet system is a packet switched system. All data are organized into packets and switched around the network based on information encoded and added to the data payload of the packet. In a packet environment, the terminal server must wait to receive the entire frame from a device before it can transfer the frame over the network. The result is the SCU does not hear the LCU begin to respond, but rather the response arrives all at once. The impact is that the SCU

must wait longer for the LCU to respond than in a normal direct RS-232 installation. Wait times need to be increased to allow for the LCU frame to fully transfer into the terminal server and for the created Ethernet packet to traverse the network. There are other time-consuming tasks involved, but their required time is small compared to the two identified. The largest time delay is the time for a complete frame to input into the terminal server. Each character input into the terminal server consumes approximately 1 millisecond. An LCU response to a function code 18 request produces a 203 character frame and therefore approximately 203 milliseconds of time to input into the terminal server. The time to traverse a network is very small for an all fiber network, but it can grow to more than 100 milliseconds if public network services such as ISDN or CDMA are used.

To compensate for a packetized system, the SCU wait timer must be increased to allow for a longer LCU response time. Realistically, a timing study can be done by pinging terminal servers to reveal network travel time. Network delay time for systems connected using the public telephone system (POTS, ISDN) is generally held to less than 150 milliseconds for voice quality assurance reasons. Adding this time to the framing timing, a total of 450 milliseconds can be consumed by an LCU frame moving over a public switched network to reach the SCU. A minimum of 300 milliseconds appears appropriate for low delay networks, but times as much as a full second may be more appropriate.

To test the packetization and delay time theory, the project team assembled multiple LCUs and connected them to a single SCU. The LCUs were all connected to terminal servers that had a path to the SCU. The path length and type varied to help us better understand the associated delays. One LCU was located in an operational field cabinet in Austin and connected via an Internet link, and a pair of LCUs were located in a field cabinet in College Station and connected via a local Internet connection and a leased ISDN connection, respectively. The remaining five LCUs were located in the TransLink Laboratory on an experimental network. The SCU was also connected to a terminal server. In addition, a software application was created to monitor the packet communications between the SCU terminal server and the terminal servers of each of the eight connected LCUs (creating a full SCU channel load). The arrangement is shown in Figure 6. The SCU terminal server was configured to send a copy of all commands to all LCUs and the monitor. The LCUs were similarly configured. The software

application recovers each command sent between the SCU and LCUs and timestamps their arrival.

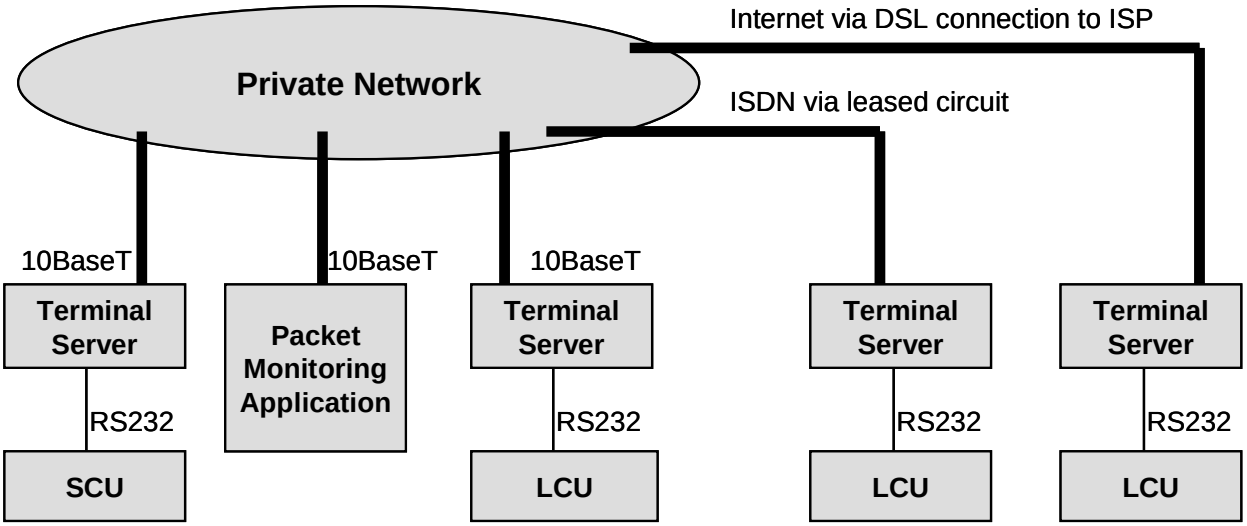

Figure 6. SCU/LCU Experimental Setup.

In addition, the program records the number of bytes in the transmission and uses it to determine if the packet contained a full data frame. Terminal server 'wait for end of frame' timing can be tuned to reliably capture a full frame per packet. A short wait time will create more than one packet per data frame, which is an undesirable condition. Table 1 represents a full 20 second cycle of communications between an SCU and eight connected LCUs. Each row represents a packet arriving at the monitor station. The device sending the packet is identified as either the SCU or one of eight LCUs based on the LCU thumbwheel setting. SCU messages contain a function code to inform the LCU how to decode the message. LCU responses do not contain a function code, but each is assumed to be the response from the last issued SCU command.

Table 1. SCU/LCU Message Delay.

Full 20 Second Cycle of SCU - LCU Channel Communications					
Device	Function Code	Bytes	Relative Time(msec)	Delta Time(msec)	Clock Time
SCU	17	6	0		3:04:17 PM
SCU	18	6	2000	2000	3:04:19 PM
LCU(1)	^^	203	2328	328	3:04:19 PM
SCU	18	6	2562	234	3:04:19 PM
LCU(3)	^^	203	2812	250	3:04:20 PM
SCU	18	6	3062	250	3:04:20 PM
LCU(4)	^^	203	3312	250	3:04:20 PM
SCU	18	6	3562	250	3:04:20 PM
LCU(5)	^^	203	3812	250	3:04:21 PM
SCU	18	6	4062	250	3:04:21 PM
LCU(6)	^^	203	4328	266	3:04:21 PM
SCU	18	6	4562	234	3:04:21 PM
LCU(7)	^^	203	4828	266	3:04:22 PM
SCU	18	6	5062	234	3:04:22 PM
LCU(8)	^^	203	5328	266	3:04:22 PM
SCU	18	6	5562	234	3:04:22 PM
LCU(2)	^^	203	5953	391	3:04:23 PM
SCU	6	6	6218	265	3:04:23 PM
LCU(1)	^^	47	6328	110	3:04:23 PM
SCU	6	6	6390	62	3:04:23 PM
LCU(3)	^^	47	6562	172	3:04:23 PM
SCU	6	6	6640	78	3:04:24 PM
LCU(4)	^^	47	6765	125	3:04:24 PM
SCU	6	6	6843	78	3:04:24 PM
LCU(5)	^^	47	6968	125	3:04:24 PM
SCU	6	6	7047	79	3:04:24 PM
LCU(6)	^^	47	7172	125	3:04:24 PM
SCU	6	6	7234	62	3:04:24 PM
LCU(7)	^^	47	7359	125	3:04:24 PM
SCU	6	6	7437	78	3:04:24 PM
LCU(8)	^^	47	7562	125	3:04:24 PM
SCU	6	6	7640	78	3:04:25 PM
LCU(2)	^^	47	7843	203	3:04:25 PM
SCU	6	6	10000	2157	3:04:27 PM
LCU(1)	^^	47	10203	203	3:04:27 PM
SCU	6	6	10281	78	3:04:27 PM
LCU(3)	^^	47	10422	141	3:04:27 PM
SCU	6	6	10484	62	3:04:27 PM
LCU(4)	^^	47	10625	141	3:04:28 PM
SCU	6	6	10687	62	3:04:28 PM
LCU(5)	^^	47	10828	141	3:04:28 PM
SCU	6	6	10890	62	3:04:28 PM
LCU(6)	^^	47	11015	125	3:04:28 PM
SCU	6	6	11093	78	3:04:28 PM
LCU(7)	^^	47	11218	125	3:04:28 PM
SCU	6	6	11297	79	3:04:28 PM
LCU(8)	^^	47	11422	125	3:04:28 PM
SCU	6	6	11500	78	3:04:28 PM
LCU(2)	^^	47	11672	172	3:04:29 PM
SCU	17	6	20000	8328	3:04:37 PM
SCU	18	6	22000	2000	3:04:39 PM

"^^" indicates an LCU response with no listed Function Code.

In this example the SCU was located in the TransLink Laboratory on a network with access to a private Ethernet network as well as a link to the Internet. LCU(1) was located at the Highway 6 Testbed and linked in via a DSL connection to the Internet. LCU(2) was located at

the same site but linked in via a leased ISDN circuit. LCU(3) through LCU(7) are virtual LCUs running on a laptop and located on the private network in the TransLink Laboratory. LCU(8) was a virtual LCU located in Austin and linked in via a DSL connection to the Internet. The tested arrangement utilized LCUs on each of the three example field cabinet architectures: private, leased, and public.

The fifth column, delta time, represents the amount of time between packets arriving at the monitor application. Another way of viewing the column is to consider the delta time to be the time delay between an SCU message and the response from the addressed LCU. The LCU communicated at 9600 baud over its RS-232 port. At 9600 baud, 203 bytes consumes approximately 211 milliseconds of time to transfer from the LCU into the terminal server. The terminal server waits for a certain amount of silence, in this case 40 milliseconds, to trigger the packetization of the inbound data frame. Approximately 250 milliseconds of time is required by the terminal server to receive the LCU data frame and pass it out as an Ethernet packet, assuming the time to create the packet is negligible in comparison to the RS-232 transfer time. Network delay time must also be factored in. The time to travel through the private, fiber network is less than 1 millisecond. The time for a round trip over the network segment created by the ISDN circuit is approximately 115 milliseconds. Under further inspection, approximately 100 milliseconds of delay were incurred by the ISDN circuit and approximately 15 milliseconds were incurred by the asynchronous bridges processing time. Travel time over the Internet connection to Austin was approximately 47 milliseconds.

SCU dwell time delays must take into account all the timings discussed. If the LCUs are all connected with a private fiber network, the travel time is very small compared to the terminal server time. On network segments running over leased or public services, the delay time should be tested and used in SCU timing calculations. In the test case, the slow link was the network segment provided by the ISDN circuit. The delay time from the SCU to LCU response was measured at 391 milliseconds for a 203 byte response message. The Internet connected LCU times were smaller and tested at 330 milliseconds from Austin. The SCU dwell time must be set to accommodate the largest byte count message for the LCU on the slowest link.

CHAPTER 4:
SYSTEM MANAGEMENT

As identified in previous sections, the advantages of Ethernet are numerous. These advantages include being a worldwide open standard, as well as possessing the flexibility, scalability, and reliability to handle installations as diverse as a simple home network to the largest multi-national corporations in the world. These advantages have resulted in Ethernet having the largest installed base of any networking system in the world and the lowest cost for providing network capability.

Apart from the advantages, implementation of an Ethernet communications system requires an understanding of how to manage the system to achieve and maintain optimum communication flows, which translates to the data getting to where it needs to be, when it needs to be there. While this is true for any type of communication system, there are several aspects of Ethernet that call for active management solutions.

One important aspect of Ethernet to understand is that it is a connectionless oriented system. Ethernet does not assign bandwidth to a particular device; rather, it provides bandwidth as a backbone capability that all devices share. This can lead to situations where one or more devices could consume a large amount of the available bandwidth. While this typically only happens as a result of a malfunctioning device or a network configuration error, the consequences will be immediately felt when data reception is compromised.

Another important aspect of Ethernet is that because it is an open standard, with more than 30 years of devices and installations, it is well known and understood. That means that most network devices and equipment support Ethernet and can be plugged into an existing network with a minimum of configuration. This ease of functionality represents a double-edged sword. While network administrators can add to their network quickly and easily, this capability is also an open door with respect to security that must be firmly closed. Overall, these caveats are not a cause for significant alarm, but rather a call for sound management strategies. Sound system management strategies should be used to monitor any deployment, to ensure that it meets the requirements for reliable, consistent, affordable data communications.

System management generally consists of both performance monitoring and security monitoring. While these activities are related and, in fact, utilize some of the same information, it is appropriate to consider each aspect of system management by itself, in order to thoroughly

understand both the impacts and solutions available to accomplish these activities. Each of these topics will be addressed in the following sections, followed by a discussion of how monitoring capabilities are typically implemented in today's modern Ethernet-based networks.

PERFORMANCE MONITORING

System administrators have been monitoring networks since they were first invented and deployed. While the initial monitoring capabilities were crude and focused solely on whether or not a device was operational, modern, sophisticated monitoring techniques can examine several aspects of a network simultaneously.

At the highest level, performance monitoring focuses on the overall state of devices in the network. Some typical questions that are examined via monitoring techniques would include:

- What devices are present in the network?
- Is each device responding to a health check?
- Is each device responding in a timely fashion?
- Is each device sending and receiving data?

At the next level, information about the components of a particular device can be obtained. For example, a network switch placed in a traffic signal cabinet must endure very high temperatures. A typical piece of information that the system administrator can monitor is the internal temperature of the device to ensure that it is not operating beyond its threshold. Other items that could be monitored on some devices may include the status of various components, such as cooling fans or power supplies. Together, these items provide a wealth of information pertaining to the physical health of the network. Ensuring the optimum health and consistent operation of network devices is a critical part of maintaining a reliable and consistent communications.

Beyond the basics of device health, the concept of performance monitoring is used to determine how well the network is functioning. While a device may be alive in the network and responding to a basic query technique, such as a ping, there may be other issues affecting the device and its communications that will not show up in health investigations. At this level, performance monitoring seeks to identify issues such as:

- high latencies (lag time) in data stream communications
- high rates of packet (information) loss

- levels of bandwidth being consumed by devices
- memory usage and CPU load on devices
- input, output or buffer errors on devices
- errors that devices can log to a central server.

By identifying these and other types of problems early on, troubleshooting can pinpoint the cause and resolution so that the situation does not cascade across the entire network.

As a practical example of using performance monitoring, consider a situation where a video encoder was added into a network and set to operate in a multicast environment. However, a device configuration error has all of the packets being sent out with a time-to-live (TTL) value of 1. The end result of that error is that these packets will be sent to every port on every device across the network, a situation commonly referred to as flooding. Performance monitoring applications can track not only the changes in the physical network, by the addition of devices, but also the bandwidth that is being sent to and from every device. Having the capability of determining and using this information can help system administrators find and fix the problem.

SECURITY MONITORING

With the ever increasing threats of virus attacks, hackers, denial of service attacks, and more, network security concerns have never been more in the forefront of communications than they are today. Indeed, there is an entire section of the networking industry that is focused exclusively on security. Solutions exist in the form of both software and hardware, each with their own set of pros and cons.

Even though sophisticated and expensive solutions for supplementing the security of networks exist, a great deal of security can be accomplished by using common sense and some well-established guidelines for network security. These practices include items such as:

- Establish consistent and non-standard passwords.
- Do not support network protocols that are not in use.
- Monitor all network devices for intrusions.
- Establish a change management system.

While these may sound like complicated items to accomplish, the reality is that system administrators can quickly and easily accomplish all of these items and more. The statement pertaining to the support of protocols simply means that a system administrator should turn off

those items that will not be used. If all devices will be managed via a standard telnet interface, then turn off the web interface that uses the hypertext transfer protocol (HTTP), as this is one of the most prevalent forms of attacks currently known. Likewise, a change management system, while sounding complicated, may be nothing more than a notebook containing logs of when a device configuration was changed, by whom, and for what reason. Having the ability to backtrack why a change was made can quickly and easily answer most questions that are raised pertaining to device configurations.

One thing that must be understood about security is that devices vary widely in how much security they support, even across the same vendor. A typical managed switch, for example, will contain a large number of considerations for addressing security, whereas a wireless device will likely contain a far smaller set of security considerations. Wireless components bring a whole other set of security considerations to the forefront, as the access is now 'in the air' and requires no special access. In fact, most wireless deployments that utilize only the minimum security settings can be hacked within 4 to 8 hours.

While it is true that a complete set of security considerations can only be developed given the knowledge of a particular deployment, many standard practices can be identified and provide a solid impetus to establishing network security as ongoing consideration in deployment. Table 2 details the best practices for establishing minimum security in deployments.

SNMP Applications

Simple network management protocol is a set of standards for communication with devices connected to a TCP/IP network. In particular, SNMP provides administrators with a set of operations that allows them to remotely monitor and manage devices. As an example, not only can SNMP check the bandwidth in use on any port on a switch, it could also be used to change the configuration of the switch to shut down a port, if the bandwidth needs are excessive. It could also be used to determine the media access control (MAC) address of all devices on that switch and what protocols are being sent through each and every port on the switch, and produce a continuous graph of all of this information. In short, SNMP provides a very powerful and standards-based method for accomplishing not only performance monitoring, but most aspects of security monitoring as well.

Table 2. Best Security Practices.

1. Establish a user management system that identifies who requires access to network devices and at what level (administrator or user).
2. Establish a change management system for all device configurations.
3. Establish a network time server so that all devices and events are operating on the same clock, to help in tracking events.
4. Establish a network logging server to receive and forward critical device alerts.
5. Establish routine network security checks and polling intervals to aid in catching any network changes.
6. Apart from physical security, set up an application level user identification and password for authentication for all devices.
7. Remove standard application level user-IDs such as "Guest" or "Anonymous."
8. Reduce hacking capabilities by using passwords 6-8 characters in length with non-standard characters.
9. Remove protocols that are not necessary for network operations.
10. Protocols to consider for their need / security aspects include: a. Center to Center (C2C) b. Hypertext Transfer Protocol c. File Transfer Protocol (FTP) d. Simple Mail Transfer Protocol (SMTP) e. Terminal Emulation (Telnet) f. Simple Network Management Protocol g. Network Time Protocol (NTP) h. Trivial File Transfer Protocol (TFTP)
11. Where SNMP is used, utilize non-standard community strings for Read and Read/Write operations.
12. When possible, use private network addressing schemes to minimize entry points into the network infrastructure.
13. When using wireless devices, at minimum: a. Disable the service set identifier (SSID) broadcast. b. Enable MAC address filtering. c. Enable wired equivalent privacy (WEP).
14. Establish a security response to any threats. This response should at minimum: a. Isolate the device. b. Restrict normal user access. c. Decrease polling intervals. d. Compare current configurations to last known good configuration from change management system.

Because SNMP is a well-known and standardized protocol, numerous applications exist to accomplish the diverse set of needs within the arena of performance and security monitoring. While the discussion of individual applications or packages is beyond the scope of this report, feature-rich applications can be obtained for as little as several hundred to one thousand dollars. Accomplishing routine monitoring does not have to be an extensive or expensive addition to any deployment.

Figure 7 shows a sample from an SNMP package that is monitoring device status. Each device is tested on a 2 minute cycle to ensure proper response. If a response is received, the poll simply occurs at the next cycle. Polling can be set at any level desired by the administrator. Obviously, in a situation where bandwidth is at a premium for data transfer, the polling intervals would be increased (less frequent) to minimize data disruption.

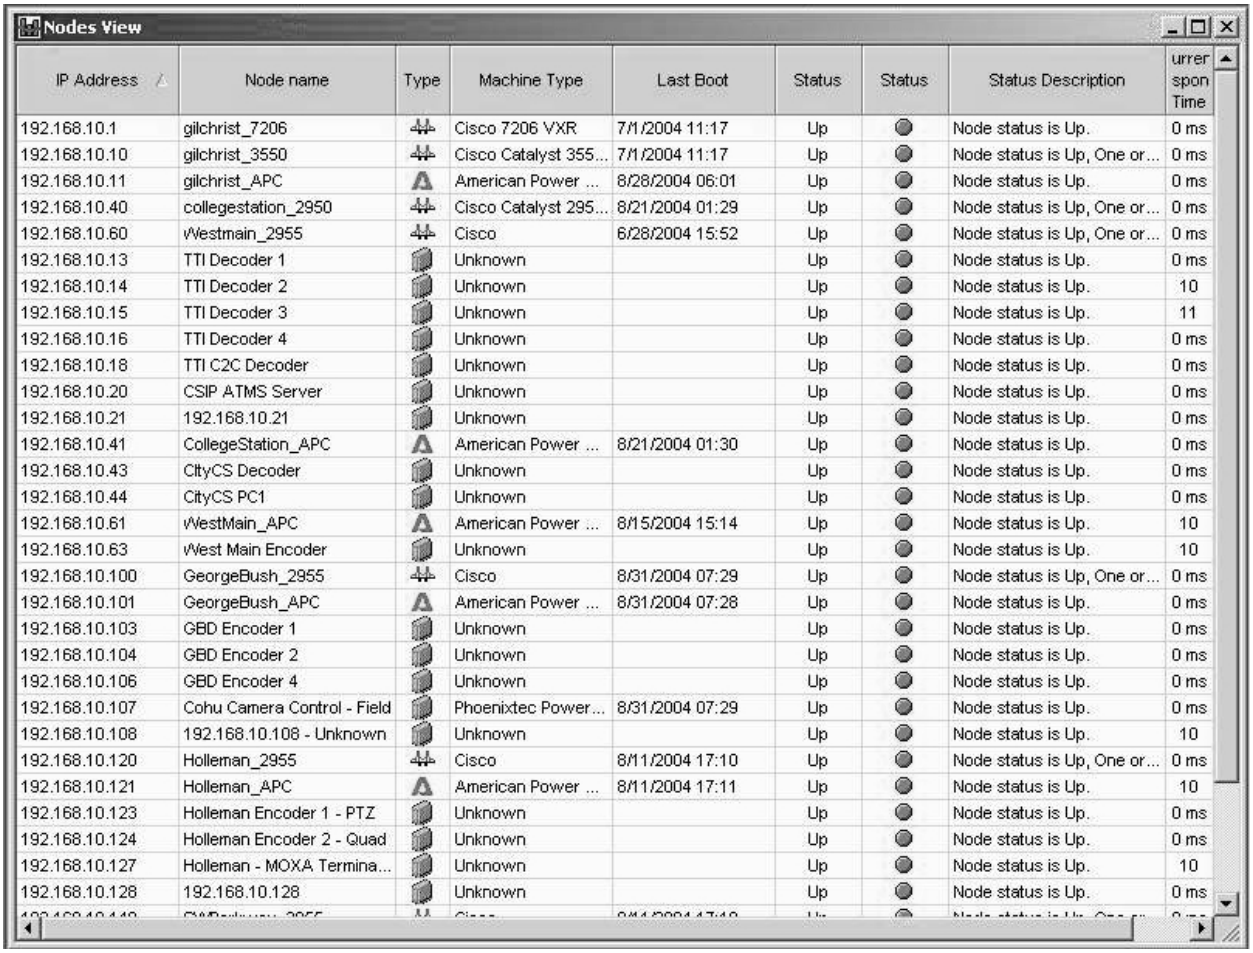

IP Address	Node name	Type	Machine Type	Last Boot	Status	Status	Status Description	urrer spon Time
192.168.10.1	gilchrist_7206		Cisco 7206 VXR	7/1/2004 11:17	Up		Node status is Up.	0 ms
192.168.10.10	gilchrist_3550		Cisco Catalyst 355...	7/1/2004 11:17	Up		Node status is Up, One or...	0 ms
192.168.10.11	gilchrist_APC		American Power ...	8/28/2004 06:01	Up		Node status is Up.	0 ms
192.168.10.40	collegestation_2950		Cisco Catalyst 295...	8/21/2004 01:29	Up		Node status is Up, One or...	0 ms
192.168.10.60	Westmain_2955		Cisco	6/28/2004 15:52	Up		Node status is Up, One or...	0 ms
192.168.10.13	TTI Decoder 1		Unknown		Up		Node status is Up.	0 ms
192.168.10.14	TTI Decoder 2		Unknown		Up		Node status is Up.	10
192.168.10.15	TTI Decoder 3		Unknown		Up		Node status is Up.	11
192.168.10.16	TTI Decoder 4		Unknown		Up		Node status is Up.	0 ms
192.168.10.18	TTI C2C Decoder		Unknown		Up		Node status is Up.	0 ms
192.168.10.20	CSIP ATMS Server		Unknown		Up		Node status is Up.	0 ms
192.168.10.21	192.168.10.21		Unknown		Up		Node status is Up.	0 ms
192.168.10.41	CollegeStation_APC		American Power ...	8/21/2004 01:30	Up		Node status is Up.	0 ms
192.168.10.43	CityCS Decoder		Unknown		Up		Node status is Up.	0 ms
192.168.10.44	CityCS PC1		Unknown		Up		Node status is Up.	0 ms
192.168.10.61	WestMain_APC		American Power ...	8/15/2004 15:14	Up		Node status is Up.	10
192.168.10.63	West Main Encoder		Unknown		Up		Node status is Up.	10
192.168.10.100	GeorgeBush_2955		Cisco	8/31/2004 07:29	Up		Node status is Up, One or...	0 ms
192.168.10.101	GeorgeBush_APC		American Power ...	8/31/2004 07:28	Up		Node status is Up.	0 ms
192.168.10.103	GBD Encoder 1		Unknown		Up		Node status is Up.	0 ms
192.168.10.104	GBD Encoder 2		Unknown		Up		Node status is Up.	0 ms
192.168.10.106	GBD Encoder 4		Unknown		Up		Node status is Up.	0 ms
192.168.10.107	Cohu Camera Control - Field		Phoenixtec Power...	8/31/2004 07:29	Up		Node status is Up.	0 ms
192.168.10.108	192.168.10.108 - Unknown		Unknown		Up		Node status is Up.	10
192.168.10.120	Holleman_2955		Cisco	8/11/2004 17:10	Up		Node status is Up, One or...	0 ms
192.168.10.121	Holleman_APC		American Power ...	8/11/2004 17:11	Up		Node status is Up.	10
192.168.10.123	Holleman Encoder 1 - PTZ		Unknown		Up		Node status is Up.	0 ms
192.168.10.124	Holleman Encoder 2 - Quad		Unknown		Up		Node status is Up.	0 ms
192.168.10.127	Holleman - MOXA Termina...		Unknown		Up		Node status is Up.	10
192.168.10.128	192.168.10.128		Unknown		Up		Node status is Up.	0 ms

Figure 7. Example of Routine SNMP Polling of Network Devices.

In most SNMP applications, if a change is detected, the application would go into a rapid polling situation, where the timeframe between communications to that particular device is shortened in order to capture any changes. After a defined number of failed attempts, the application would send an alert to the system administrator identifying the device failure. The power of this type of monitoring is two-fold. First, after the initial setup, all of this is completely automatic and can take place without administrator intervention on a 24-hour a day, 7-day a week timeframe.

The second powerful aspect of monitoring applications is that virtually any facet of device communications can be monitored via SNMP. Many of these were mentioned earlier, such as bandwidth, latency, packet loss, and temperature. Whatever the device manufacturer exposes to the SNMP protocol, a standard SNMP application can find and track the information and act on it. Similar to the monitoring shown in Figure 7, these communications are completely automatic and operate on a continuous basis.

One of the most common ways in which an SNMP application would act on a change in expected conditions would be to provide an alert to the system administrator. Here again, the power of standards based applications is seen, with feature-rich methods of getting information from an application to an administrator. Figure 8 shows a listing of the typical methods in which any abnormal conditions can be communicated, including e-mail, page, logs (of several types), execution of another program, sounds, and more. Most SNMP applications support a wide variety of notification routines, even supporting the Short Message System (SMS) common in most cell phones today.

As mentioned previously, another aspect of monitoring is examining the security of a network. Figure 9 shows an application where a port scanner is being used to test the commonly targeted ports on a network. The application reports if the port is open, closed, or if there is no reply. Scans like this can often be automated on a defined user interval and then e-mailed or exported for storage. Periodic comparison against previous scans can show any differences in device configurations and alert administrators to any attempts to compromise network devices.

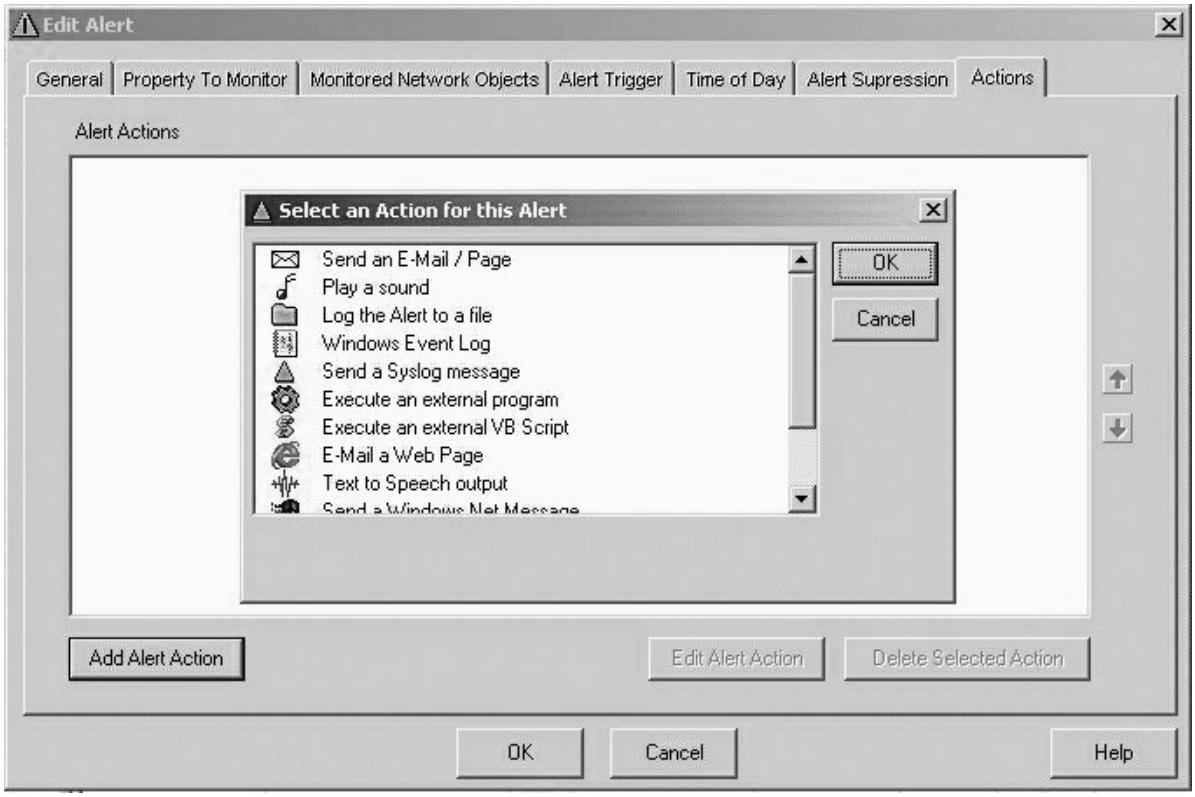

Figure 8. Example of SNMP Notification to System Administrators.

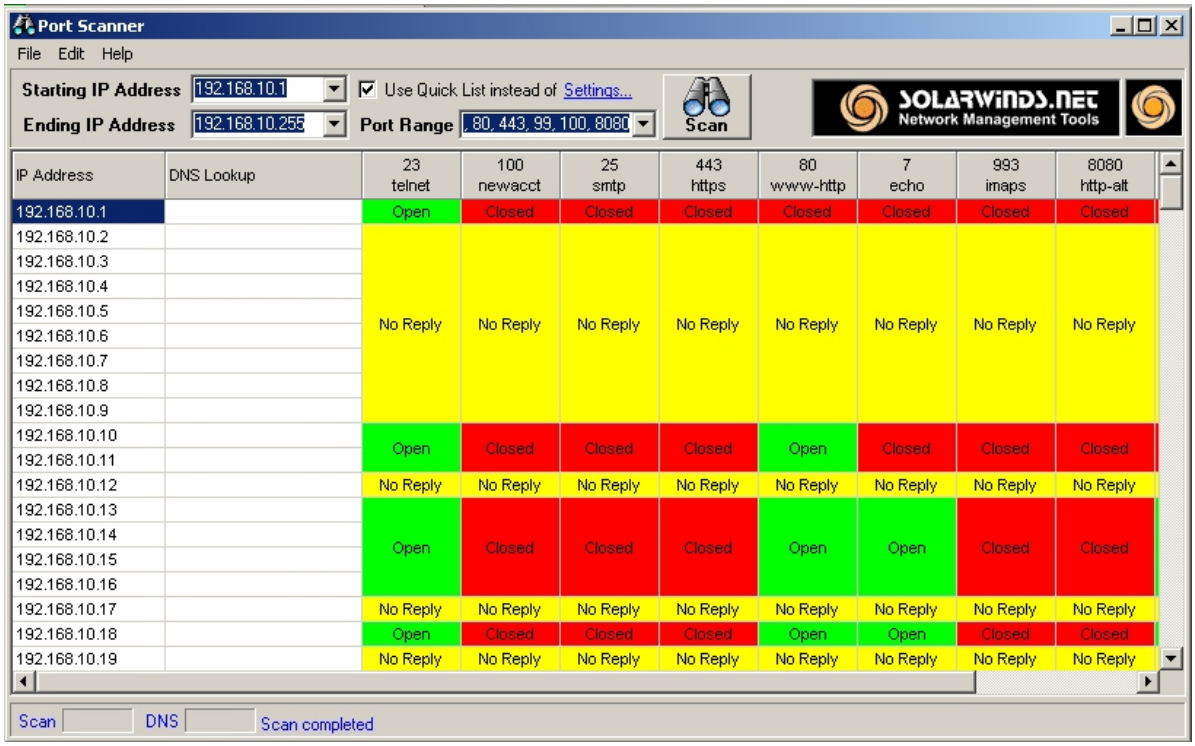

Figure 9. Network Port Scanner.

Summary

Monitoring of any network implementation is not only desirable, it is necessary to maintain the reliability, stability, and security of the network. Most monitoring is either performance or security based and relies on the SNMP support in individual devices. The use of SNMP also provides a standard methodology for accomplishing monitoring activities, and it is widely implemented in cost-effective applications from multiple vendors. It should be noted that there are numerous free SNMP applications available, in addition to the commercial products. The choice of which tool would perform best is based on need and is left to the reader. Regardless of the tool, the use of standard monitoring techniques, along with an understanding of best practices for security and common sense, means that most networks can operate efficiently and reliably, and be reasonably secure with a minimum of effort. This is the true power of SNMP and standards-based implementations.

CHAPTER 5:
DEMONSTRATION NETWORK

To further investigate the concept of Ethernet as a center to field communication option, an experimental Ethernet network was constructed in the TransLink Laboratory and also included external links to field cabinets. The network was specifically designed to include all three architectures as previously discussed. A small private fiber ring network was assembled in traffic signal cabinets residing in the TransLink Laboratory with hardened field-ready switches. Links to operational field cabinets over ISDN and the public Internet extended off this private network ring. The design allowed live traffic sensor data to be used for demonstration and experimental purposes.

EQUIPMENT SELECTION

The project budget allowed for the selection and purchase of equipment to support an ISDN network link and to provide terminal servers for traffic sensor integration. Numerous network product and sensor vendors were contacted and asked to participate by providing demonstration equipment for the network on a loan agreement. Several responded and generously provided demonstration equipment for the network build out. Other segments of the total network were already in place. The segments include Internet links to field cabinets and a broadband gigabit backbone and field cabinet network, as well as a TMC office network. Many different products were used in creating the total network.

The purpose of the project was to demonstrate that an Ethernet-based system, with all the required components, is available today to fulfill the requirements of an ITS field data network. Product selection was based mainly on capability and cost. The equipment chosen must provide the functionality required yet not necessarily be the absolute best choice of available products or techniques. None of the products selected were extensively tested nor were comparisons made among similar products. Product selection should not be construed as a product endorsement.

An exhaustive search for low cost ISDN network equipment revealed that there are few ISDN modems in the marketplace. ISDN routers and router blades are available for enterprise class network devices, but the cost and temperature specifications made them less attractive for the project. The asynchronous bridging technique was chosen based on cost, capability, and ease

of integration. The bridge can be used with any media that can transport RS-232 level signals. Media choices can be wireline as well as wireless and need only a simple RS-232 cable to connect the bridge to the RS-232 device. The bridges can be managed or unmanaged and were very simple to configure. An Adtran model 3000 ISDN modem pair was selected mainly for their ability to automatically bind both "B" channels of the ISDN line and to auto redial upon loss of connection. The modems delivered 115,000 bits per second continuous RS-232 connection for the bridges to operate on. The available network speed enables multiple simple sensors such as an LCU and traffic detectors to be serviced by a single connection. The connected devices bandwidth needs require close consideration in order to maintain a reasonably quick response time. Bandwidth 'breathing room' needs to be maintained on the link to insure acceptable operation.

Terminal servers were required to encapsulate RS-232 data frames from sensors for transport over the network. Multiple terminal server manufacturers are in the marketplace, but few are offering truly hardened, manageable equipment. Terminal servers from Digi International were selected based on their specifications, flexibility in operating modes, and manageability via SNMP. In addition to the purchased items, Digi International provided evaluation equipment as part of the project.

Industrial high speed Ethernet is making great strides in the marketplace with multiple vendors offering products that are temperature hardened and field ready. Several vendors of 100 megabit product lines were contacted and asked to participate in the project by providing demonstration / evaluation equipment. GarrettCom Inc. agreed to participate and provided a full compliment of hardened fiber Ethernet switches to construct a four-node field cabinet demonstration network. The equipment was available (as was the evaluation equipment from Digi International) for a limited, but sufficient, time to construct and operate a demonstration network.

NETWORK DESCRIPTION

The demonstration network constructed during the project can be divided into five subsections. Two sections physically reside inside the Texas Transportation Institute's (TTI) TransLink Laboratory. Two sections reside at TTI's Highway 6 Testbed, and the fifth section is located in Austin, Texas, at TTI's I-35 Test Site. Figure 10 displays an overall view of the

demonstration network. Two segments connect via DSL links routed over the Internet. One segment connects via a point-to-point ISDN connection. The demonstration network transports Ethernet-encapsulated RS-232/422 channels, video streams, and general network data. With the exception of the DSL and ISDN lines, the majority of the network relies on 10/100 Mbps connections. The network contains fiber, wired, and wireless interconnections.

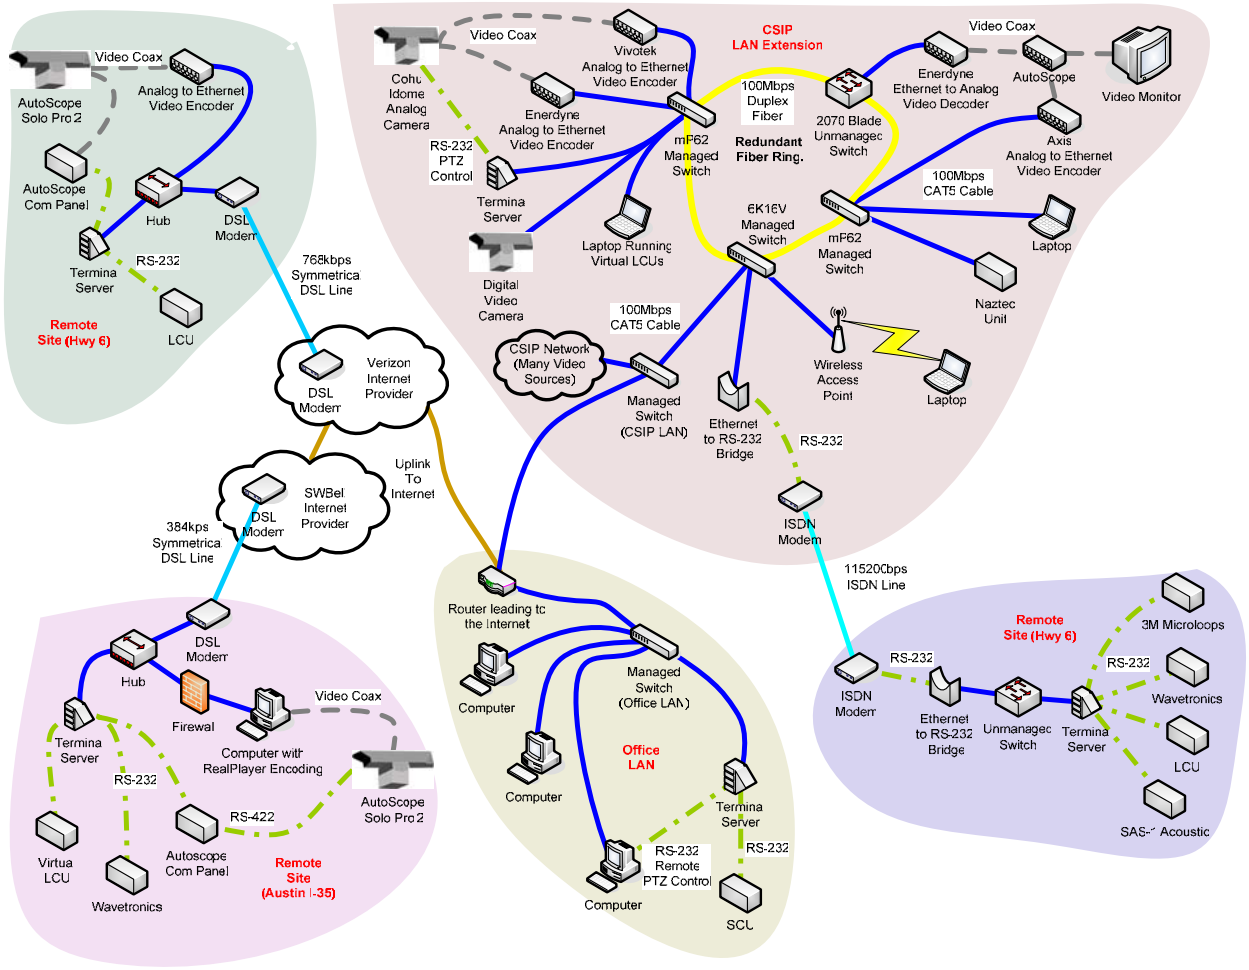

Figure 10. Demonstration Network Overview.

The demonstration network expands from a previously existing network built to support the College Station Integration Project (CSIP). The CSIP network includes a series of Ethernet-connected traffic cabinets along a corridor in College Station and transports a collection of traffic video sources to the TransLink Lab via Ethernet. The demonstration network utilized those CSIP video sources during the analysis stage of the effort. The CSIP network plays the role of the TMC to field satellite network for the ensuing experiments and discussion. The ability to

interface easily with a preexisting network demonstrates an advantage of using Ethernet equipment.

CSIP LAN Extension

Figure 11 displays a collection of the demonstration network located in the TransLink Lab at TTI's Gilchrist building. The main elements consist of two traffic cabinets and a back rack with interconnects to the CSIP network, the Internet, and an ISDN connection to the Highway 6 Testbed.

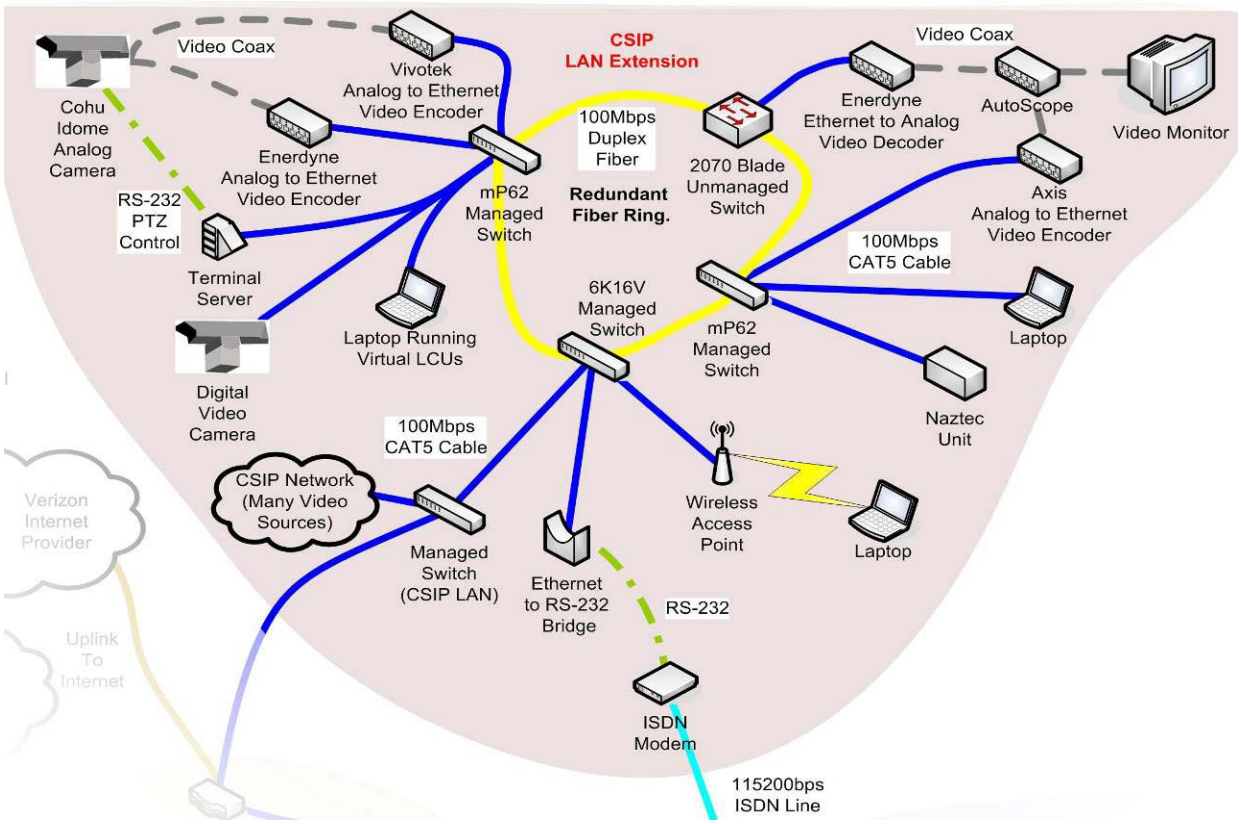

Figure 11. Test Network CSIP Extension.

The left top section of Figure 11 represents the "left" traffic cabinet. The right top section of Figure 11 is physically located in the "right" traffic cabinet. The bottom sections of Figure 11 are located on a "back rack" of equipment in another room. Single mode fiber optic cable connects all three areas, delivering a redundant ring physical layout. The collection of traffic signal cabinets represents the satellite to field cabinets in the previous architecture discussion. Figure 12 displays photographs of the two traffic cabinets ("left" and "right") and

the rack of equipment in a back room that ties the ISDN, Internet, and office extensions to the demonstration network. The rack also hosts the traffic cabinet ring's connection into the CSIP network. The rack area represents the functionality of the field satellite and is the point where a field cabinet network connects into the TMC broadband network. In the case of the CSIP network, the broadband is gigabit Ethernet. In some TxDOT implementations, the connection will likely be Ethernet over ATM. In either case, the function is to collect data from a field node network and multiplex it over a high bandwidth link back to the TMC.

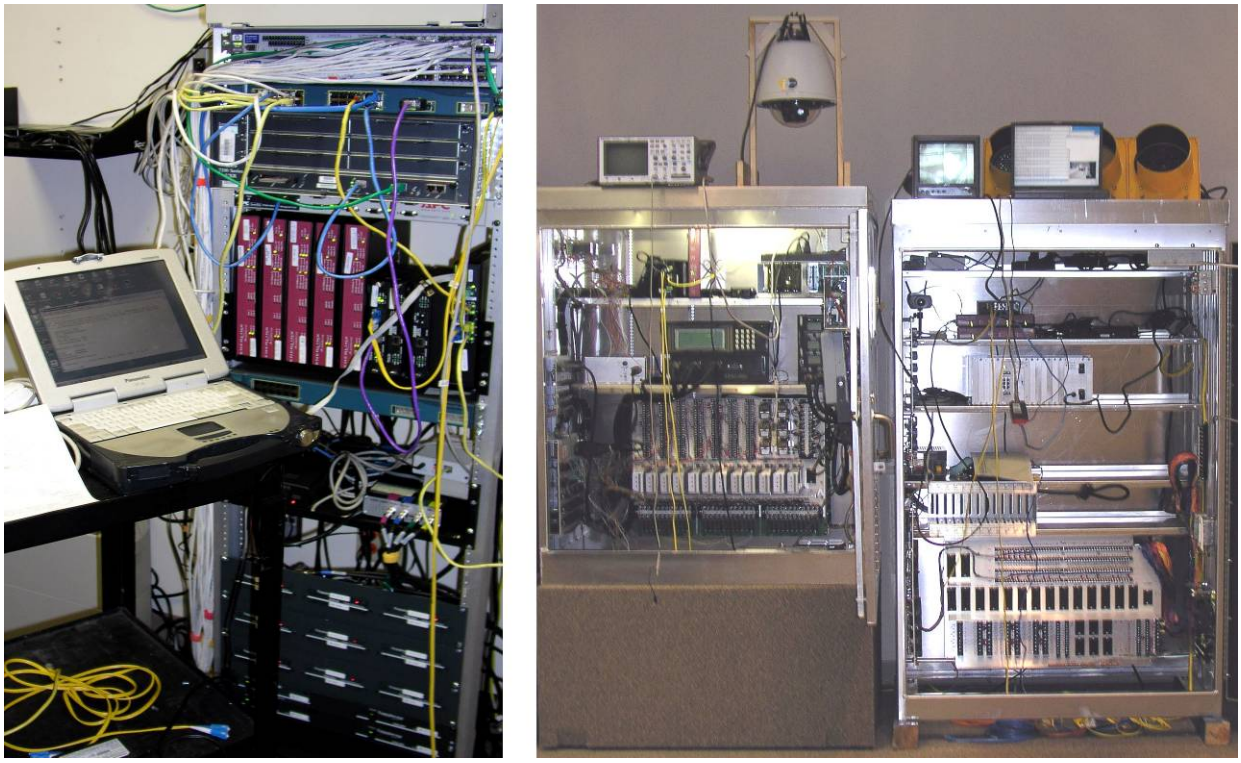

Figure 12. CSIP Network Extension Photographs.

Center Ring

Figure 13 highlights the center of the CSIP LAN extension and represents the field cabinet ring. Several vendors were contacted and asked to participate in the demonstration by providing network equipment to construct a 100 Mbps Ethernet fiber ring. GarrettCom Inc. responded and provided all the hardened Ethernet equipment to construct the traffic cabinet network. The cabinet network is comprised of three managed hardened Ethernet switches and one unmanaged hardened Ethernet switch blade installed in a 2070 traffic signal controller.

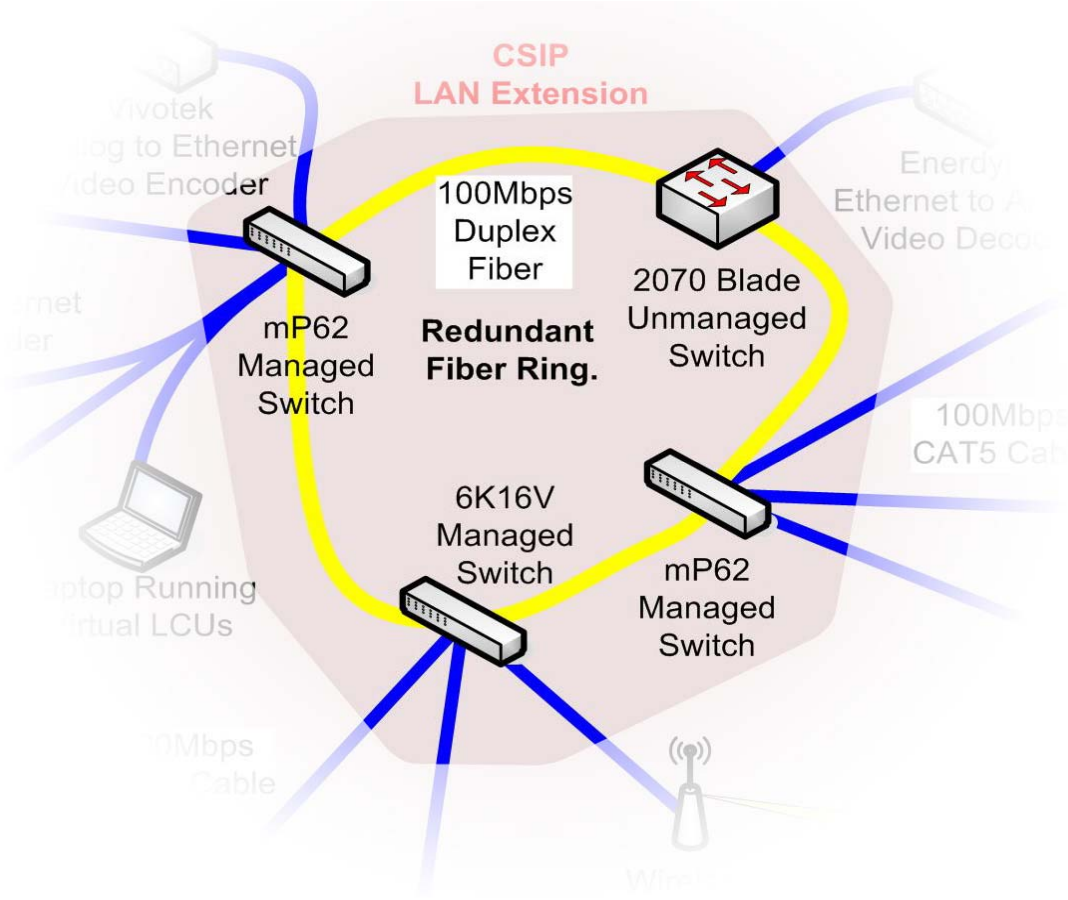

Figure 13. Redundant Fiber Ring.

The four switches are connected in an arrangement forming a redundant ring. Ethernet switches must be configured to operate properly in a ring. An administrator must connect to the switches and configure them for advanced networking. The unmanaged switch has no capability for an administrator to alter its operation and, therefore, operates with factory settings. All of the switches are configured to engage a protocol called spanning tree. The spanning tree protocol manages redundant physical connections between network equipment, and it allows for correct operation of the network. Even though the switches are physically configured in a ring, logically the network must operate as a bus or stub. The switches create the bus topology by logically severing the network but keeping the connection alive for future use if required. The logical cut is created by not forwarding packets from one switch to another over an active link. The switches share administrative information but no data flows over this severed path.

Figure 14 shows the mP62-5V ® managed switch. Both traffic cabinets contain mP52-5V switches. These particular models (including the power supplies) are hardened to accommodate environmental temperatures ranging from −40° to 170° F (−40° to 75° C) and are fully sealed, small, and durable.

Figure 14. GarrettCom mP62-5V Managed Switch.

They consume 10 watts of power during operation and are able to receive voltage ranging from 20 to 60 volts. The mP62-5Vs have two fiber optic modules installed on them as well as six RJ-45 10/100 Mbps connections. The mP62-5Vs are accessible to an administrator via a DB-9 console serial port attached directly to the switch. It is also possible to affix an IP address to the mP62-5V and administer it remotely via telnet.

The unmanaged switch is located in the right cabinet as well. Fiber loops link it to the mP62-5V managed switch in the right cabinet, and to the 6K16V managed switch in the back rack. Figure 15 shows the GarrettCom Magnum ITS 2070 Blade Unmanaged Switch ®.

This Ethernet switch is an insertable card that fits into a 2070 chassis. The device only receives power from the 2070 chassis. All other connections are external. In this case, there are two fiber ports and six 10/100 RJ-45 ports on the switch. Similar to the mP62-5Vs, the 2070 switch is temperature rated for −40° to 185° F (−40° to 85° C). It was designed to exceed CalTrans and advanced traffic controller (ATC) requirements and Telcordia GR-63-CORE

Sections 4.3.1 and 4.4.3. The 2070 has a power consumption of 8 watts, with maximum peaks of 10 watts. It is powered by a 24 VDC and −48 VDC internal power supply. The space constraints of a 2070 card prevented the implementation of a managed switch. Currently, the only available switch from GarrettCom Inc. with the 2070 footprint is an unmanaged model. This unmanaged model did include the necessary spanning tree algorithm for handling redundant physical paths between neighboring switches. A GarrettCom 6K16V ® managed switch was the anchor node and provided the uplink to the CSIP network. It was located in the rack of equipment in the back room.

Figure 15. GarrettCom Unmanaged Switch, 2070 Card Footprint.

The 6K16V (Figure 16) is a hardened switch with environmental temperature ratings of from −40° to 140° F (−40° to 60° C). It consumes 50 watts during normal operation and requires a standard 115 VAC power plug.

Figure 16. GarrettCom 6K16V Managed Switch.

The 6K16V is intended for use as the main switch in a network. It has the highest performance of all the switches mentioned here, and it supports the widest range of advanced network options. Unlike the 2070 or mP62-5V switches, the 6K16V supports the rapid spanning tree protocol. A full network of switches running the rapid spanning tree protocol is able to recover from severed physical connections in a matter of seconds rather than minutes. Like the mP62-5Vs, the 6K16V can be administrated from a serial console port directly connected to the chassis, or it can be assigned an IP address and be remotely administrated via telnet client software.

Left Traffic Cabinet

Figure 17 displays one-third of the CSIP LAN extension. This resides in the "left" of two traffic cabinets located in the TransLink Lab.

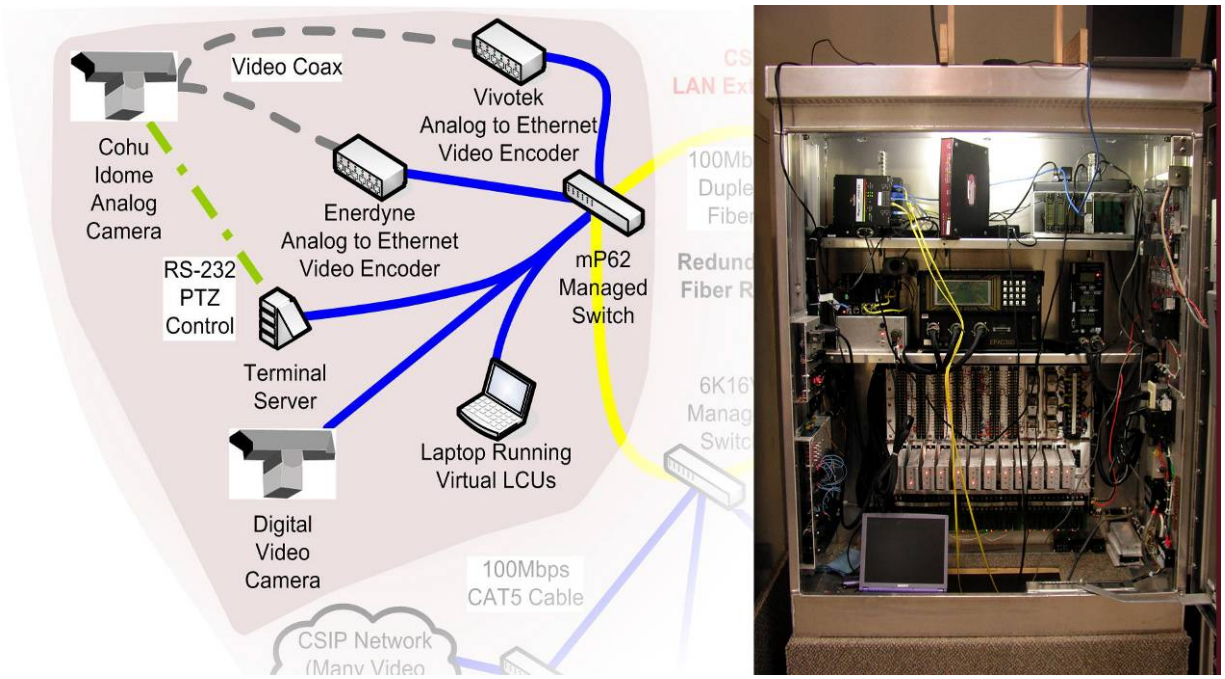

Figure 17. Left Cabinet Topology and Photograph.

An analog Cohu Idome video camera produces a video output and accepts pan-tilt-zoom commands over an RS-422 serial interface. An Enerdyne LNX7000-01 ® MPEG2 video encoder and a Vivotek ® analog to Ethernet motion Joint Photographic Experts Group (JPEG) video encoder both utilize the Cohu camera's video output as their video input source. The video encoders were used to create broadband, delay sensitive traffic on the network. The RS-

422 serial control of the Cohu camera is routed into a Digi International Portserver TS 4 ® serial to Ethernet terminal server. This device encapsulates serial RS-422 signals into Ethernet packets and routes them to another terminal server located on the office LAN. Commands sent by the remote end camera control software package are routed to this Digi device, which in turn converts the commands into RS-422 signals for the camera's controller.

A D-Link Moving Picture Experts Group Standard 4 (MPEG-4) video camera stationed in the left traffic cabinet serves as a second video source. The D-Link camera does not produce analog video outputs. The camera encodes its video directly to a digital stream, provided to the user by a web server running on the camera. This enables any computer with an Internet browser to display the video by typing in the camera's webpage address. The camera was mainly used to create time sensitive network traffic and would not be a normal device deployed by TxDOT.

The last network host contained in the left traffic cabinet is a laptop computer. This laptop runs a program that emulates the functionality of multiple LCUs. This particular laptop emulates five unique LCUs through a single serial port. The LCU RS-232 data are sent to a channel on the Digi Portserver for encapsulation. The LCU encapsulation utilizes UDP packets with routing to a terminal server channel connected to an SCU. By encapsulating and transporting serial data via Ethernet, the LCUs and SCU remain in serial contact despite their distance. Notice that both video and serial data are using the same, shared Ethernet resource to transport data back and forth.

Right Traffic Cabinet

Figure 18 shows the next segment of the CSIP test extension. The "right" traffic cabinet contains two switches. Functional necessity requires only one switch to serve the cabinet; but for the sake of testing, the unmanaged 2070 switch was placed inside and connected as part of the fiber ring.

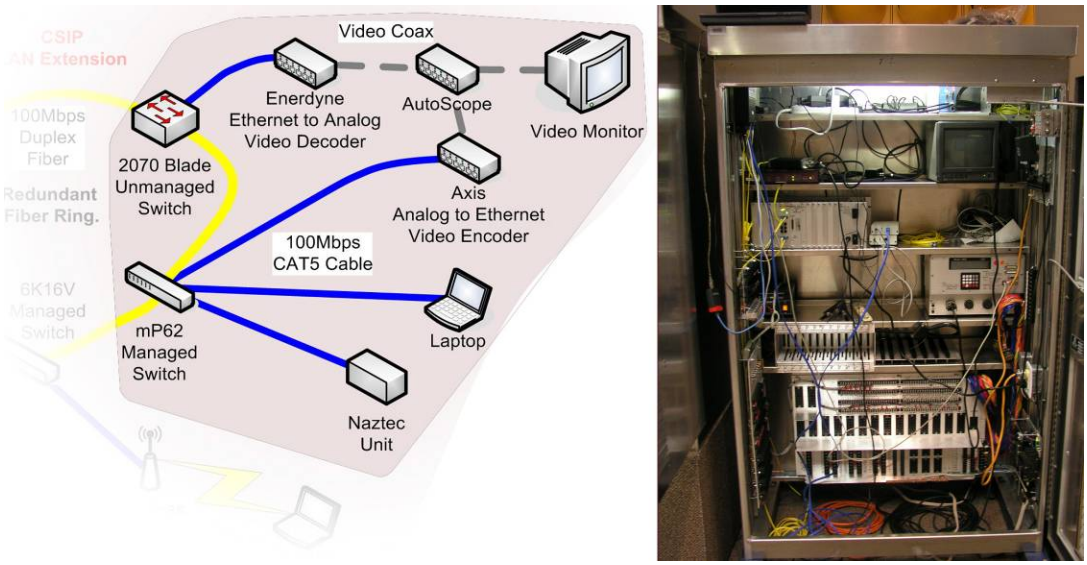

Figure 18. Right Traffic Cabinet Topology and Photograph.

The right cabinet contains an Enerdyne LNX7000-01 video decoder. The Enerdyne decoder is able to "tune in" on any of five different multicast streams of video data provided by the CISP network, and one stream provided by the Enerdyne encoder in the left cabinet. The analog video signal produced by the Enerdyne decoder connects to an AutoScope Solo Pro 2 vehicle detection device and to an Axis ® analog to Ethernet Motion JPEG video encoder. The Axis video encoder operates in a method similar to the D-Link MPEG-4 camera. The Axis encoder accepts up to four analog video signals. It encodes these signals digitally and serves these four streams from a webpage server running on the device. The user types the Axis encoder's IP address into an Internet browser and accesses these four video feeds singularly, or in a two by two composite. The user also has the option to integrate the Axis' video feeds into other webpage documents by including a specific text string that summons a video stream to the third-party webpage. The Axis encoder can produce a video data stream of multiple megabits per second and is used to create additional network traffic. The final device located in the right cabinet is a Naztec Traffic Signal Controller. The controller contains an Ethernet port that directly connects the controller to the LAN and provides for remote management of the device.

Back Rack Equipment

The "back rack" area (Figure 19) serves as an interconnection point for the TTI office LAN, the Highway 6 ISDN link, the main CSIP network, the Internet that leads to the Austin and

Highway 6 DSL segments, and the CSIP extension network constructed for this project. A Cisco Aironet 350 ® 802.11b Wireless Access Point (WAP) attaches to the main 6K16V managed switch. The Aironet WAP serves a collection of laptops with wireless cards. This allows for physical portability of the laptops during times of network maintenance or general use of the network such as using a wireless laptop to display a video stream served by the Axis video encoder.

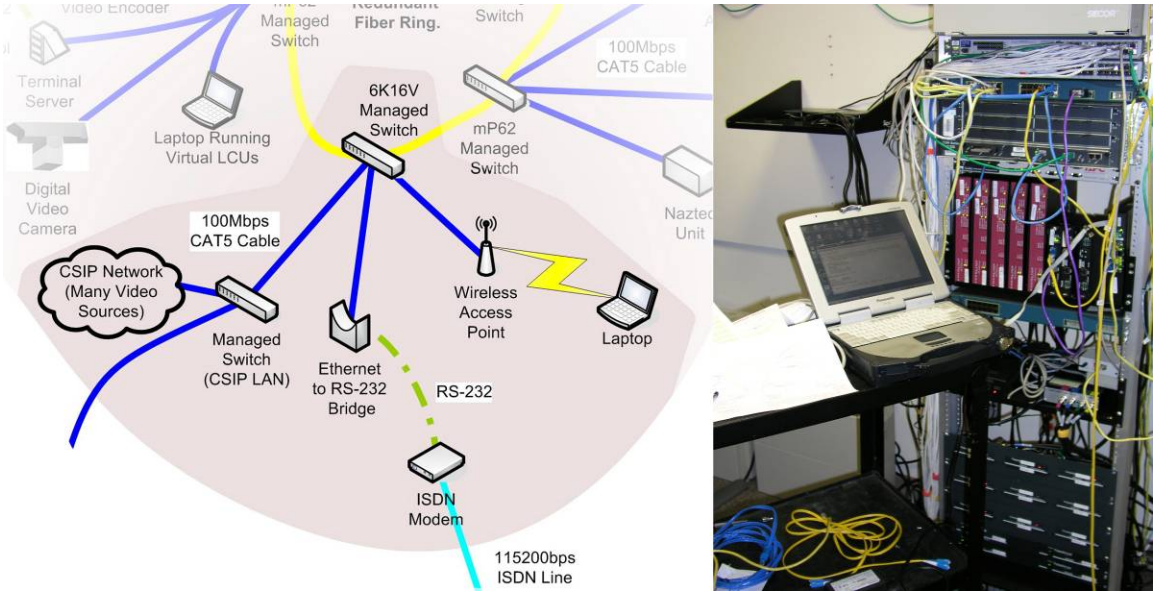

Figure 19. "Back Rack" Equipment.

The main 6K16V managed switch uplinks to the main switch for the greater CSIP network. The modularity of Ethernet allows both networks to share their resources after the simple act of connecting a twisted-pair CAT5 crossover cable between them. The two switches talk to each other and soon begin to forward data between both the CSIP network and all of the LAN extensions constructed for this project. The ISDN connection from Highway 6 routes to the 6K16V managed switch. This link is established by using a DataComm for Business (DCB) asynchronous bridge coupled with an Adtran Express 3000 ® ISDN modem, as shown in Figure 20.

Figure 20. DCB Ethernet Asynchronous Bridge (left) and Adtran Express 3000 ISDN Modem (right).

TTI Office LAN

Extending from the interconnections in the back rack is the TTI Office LAN (Figure 21). The Office LAN represents the LAN inside the TMC. All data from the field and the experimental fiber ring in the test cabinets are logically routed here. Any computer or device on the office LAN is able to view content from the field and from the test cabinets. A terminal server is such a device and was used to interface to the Cohu camera control software and to an SCU. One channel of the SCU was fully populated with real and emulated LCUs, all connected via terminal servers and Ethernet. LCUs were located on Internet links (Austin and College Station), an ISDN link (Highway 6 Testbed), and emulated LCUs in a test cabinet. Figure 22 displays the SCU and associated terminal server connected to the TTI Office LAN segment.

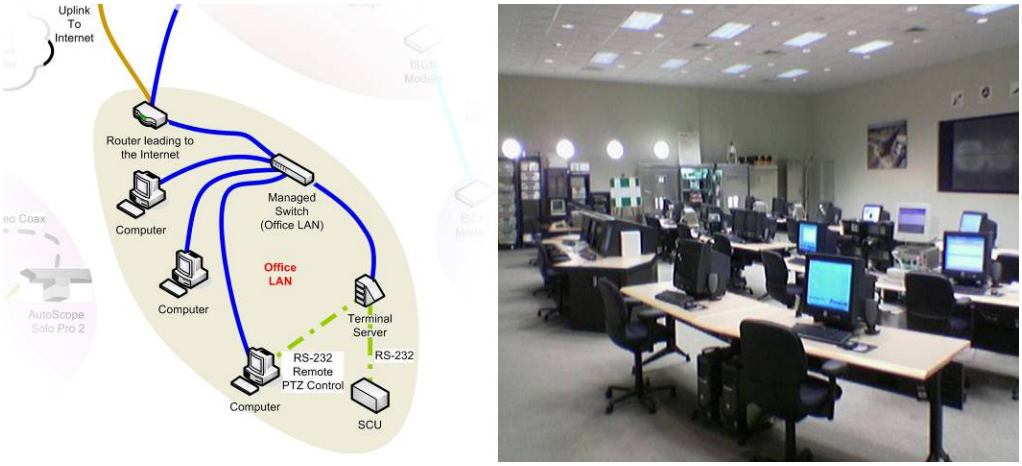

Figure 21. TTI TransLink Lab, Endpoint of Most Test LAN Devices.

Figure 22. TTI Office LAN, SCU, and Digi Terminal Server.

Highway 6 ISDN LAN Segment

There were two links connecting the Highway 6 Testbed to the Test Network. Functionality requires only one link, but two exist for the purpose of testing. One of those two links is an ISDN connection. The ISDN link begins from the back rack. There an Ethernet asynchronous bridge converts Ethernet traffic into 115,200 bps RS-232 data. These serial data feed into an Adtran Express 3000 ISDN modem. The modem transfers the serialized network data through a Verizon ISDN service and onto a matching Adtran ISDN modem located at the Highway 6 Testbed (Figure 23).

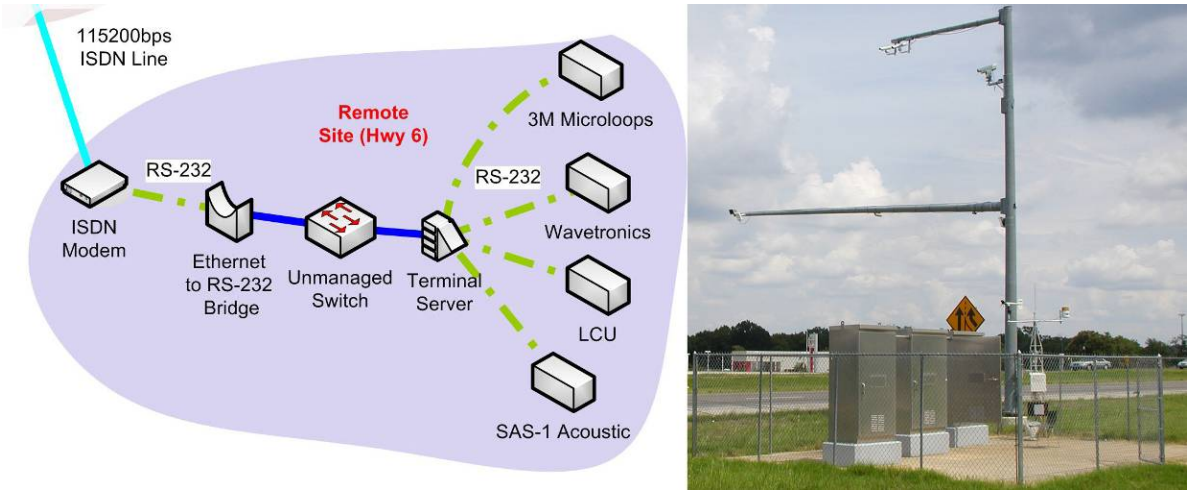

Figure 23. Highway 6 ISDN LAN Segment.

The Highway 6 modem channels the serial data back to a DCB asynchronous bridge (Figure 24, 4). The DCB unit forwards its Ethernet output to a Linksys ® unmanaged switch (Figure 24, 2) that distributes the network data to connected network hosts at the test site. Spanning from the Linksys switch are a set of traffic detection devices such as 3M micro loops, a Wavetronics ® radar unit, another LCU (Figure 24, 3), and a SAS-1 Acoustic sensor. All of the sensors forward their RS-232 data streams to a Digi terminal server (Figure 24, 1). Their data return to the back rack and then move onward to the TTI Office LAN for remote analysis and testing.

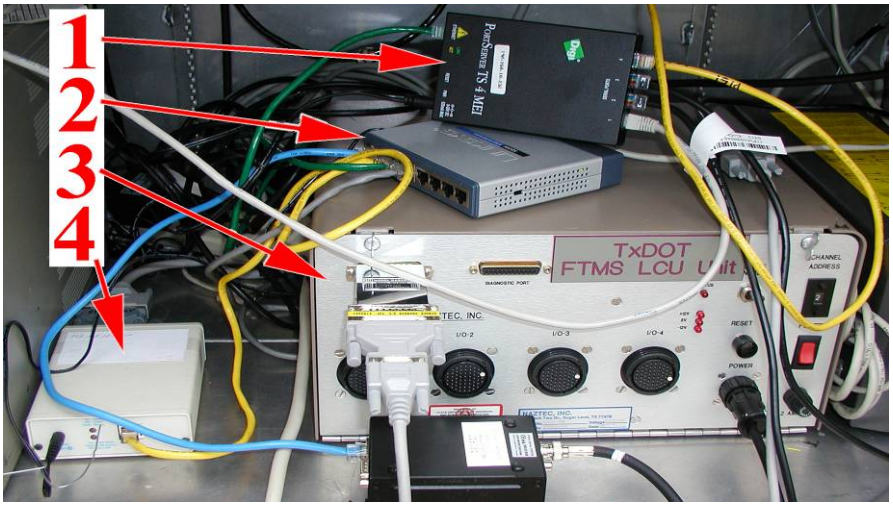

Figure 24. Highway 6: Digi 4-port Terminal Server (1), Linksys Unmanaged Switch (2), LCU (3), DCB Asynchronous Bridge (4).

Internet Linkage

The Internet was used to bring back information from two remote locations (Figure 25). The office LAN connects to a router that acts as a gateway leading to the Internet. From there, the Texas A&M University network merges with the Internet. TTI bought leased DSL services from both Verizon ® and Southwestern Bell ®. Southwestern Bell provides a 384 kbps symmetrical DSL link to the Internet for TTI's Interstate 35 test site in Austin. Verizon provides a 768 kbps symmetrical DSL link to the Internet for TTI's Highway 6 Testbed.

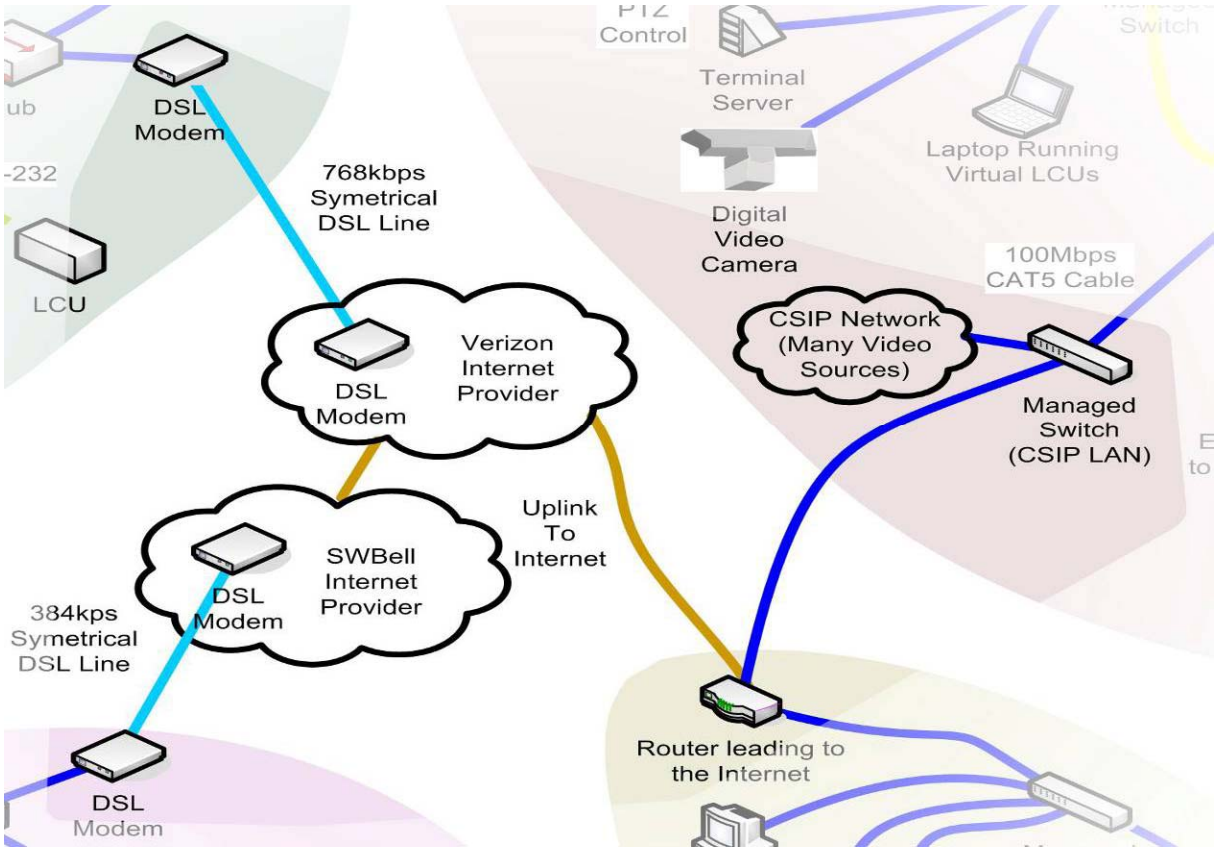

Figure 25. Internet Connections.

Highway 6 DSL LAN Segment

The second of two links connecting TTI to its Highway 6 Testbed consists of a 768 kbps symmetrical DSL connection provided by Verizon Wireless (Figure 26). Network traffic routes to the Internet from the DSL modem (Figure 27, 1) located at Highway 6, and from the Internet into the TTI Office LAN via a gateway router connected to the Office LAN. This segment contains an Axis four-port analog video to Ethernet encoder. Video from the various traffic cameras located on a nearby pole serve as Axis encoder inputs. A web server running on the Axis encoder hosts the digitized video. Typing the encoder's IP address into an Internet browser accesses the video streams. It is possible to password protect the Axis encoder's webpage to restrict unauthorized viewing or administration of its settings.

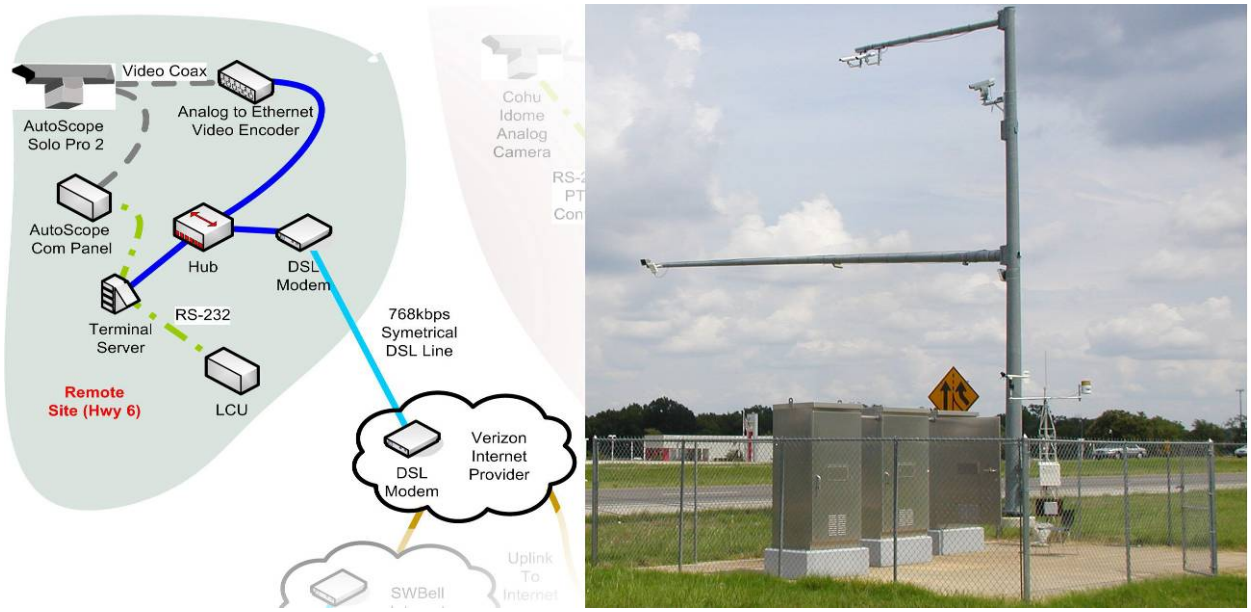

Figure 26. Highway 6 DSL LAN Segment.

Since the network traffic coming from this segment is light, an inexpensive hub (Figure 27, 2) distributes network access to the hosts in this segment. RS-232 data converts to Ethernet through a Digi portserver. In this LAN segment, serial data from an AutoScope traffic detection device and serial data from a LCU convert to Ethernet and backhaul to TTI's Office LAN. Network equipment susceptible to hacking was placed behind a firewall (Figure 27, 3) before their connections route to the hub.

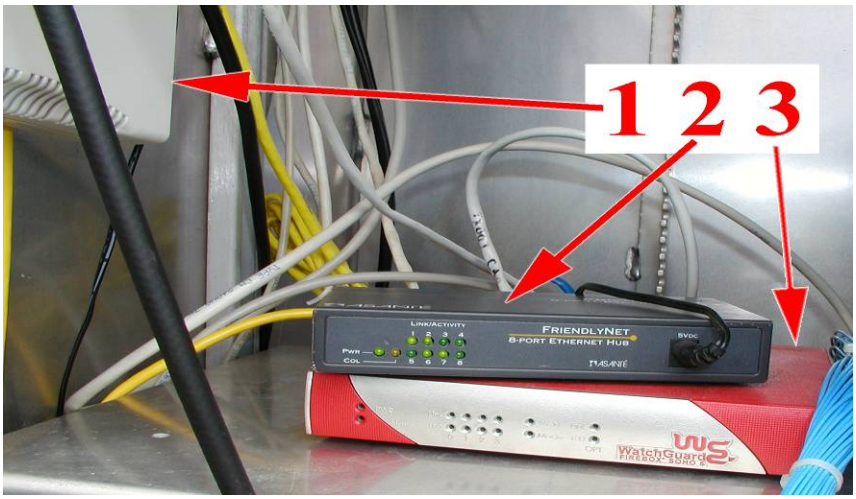

Figure 27. Highway 6: DSL Modem (1), Hub (2), Firewall (3).

It should be noted that the Digi serial terminal servers have an option where their serial ports can be directly mapped to a virtual COM port on a remote personal computer (PC). For example, one computer in the TTI Office LAN has a virtual com port #6 that binds to a com port located at the Highway 6 Testbed. Software normally intended to run on a PC with hardware serially linked nearby (such as the AutoScope) is now able to run on remote machines. To the computer, the virtual com port behaves just like a standard com port.

Austin I-35 DSL LAN Segment

As part of another project, TTI purchased a symmetrical DSL link to the Internet for its testbed located alongside I-35 in Austin (Figure 28). Located at the site are sensors whose purpose is to monitor freeway traffic. An AutoScope video detection system is installed at the site and provides traffic detection. The RS-232 output of the device connects to a Digi serial terminal server and forwards on to the TTI Office LAN for interfacing with AutoScope software running on a PC.

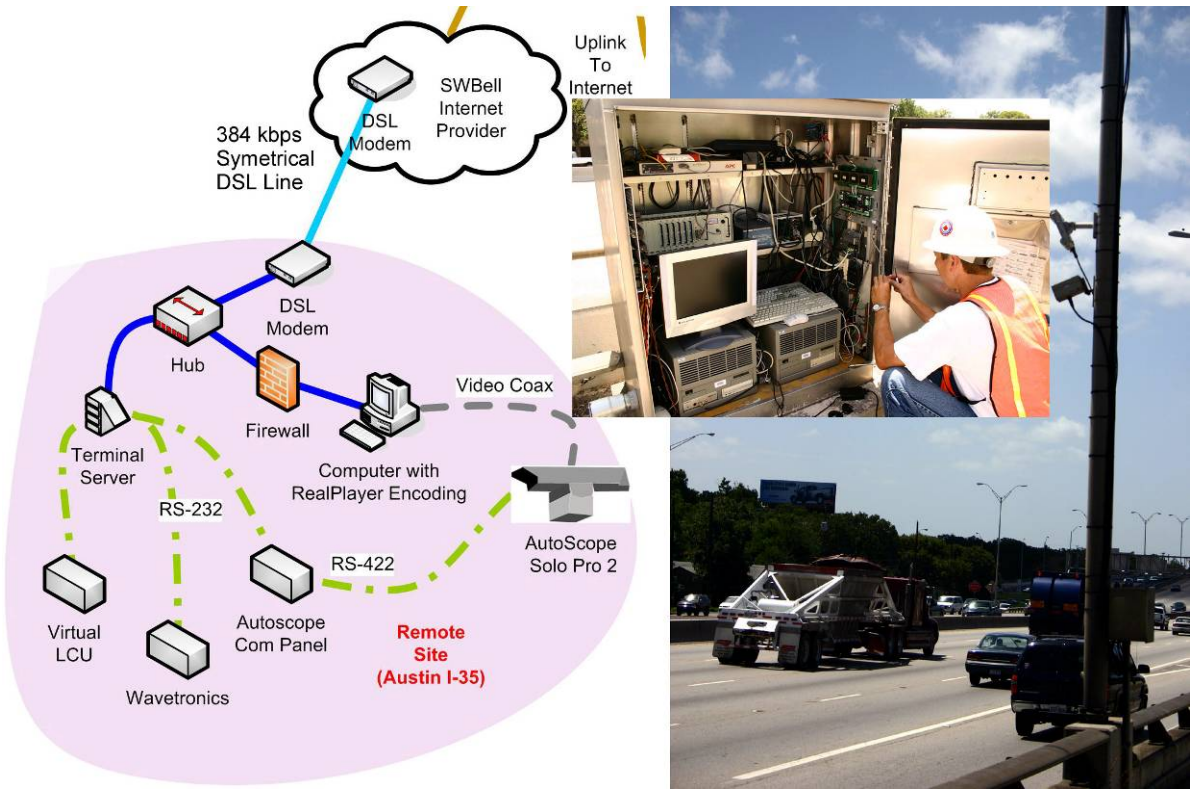

Figure 28. Austin I-35 Testbed DSL LAN Segment.

The Digi terminal server also collects serial data from a Virtual LCU and a Wavetronics radar sensor. Traffic leaving the DSL modem (Figure 29, 1) distributes through a hub (Figure 29, 4) to all attached network devices. An APC networked power bar (Figure 29, 2) connects to the hub to allow for remote rebooting of individual devices. A router/firewall (Figure 29, 3) provides network security for this site.

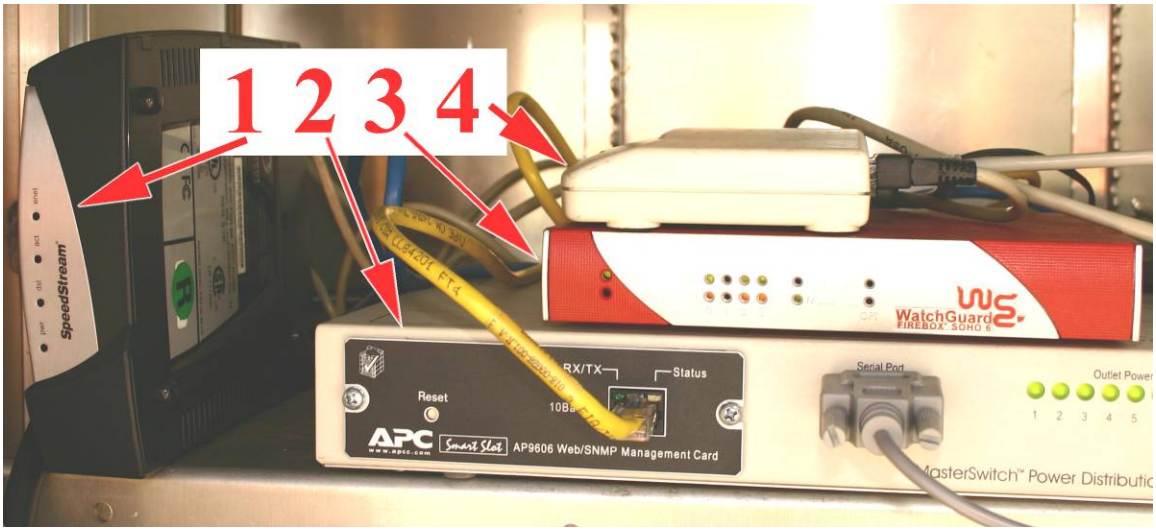

Figure 29. DSL Modem (1), Remote Rebootable Power Supply (2), Firewall (3), Hub (4).

NETWORK OPERATION

The Internet link to the Austin Test Site and the Highway 6 Testbed was in place before this project began, as well as the link to the CSIP network and the TTI Office LAN. The new portions of the ITS network consisted of the hardened fiber ring in the cabinets and the ISDN leased line link. The efforts on the preexisting network were centered on device integration and communication among networks. Numerous traffic sensors were integrated onto the experimental network both in the remote locations and in the lab experimental cabinets. The network as described above was operated in a manner similar to that expected for an ITS network.

Looking back at the private network and backbone architecture, the Office LAN represents the LAN in a TMC. The TMC LAN would contain sensor management devices and

software such as a grouping of SCUs, dynamic message sign control software, camera control software, protocol converters, etc. The demonstration TMC network hosted:

- an SCU with a terminal server to connect a full channel of LCUs
- a computer running camera control software to operate a Cohu camera located near one of the cabinets
- a computer to host Naztec traffic signal controller management software
- various computers to host sensor vendor software and to view digital video
- a computer to monitor the network using an SNMP based management package

The Office LAN connected to the CSIP network, which provided a broadband conduit to the experimental fiber ring network. The CSIP network performed the role of a broadband network backbone connecting multiple field satellites. In the project, the broadband backbone role was provided by a gigabit Ethernet network. The TxDOT equivalent will likely be an ATM backbone that transports video on native ATM from field satellites as well as supporting a broadband Ethernet switch in the satellite. The satellite functionality is represented by the back rack area. The satellite is the location for the physical connection between the field cabinet network and the broadband backbone. In the demonstration network, the anchor switch for the field cabinet ring network uplinked to the CSIP switch in the back rack (field satellite). The demonstration fiber ring network created from the vendor-loaned hardened Ethernet switches represents a satellite to field cabinet, daisy chained ring network. The ring network represents physical cabinets with communication drops located in the field near the sensors. Finally, the ISDN and Internet network segments provide connectivity to sites not located on a fiber infrastructure and represent purchased network services.

Upon arrival of fiber network equipment, the network pieces were quickly deployed. The loan time was limited; therefore, all network setup, testing, evaluation, and demonstration had to be completed in a matter of weeks. It is important to restate that the demonstration network was not operated for an extended period of time. The network was established to show the feasibility of the technology and, therefore, long term analysis is outside the scope of the effort. The network was operated successfully for 2 weeks after the initial setup and debug phase. The Ethernet network showed to be a reliable means for moving sensor data as well as video, although video experiments were not pursued extensively. The fiber ring cabinet network was shown to recover from a loss of a fiber link, simulating a fiber optic cable break in the field. The

network recovered by using the spanning tree protocol. Although the recovery time was not instantaneous, it was reasonable for ITS operations. Latency across the fiber network was measured to be very low (less than a millisecond) and did not create any issues with device communication or control. Latency increased for Internet connections but was still very reasonable at 30 to 40 milliseconds for a round trip. The most latency intensive network segment was the ISDN link with most of the delay being attributed to the telephone provider portion of the network path. Delay across the ISDN line alone was approximately 100 milliseconds.

Many sensors were integrated onto the network, and fully functioning connections were operated to the TMC for periods of days. The Ethernet solution, utilizing hardened terminal servers, proved to be a reliable means for moving legacy RS-232 communication channels over an Ethernet network. All of the equipment deployed, both network fabric and terminal servers, included an SNMP agent onboard making network management quite easy with a standard SNMP monitoring package. A computer in the TMC was assigned to monitor the network. The application essentially monitored the health of the network and served as a tool to diagnose issues during the construction and configuration of the new equipment. The software tracked device availability, response time, and bandwidth utilization.

TECHNOLOGY DEMONSTRATION

On August 30, 2004, a half-day technology demonstration was held for the TxDOT project management staff. A presentation reviewing center to field communication needs and how Ethernet technologies addressed the issues was given to the attendees. We discussed sensor interfaces and bandwidth needs, legacy equipment integration issues (specifically the LCU) and system management tools and techniques. The network as described above was fully operational, and numerous discussions and demonstrations were conducted to exercise the network and to provide live, hands-on experience with the technology.

APPENDIX A – ETHERNET CABLING

There are four major types of Ethernet cabling in use today, two of which are in decline.

- RG-11 coaxial cable
- RG-58 coaxial cable
- unshielded twisted pair (UTP) cable
- fiber optic cable

Table 3 displays a comparison of these media types.

Table 3. Ethernet Cabling Comparison.

Type	Physical Medium	Max Run Per Segment	Max Connections Per Segment	Transmission	Topology
10base2	RG-58	605 Feet	30	Half Duplex	Bus
10base5	RG-11	1640 Feet	100	Half Duplex	Bus
100baseT	UTP	328 Feet	2	Full Duplex	Star
10baseT	UTP	328 Feet	2	Full Duplex	Star
10baseFL	Fiber	6560 Feet	2	Full Duplex	Star

Coaxial cable can be used to create an Ethernet LAN, but it is not commonly used today because of its half-duplex limitation, higher expense per unit length, and bandwidth limitations and because it is harder to handle in tight spaces like a traffic cabinet.

Currently, UTP cable (Figure 30) is the most popular media for Ethernet transmission. Most applications utilize a transmit pair and a receive pair of wires. This use leaves the remaining two pairs available for other uses such as power, or double channeling. Each of the four pairs of wire are twisted together to make the cable less prone to interference from other electrical signals. Category 5 (Cat 5) is the most common UTP cable. It rates for transmissions up to 100 Mbps. Other cable types will work, but they are rated for less bandwidth. Table 4 lists some common UTP cable types.

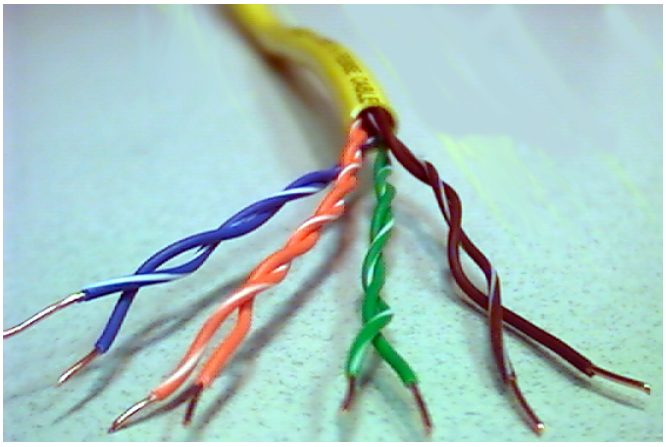

Figure 30. Cat5e UTP Cable.

Table 4. Twisted Pair Cable Types.

ANSI/EIA 586 Twisted Pair Cabling Standards		
Type	Bandwidth	Common Application
Category 1	4 Mbps	POTS, ISDN
Category 2	4 Mbps	IBM Token Ring
Category 3	16 Mbps	10BASE-T Ethernet
Category 4	20 Mbps	16 Mbps Token Ring
Category 5	100 Mbps	100BASE-T Ethernet
	100 Mbps	100 Mbps TPDDI
Category 5e	1000 Mbps	155 Mbps ATM
	(4-pair)	Gigabit Ethernet

For applications inside a traffic cabinet, Cat 5 is the best fit. The two unused wire pairs inside each cable are useful when an extra signal path is needed between various devices. Some devices located at the top of a traffic pole will have a single Cat 5 cable leading to them. Back in the cabinet, a special adapter provides an Ethernet connection to the remote device as well as sourcing the device's power through the unused wire pairs in the same cable. This setup eliminates the need for separate power and communications cables. Other situations have occurred where an extra pair on a Cat 5 cable leading to a device was used to toggle a remote reboot of the target device.

The topology of UTP cable differs from coax cable. UTP Ethernet configurations utilize a star topology as seen in Figure 31. A single cable connects at one end to an Ethernet host. The other end terminates into either a hub or switch. Multiple cables connect to and expand outward from this hub or switch, thus giving it a "star" pattern. One cable on each hub/switch serves as

an uplink route to another node of larger collections *(typically a router)*. Interconnections by multiple hosts into a hub produce similar cross talk problems as a bus topology. This problem happens because the hub blindly repeats across all connections any incoming data from a single connection. Interconnections by multiple hosts into a switch are more efficient as the switch smartly routes the Ethernet frames only to the hosts that need the inbound data.

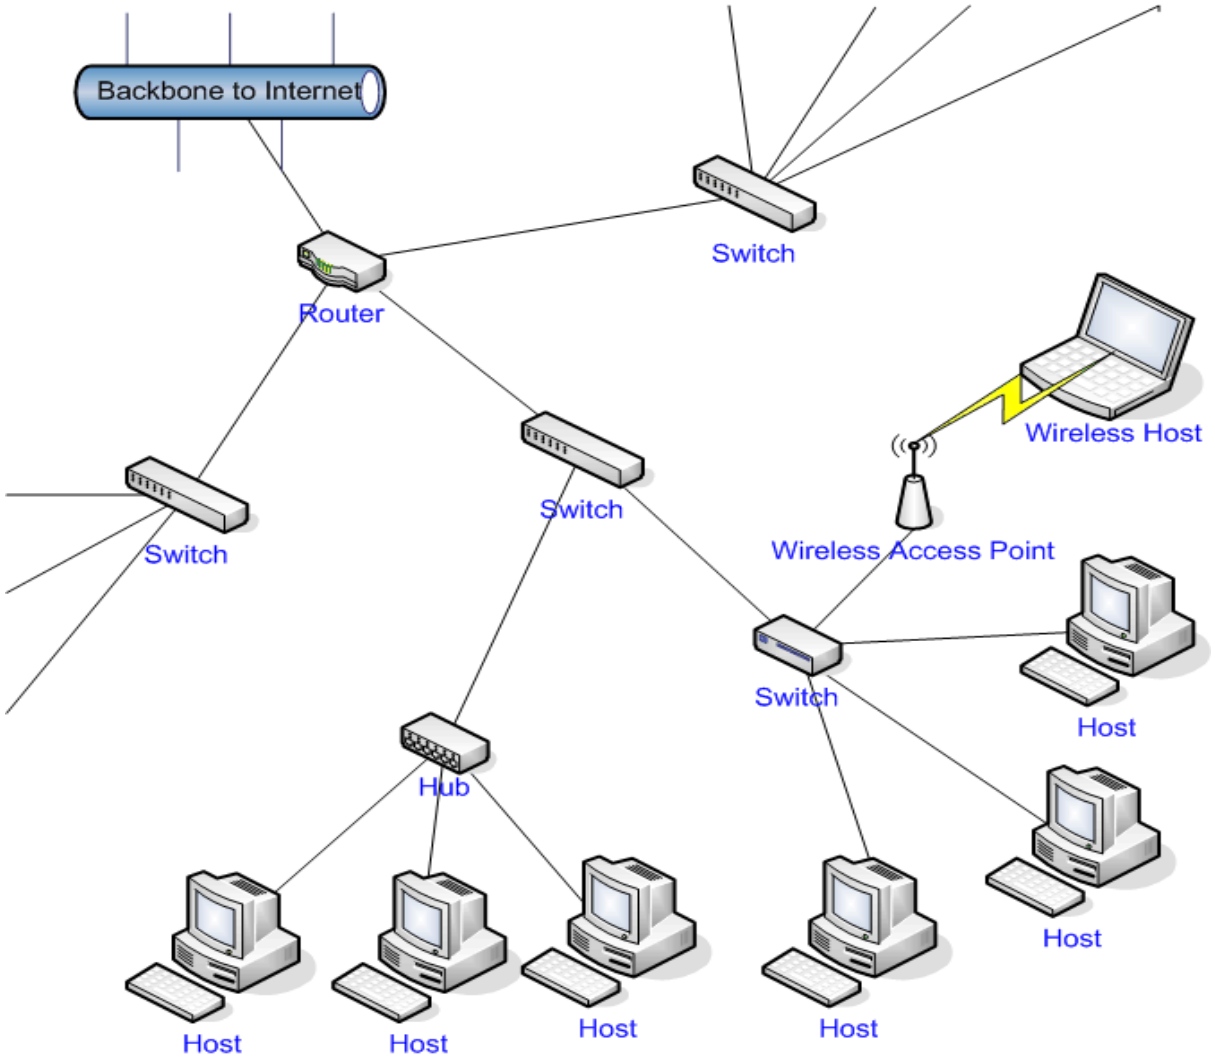

Figure 31. UTP Star Topology.

When connecting network equipment via CAT 5 cable, the gender of the equipment must be considered. Cat 5 cables have three different internal wiring patterns: straight, crossover, and rollover. The most common cable uses the straight wiring pattern. Straight cables connect networked devices of different genders. The two genders are *data terminal equipment* (DTE)

and *data communications equipment* (DCE)[1]. Networked computers and routers are DTE. Switches and hubs are generally DCE equipment. DTE to DCE connections require straight cables. Same gender connections such as DTE to DTE or DCE to DCE require crossover cables. The less common rollover cable connects a router or switch console port to a serial port on a computer. This rollover cable interface is typically used as an access channel for configuring the settings on the router or switch via a terminal emulator like Hyperterm, sourced from the attached computer. Figure 32 displays common network connections and the appropriate cable to use between them.

There is an unofficial industry habit that standard cables tend to be blue and crossover cables tend to be red. Rollover cables tend to be red, but with some distinguishing markings such as a yellow connector jacket that differentiate them from crossover cables. There is no rule that enforces this coloration, but such knowledge might be useful in a problem-solving situation.

After determining the correct connecting cable type, it is necessary to either purchase or construct the Cat 5 cable. Figure 33 shows the necessary wiring orientation between two RJ-45 connectors for straight, crossover, and rollover cabling.

Inside Cat 5 cable, there are four pairs of wires. Two industry standards (T-568A and T-568B) relate colored wire strands to specific pins on an RJ-45 connector. While it is not necessary to follow this standard, knowledge of it is useful when using a cable with only one end available. If a cable is installed in a new building, for example, there is a decent chance that the remote side was terminated using the T-568B standard. The reason for the two wiring standards was due to a correction in the T-568B standard. T-568A moves the blue and orange pairs to the center four pins of the RJ-45 plug. This allows for more compatibility with telephone voice connections.

Figure 33 is formatted for easy printing and serves as a wiring guide for CAT-5 cable with RJ-45 cable crimps.

[1] Older DCE definition: data circuit terminating equipment

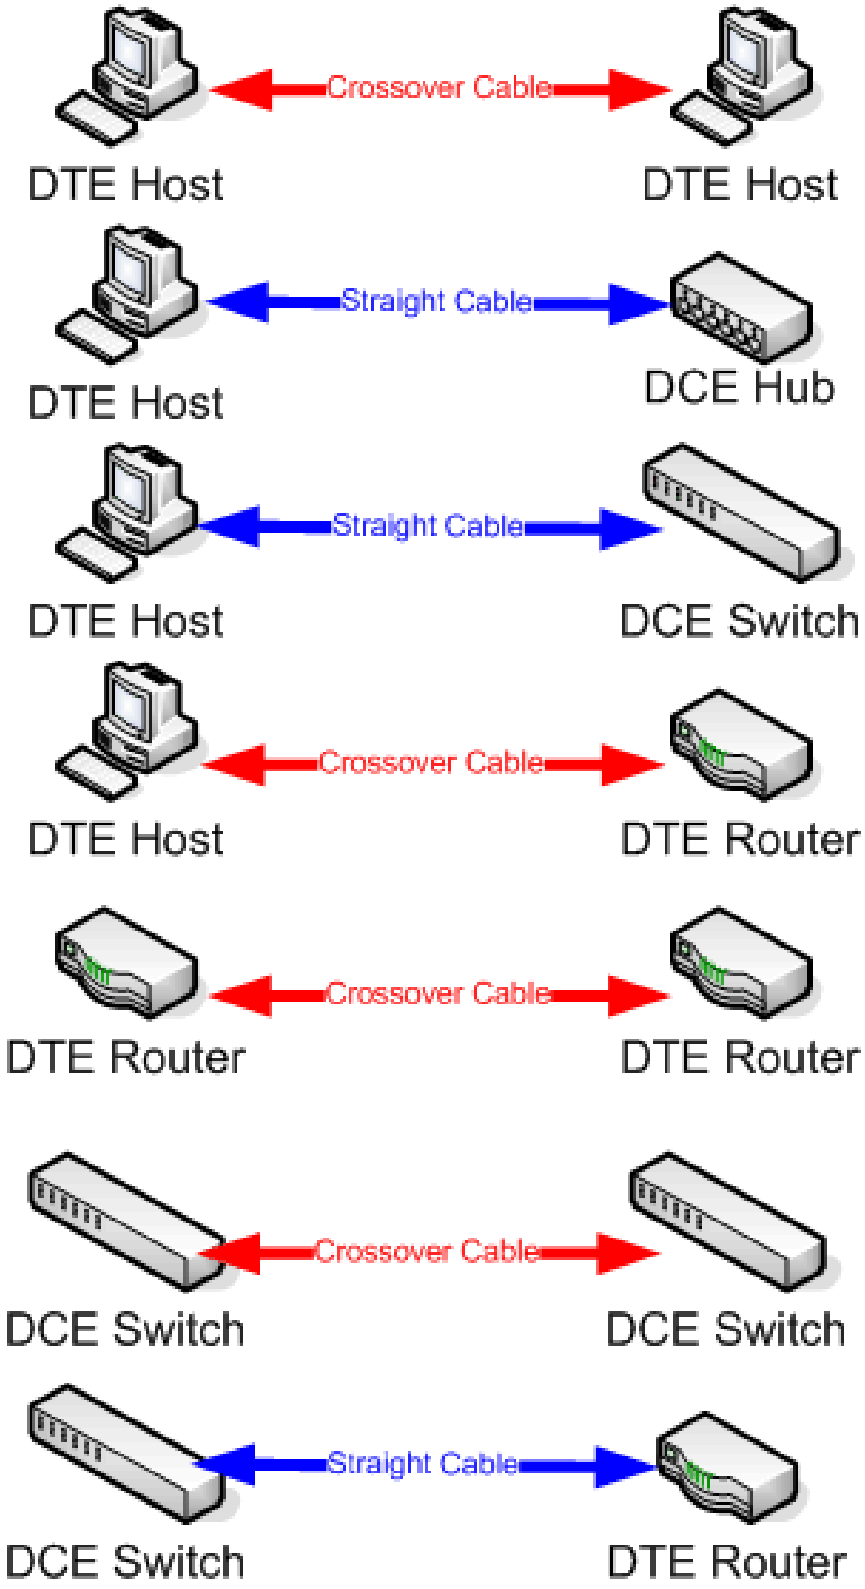

Figure 32. Crossover or Straight? How To Choose.

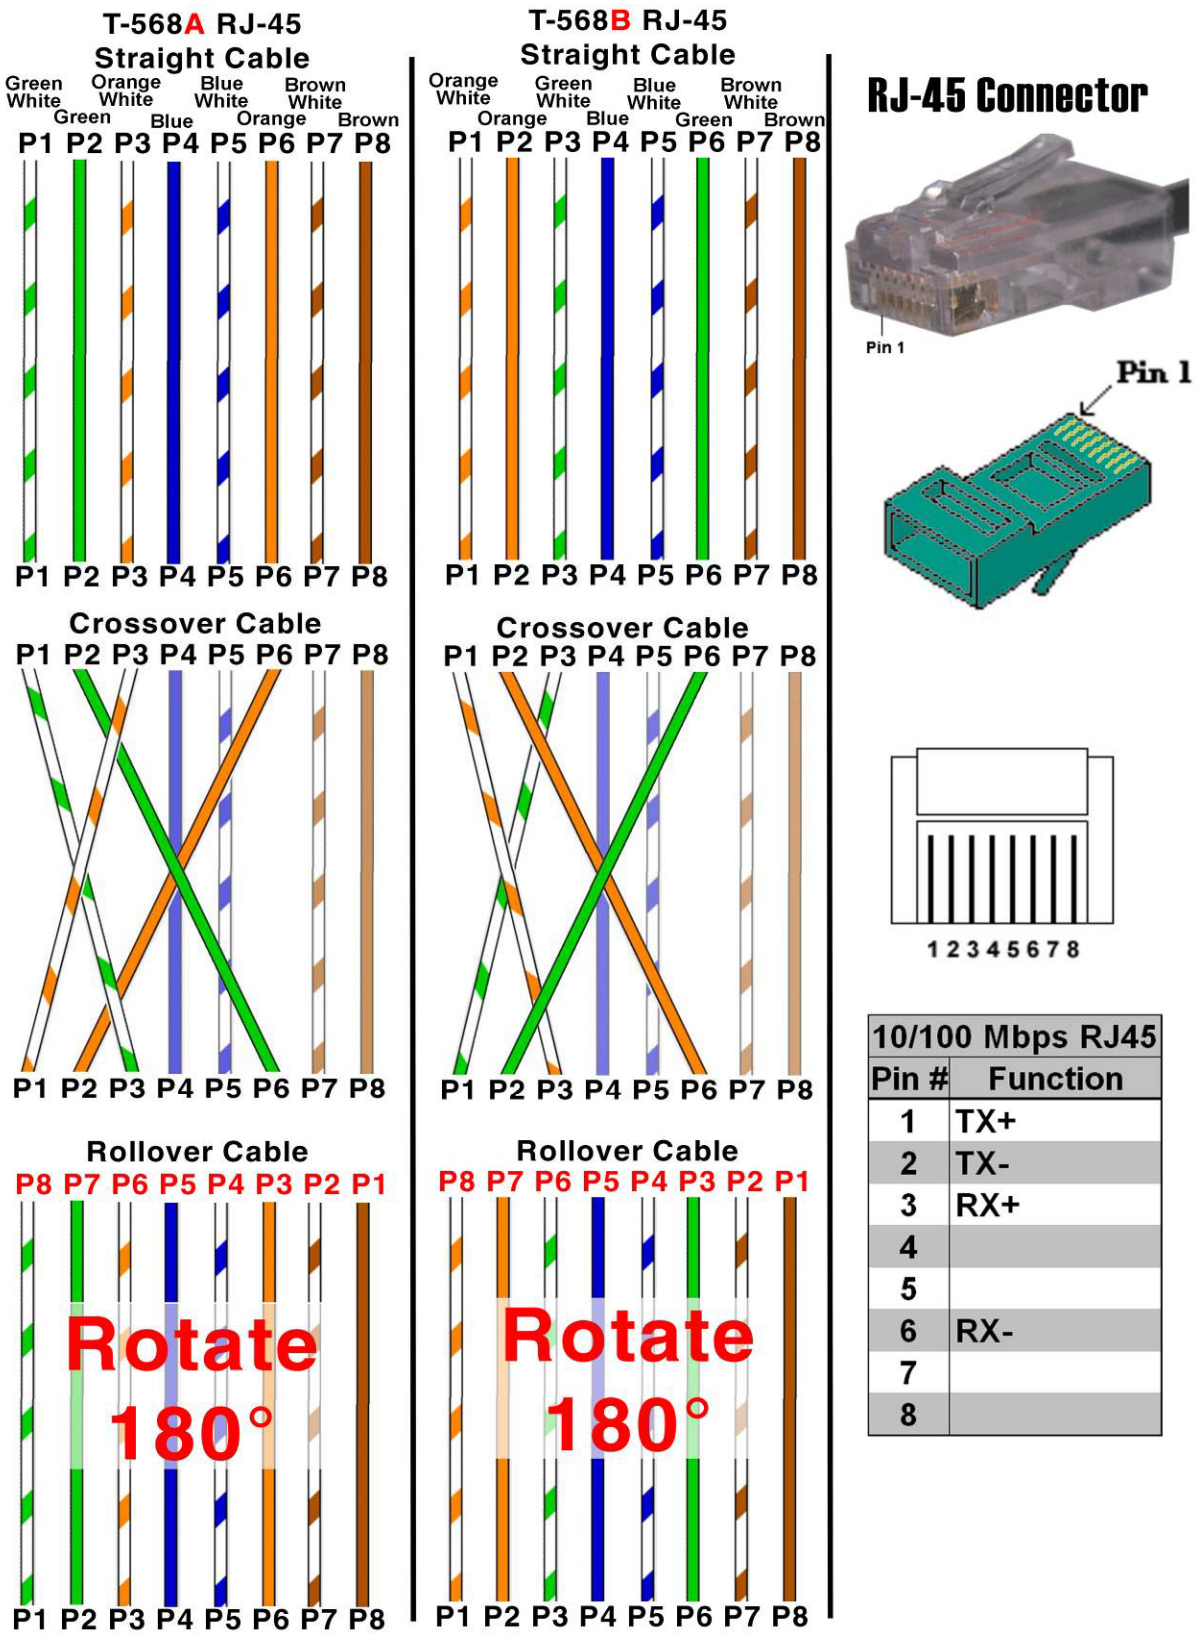

10/100 Mbps RJ45	
Pin #	Function
1	TX+
2	TX-
3	RX+
4	
5	
6	RX-
7	
8	

Figure 33. RJ-45 Cable Type Wiring Guide.

APPENDIX B – THE UTILITY OF VLANS

A VLAN is logically similar to a regular LAN. It is a collection of networked hosts that receives multicast or unicast messages from any single member of the collection. Think of it as the group of hosts that can hear what one member broadcasts. In human terms, a LAN would be any group of people who are within shouting range of someone who is speaking. To understand the utility of a VLAN, it is first necessary to understand the definition of a LAN.

Consider the networked hosts in Figure 34. The network hosts Traffic Counters 1, 2, and 3 and a Video Receiver connect to the same hub. If Traffic Counter 1 broadcasts a frame, the hub repeats that broadcast to all hosts that connect to it. This broadcast carries over to the switch on the right of the diagram. The switch forwards the broadcast to both Traffic Counter 4 and the Video Transmitter. In the Figure 34 example, all of the network hosts share the same LAN. Disconnecting the link between the switch and the hub prevents hosts on the left from hearing and broadcasting to hosts on the right. Such a disconnection would create two separate LANS.

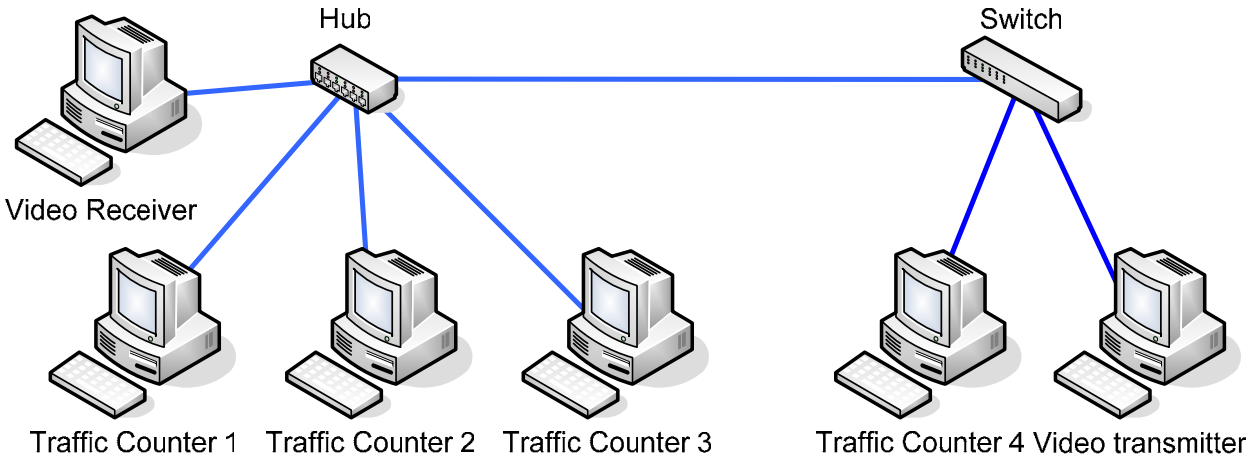

Figure 34. Two Example LANs.

Each individual host listens for broadcast frames. If the volume of broadcast frames exceeds the abilities of the host to process them, it may miss some frames purposely sent to it. As the number of hosts sharing the same LAN grows, this chance of overloading one host's network interface card increases. At some point, it becomes necessary to break hosts into smaller collections as separate LANS. Security concerns also require that LANs be broken up into task specific collections. It makes sense to separate all computers from a financial department and place them onto their own LAN for example. This division prevents

unauthorized frame snooping by a user not affiliated with finance. On the roadside, one desirable separation of hosts divides the low bandwidth hosts from the high bandwidth hosts.

Figure 35 illustrates a possible scenario where a VLAN would be useful. This example LAN has a high bandwidth Video Transmitter on the right and two Video Receivers on the right listening to the transmitter's broadcasts. Since the video transmitter is saturating the LAN with broadcast packets, Traffic Counters 1, 2, and 3 have difficulty talking to each other.

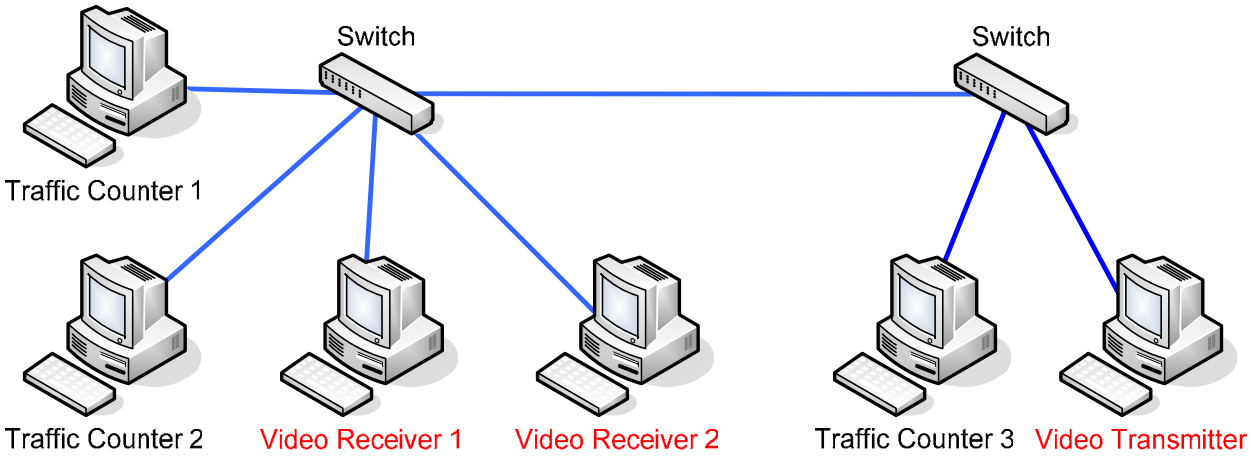

Figure 35. VLAN Application.

It would be helpful to separate the traffic counters into one LAN and place the video hosts into another LAN. This LAN separation is difficult to create because the hosts that require separation of data traffic physically connect to the same switch. VLANs overcome this physical connection limitation. High quality switches have options that allow the creation of LANs on a port-by-port basis. A VLAN is a LAN broken up on a port-by-port basis across potentially many switches. If the switches in Figure 35 support VLAN functionality, a network administrator programs the ports on the left switch that connect to the video receivers to a software label such as "VLAN 1." Likewise, the network administrator programs ports connecting Traffic Counters 1 and 2 into a second VLAN labeled "VLAN 2." On the right switch, the Video Transmitter's connecting port sets to VLAN 1 and the Traffic Counter 3's connecting port sets to VLAN2. Both right and left switches would then communicate to each other and link up similarly named VLANs. With these two VLANs linked between switches, broadcasts coming from the Video Transmitter only distribute to the Video Receiver hosts, which share the same VLAN as the Video Transmitter. The three traffic-counting hosts are now able to communicate on the much

quieter VLAN 2. With the two VLANs defined, the six hosts logically appear to connect as shown in Figure 36.

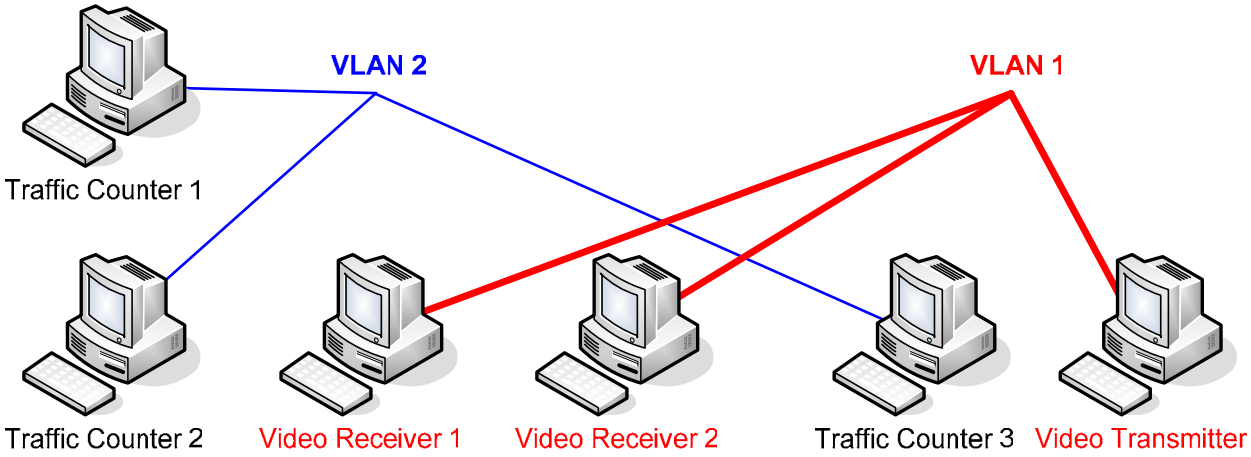

Figure 36. Logical VLAN Collections.

APPENDIX C – ROBUST NETWORKS AND THE SPANNING TREE ALGORITHM

Many networks depend on physical interconnections. Networks that link together roadside equipment have a good chance of being severed by the occasional backhoe operator. To ensure a network performs despite the failure of a physical connection, it is desirable to introduce redundancy into the network. Unfortunately, a network with redundant physical connections can cause havoc on the network traffic flowing through it. The spanning tree algorithm addresses these issues to allow the coexistence of redundant physical connections and smooth network data flow.

To understand the need for the spanning tree algorithm, consider the example network shown in Figure 37. In this example, there are two traffic cabinets, one at Main Street and the other at 1st Street. Inside these traffic cabinets are Ethernet hosts that collect video feeds coming from cameras. These video feeds are encapsulated into Ethernet frames and sent on to a video collection host located in the city office. Both traffic cabinets and the city office share the same LAN via interconnected switches. This network is redundantly connected by two links between each switch to provide an alternate path should one link fail. Without the spanning tree algorithm programmed into these three switches, this network will fail in an exponential manner.

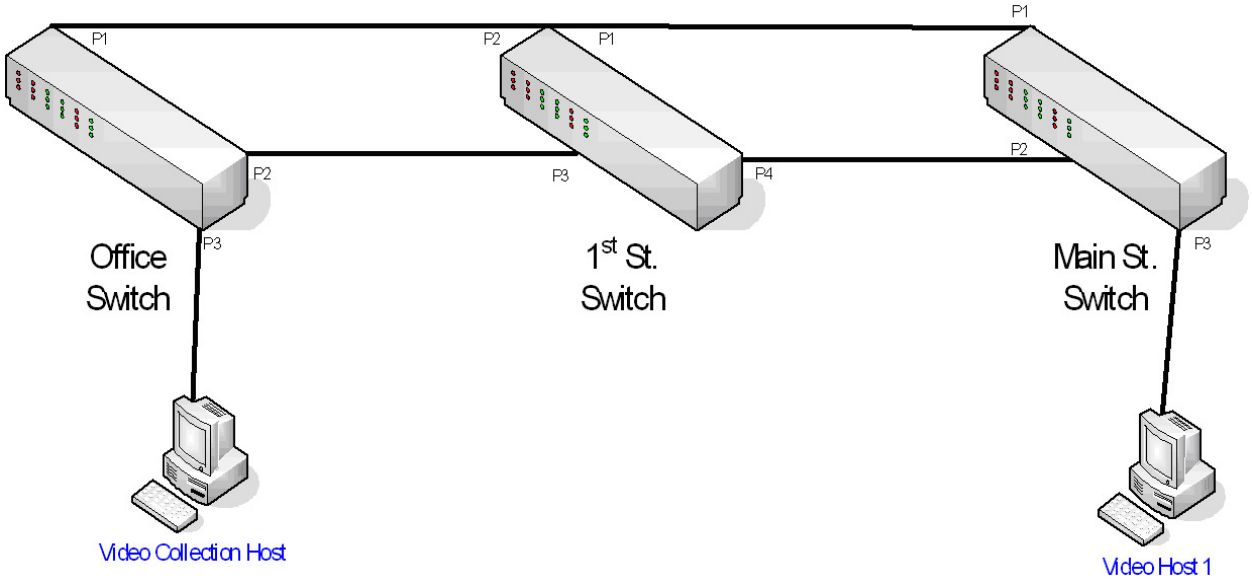

Figure 37. Layer 2 Loop Example 1a.

To illustrate the exponential problem caused by redundant data paths on a LAN, consider Video Host 1 located in the Main Street cabinet (Figure 38). Video Host 1 transmits an Ethernet frame to port 3 (P3) on the Main Street switch. The destination for the transmitted frame is the Video Collection Host in the Office. The Main Street switch notices that the destination host of this frame is not directly connected to the Main Street switch, so it forwards the Ethernet frame out all of its ports except the port from which the frame originated. This results in the transmission of duplicate frames out of ports P1 and P2 of the Main Street switch. The original frame now has doubled.

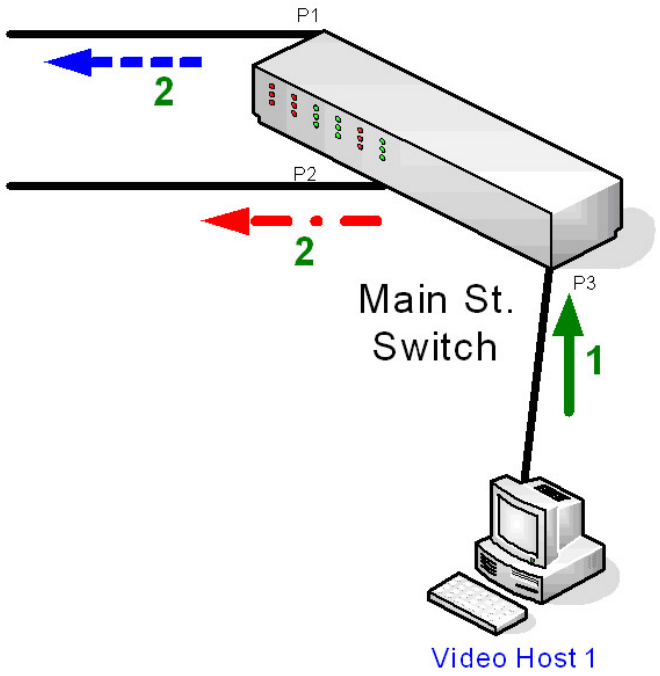

Figure 38. Layer 2 Loop Example 1b.

The two duplicate frames arrive at ports P1 and P4 of the 1st Street switch. The 1st Street switch looks at the two frames and determines that their destination host is not directly connected. Because of this, the 1st Street switch forwards the frames out all of its ports except for the port where an individual frame originated (Figure 39). This sends two new frames out P2 and two duplicate frames out P3. One frame from P4 is sent back via P1, and the inbound frame from P1 is sent back via P4. Now the original transmitted frame has six duplicates moving across the network, two of which are headed *back* to the originating switch! The beginning of a bandwidth-consuming frame loop has started.

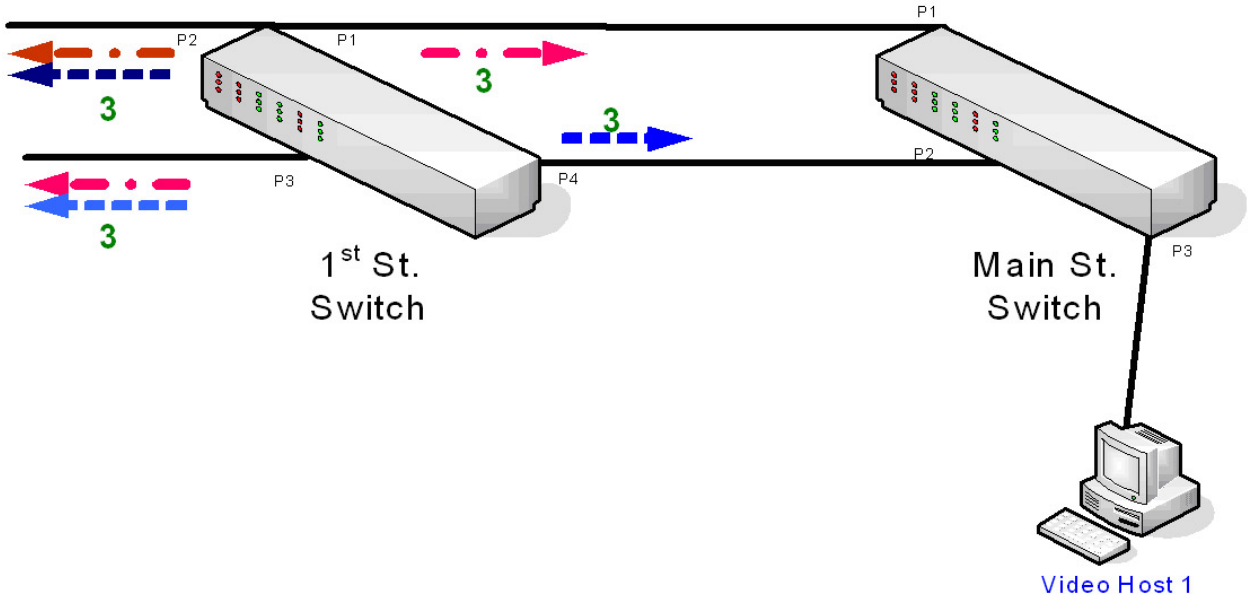

Figure 39. Layer 2 Loop Example 1c.

The switch in the office receives four duplicated frames from its P1 and P2 ports and echoes the duplicates to the destination Video Collection Host (Figure 40). It also forwards the opposite frame pairs back to the sending switch on 1st Street. Meanwhile, on the other end of the network the Main Street switch receives two useless frames and echoes them back to the 1st Street switch. Now the 1st Street switch is about to receive six redundant frames that it will echo back as 24 frames! This frame loop will grow until the switches are no longer able to keep up with the network traffic.

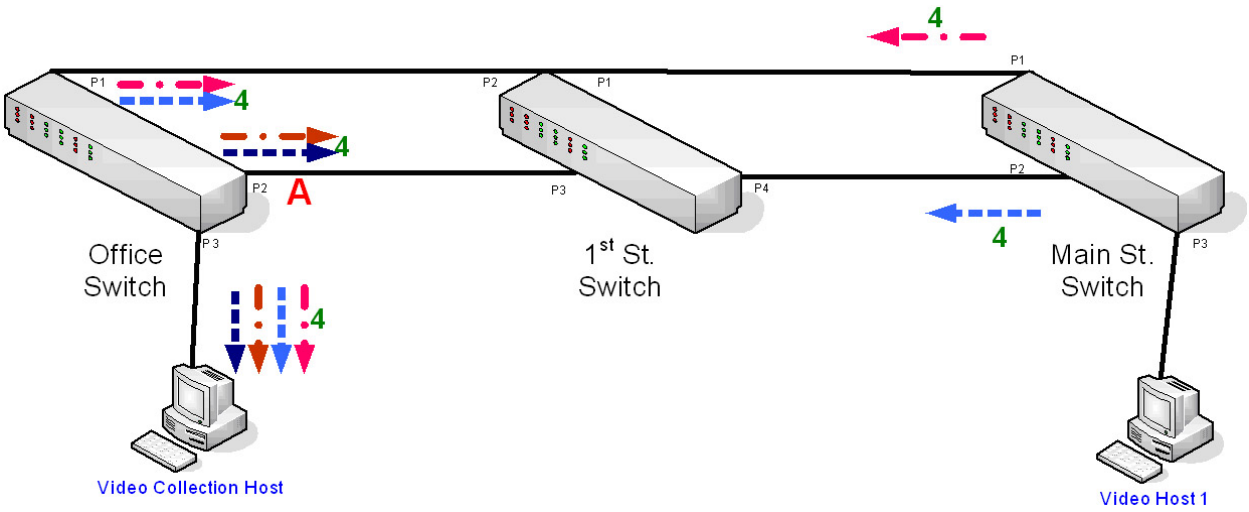

Figure 40. Layer 2 Loop Example 1d.

The main problem that the spanning tree algorithm addresses is that it prevents Ethernet frames from flowing over redundant paths between network devices. Switches running a spanning tree algorithm will talk to each other and in software shut down all but one path between each other (Figure 41). This one path between devices resolves the previously described frame loop situation. Even with a software cutoff in place, the switches continue to send administrative messages to each other through the inactive links. This allows them to quickly activate an alternate link should the primary link fail.

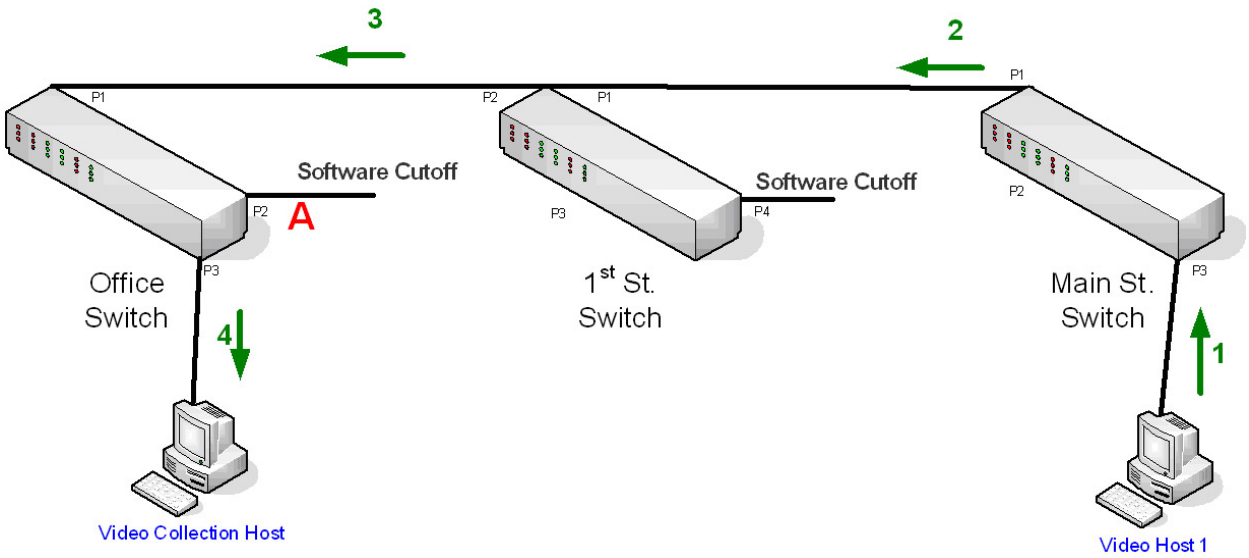

Figure 41. Layer 2 Loop, with Spanning Tree Correction.

To better illustrate the spanning tree algorithm, consider the example shown in Figure 42. Here a collection of switches (represented as circles) connects redundant links of various bandwidths (which could be different physical media such as fiber, Cat5, or coax).

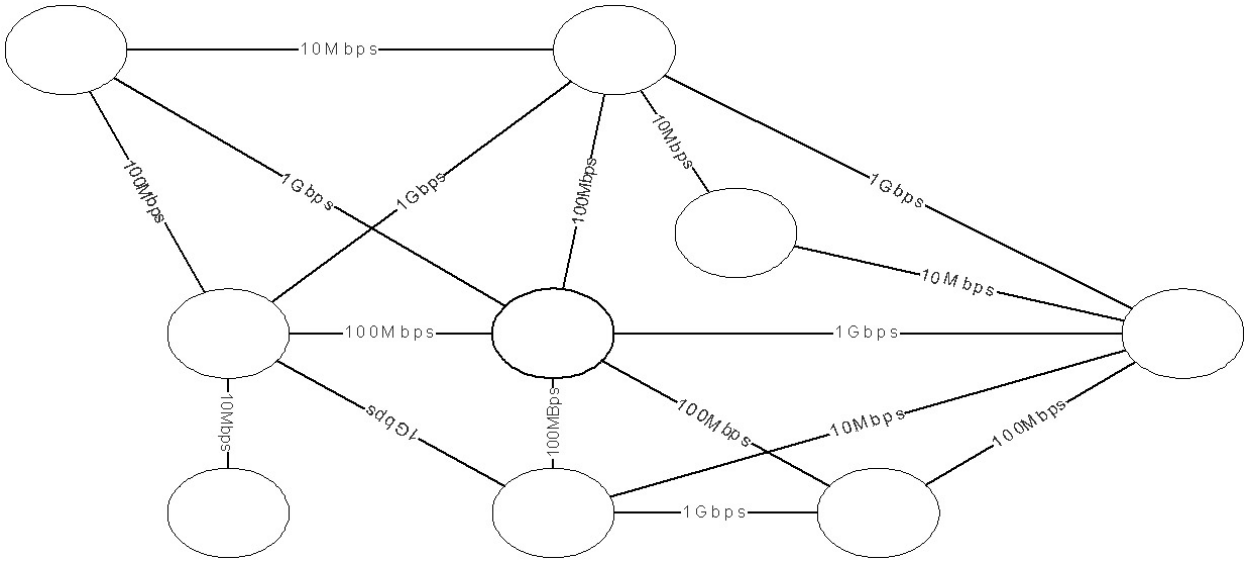

Figure 42. Spanning Tree Example 2a.

The first thing the spanning tree algorithm will do is assign all of the switches a root identification number (ID). The root IDs can be assigned by the network administrator or performed automatically by the switches. The switch assigned the lowest root ID is considered the trunk of the spanning tree (Figure 43).

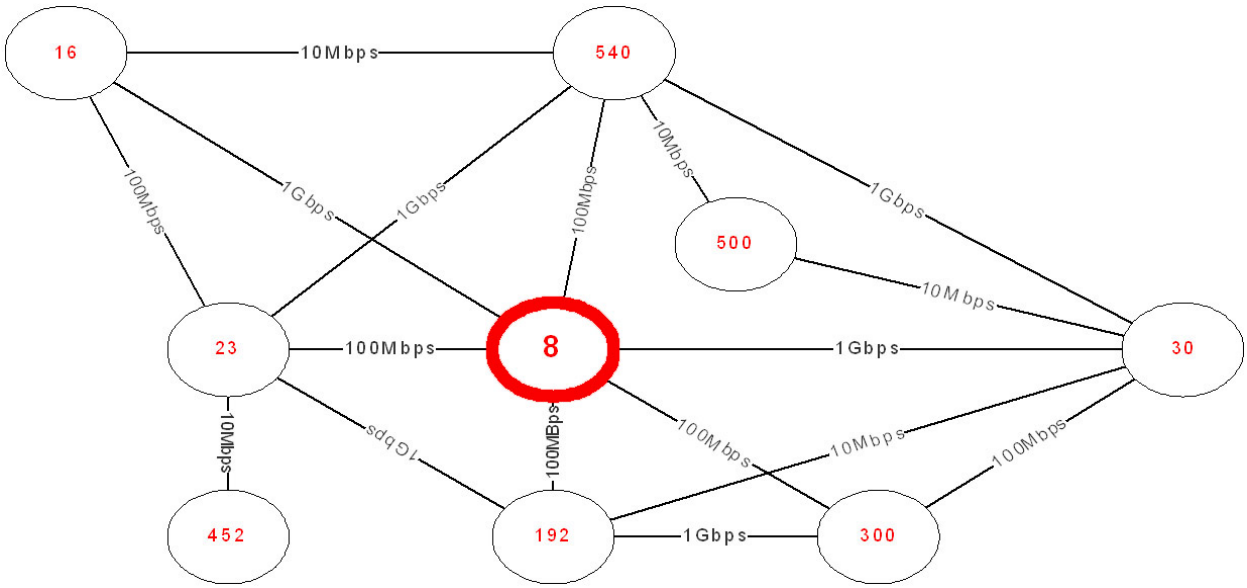

Figure 43. Spanning Tree Example 2b.

In this example, the switch with root ID of 8 is considered the main root. Think of it as the trunk of the spanning tree. From this switch, all other switches talk to each other and

calculate the highest bandwidth path to the root station. These calculations use the root ID numbers, the bandwidth of the links, and the total path cost from the target switch to the main root switch. For example, the switch with root ID of 500 near the middle top right of Figure 43 can connect to the main root switch 8 via a 10 Mbps connection to switch 540 and then a 100 Mbps link to switch 8. This path goes through a slow 10 Mbps connection followed by a medium bandwidth 100 Mbps connection. An alternate path to 8 from 30 goes via switch 30. This path travels through a 10 Mbps connection followed by a much faster 1 Gbps connection. For switch 500, the path to 8 is routed through switch 30 and the slower path through 540 is blocked by switch 500's software. Once all of the best paths are calculated, the switches route Ethernet frames only over those paths, and all other redundant paths are software blocked and remain ready to route in the event that a primary path fails. Figure 44 shows the best paths for this example network as solid and the software-blocked paths as dashed.

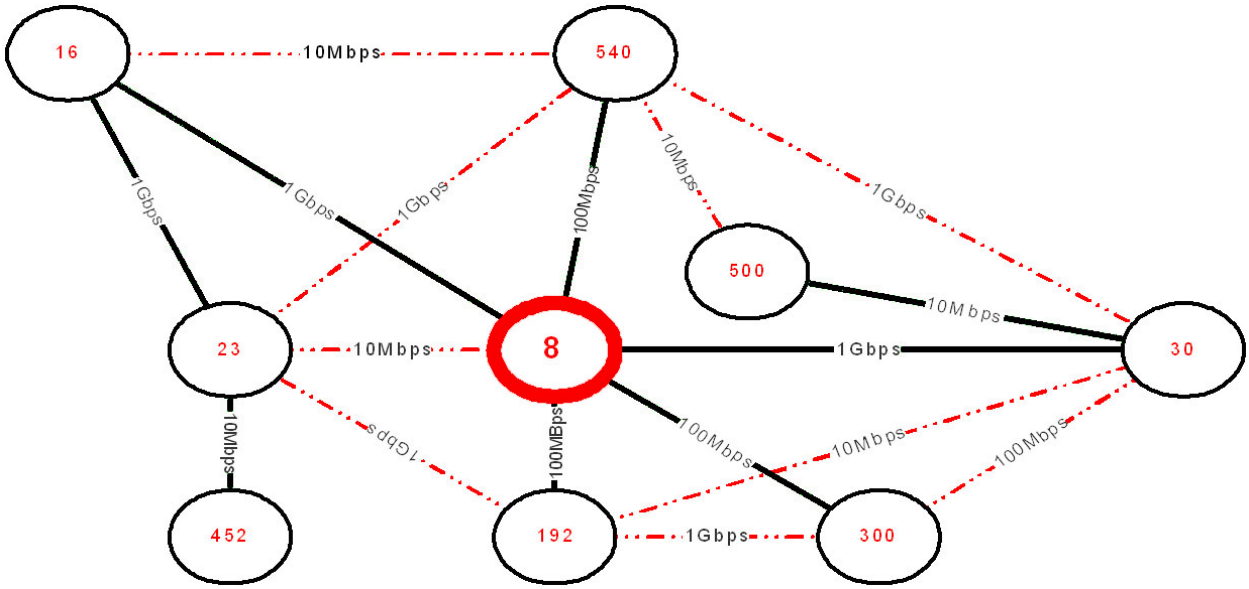

Figure 44. Spanning Tree Example 2c.

The network shown in Figure 44 is rearranged in Figure 45 to show the tree structure created by the spanning tree algorithm. With the primary paths set, there is only one path to get from any two switches. For example, there is only one active path between switch 452 and switch 540. This path starts at 452 then moves to 23, then 16, 8, and finally to 540. With the spanning tree active, there are no active network loops.

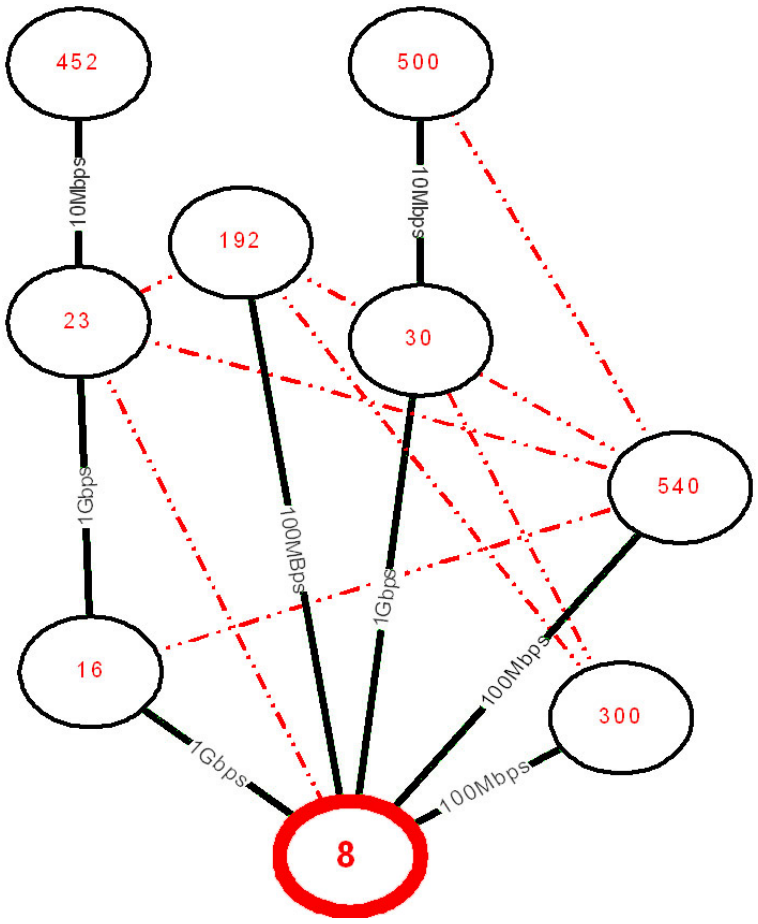

Figure 45. Spanning Tree Example 2d.

The spanning tree algorithm operates with five rules if viewed from an individual switch's perspective.

Rule 1: Determine the best path to the root switch (or the switch seen as having the smallest bridge ID) by comparing all administrative information provided by the neighboring switches. The port that has the best path is labeled the root port for the targeted switch.

Rule 2: Once the root port is determined, configure the port to forward all frames that do not have destination addresses directly connected to the target switch.

Rule 3: Any LAN segment extending from a switch that has been designated as a forwarding path for other switches, must be set to forward frames that do not have destination addresses directly connected to the target switch.

Rule 4: All connections directly linking hosts (such as a PC) must remain in a frame-forwarding mode.

Rule 5: All other ports on the target switch must be set into a blocking mode that prevents Ethernet frames from traveling over them, with the exception of switch-to-switch administrative messages.

Figure 46 displays a common LAN configuration that utilizes a redundant loop. In this example, there are Ethernet switches placed in traffic cabinets that service four city blocks. With this ring, there are two physical paths between any two locations. This setup provides a backup path that maintains connectivity during times when one segment of the LAN fails.

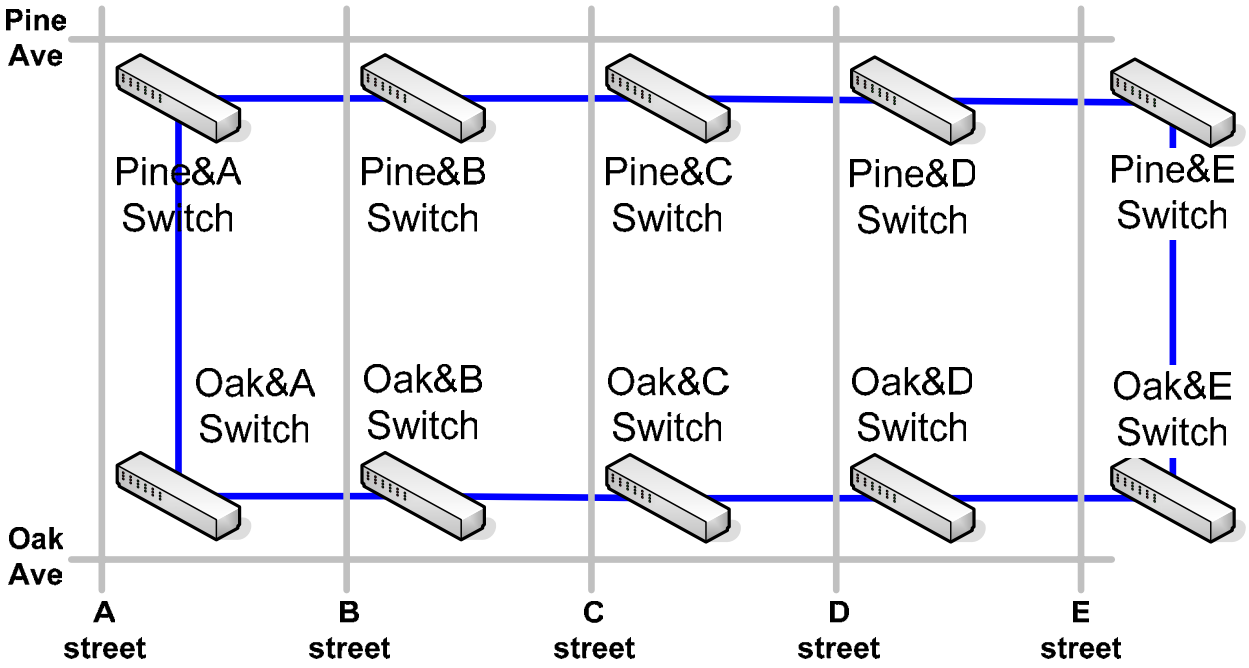

Figure 46. Ring LAN Example.

This LAN is represented logically in Figure 47. If the root switch is selected to be at Pine and C Street, the spanning tree algorithm will calculate one-to-one paths and close down one redundant segment located at the opposite end of the ring (assuming all ring segments have the same bandwidth rating). With one path blocked in software, the spanning tree algorithm creates a single path between all switches. If one of the active links fails, it will be detected and the topology will be recalculated for the network utilizing the previously blocked segment.

One final thing to know about the spanning tree algorithm is the time it takes to recalculate topologies. Every switch on the LAN that is using the regular spanning tree

algorithm will need 30 seconds to recalculate and adjust to the new topology. The LAN described in Figure 46 and Figure 47 could take up to 3 minutes to stabilize after a topology change. The 30 second per node recalculation time is restricted by an Institute of Electrical and Electronics Engineers (IEEE) standard rather than a hardware limitation. There are faster implementations of the spanning tree algorithm available in most managed switches. One such implementation is called rapid spanning tree. If the timely recovery of a failed network is important, there are many commercial options available like rapid spanning tree to address that issue. There are numerous protocols available to help a LAN meet desired performance requirements. Spanning tree is one of many solutions.

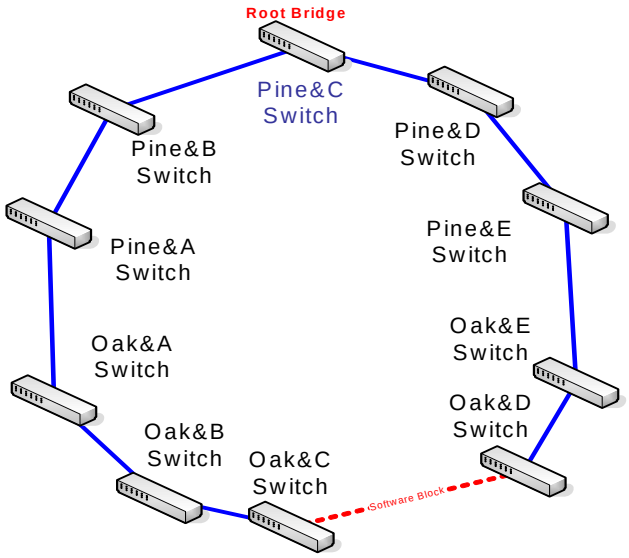

Figure 47. Ring LAN Logical Topology.

GLOSSARY OF TERMS

ATM

Asynchronous transfer mode is a network technology that transfers data in fixed-size packets or cells. ATM packets are smaller than the majority of other protocols. This smaller and constant size lets ATM devices transmit different dissimilar data such as video, voice, or computer data over the same shared network while limiting the dominance of any one data type. In other words, bandwidth heavy data like video won't flood out tiny data messages traveling over the same ATM network. ATM constructs a fixed path between receiver and transmitter. Data traveling from the transmitter to receiver all go via the same relay of interconnects. This is unlike TCP/IP, which allows packets to take multiple routes between transmitter and receiver. The advantage of ATM's same-path routing is it allows for easier network monitoring of data flow and makes for easier sequential delivery of packets. The disadvantage is ATM links have a harder time handling large bursts of data compared to TCP/IP.

Concentrator

(See hub.)

Bridge

A bridge takes two or more network segments and smartly allows each segment to send frames back and forth. The bridge will look at an incoming Ethernet frame from network segment A and determine if the frame's destination is also on segment A. If this is so, the frame is not repeated uselessly onto other connected network segments. If the bridge detects that the incoming Ethernet frame from network segment A has a destination in network segment B, it will forward that frame on to network segment B. Unlike a hub, a bridge only allows across Ethernet frames that have destinations that are not local to their originating network segment. This selective forwarding of Ethernet frames begins when the unit is activated. It starts listening to all Ethernet frames coming into its connected ports. As the Ethernet frames arrive, the bridge records the MAC addresses of the frames and constructs a table for MAC addresses for each port. When a new Ethernet frame arrives on a port, the bridge refers to these generated MAC tables to make a frame forwarding decision. Another important property of a bridge is its ability

to “break up collision domains” by buffering inbound and outbound data so that no two Ethernet frames are sent out of a port at the same time. This improves data flow through the connecting ports.

Central Office (CO)

The central office (CO) is what telephone companies call the buildings they own that contain the “other” ends of all the phone lines in their local neighborhoods. Central offices contain computers and machines that connect (or “switch”) phone lines together. When a phone call is made, the signal goes down the phone line to the central office, where the “switches” connect that line to the one dialed. This site also contains equipment that connects customers to long distance services and Internet service providers. xDSL lines running from a subscriber’s home connect at their serving central office *(1)*.

Collision Domain

A collision domain defines the set of ports between which a repeater will repeat a signal. Any time two hosts on a network segment try to transmit at the same time, they enter a collision scenario. Both hold the transmit line high for a little bit to inform all hosts sharing the same data channel. Next, both hosts stop and wait an individually random amount of time before attempting to resend their respective messages. If too many hosts share the same data channel, the chance of a collision and, thus, congested data flow is significantly increased. A bridge breaks up a collision domain into smaller domains by buffering inbound and outbound Ethernet frames. A hub does not have this ability. This means that a hub is not able to handle ports with multiple hosts as efficiently as a bridge. Always look for ways to minimize the number of hosts per collision domain. Inserting bridges between multiple Ethernet ports is one way to accomplish this.

Ethernet

Ethernet is an IEEE standard network protocol that specifies how data are sent on and retrieved from a shared transmission medium such as unshielded twisted pair cable or coax cable. It provides the underlying transport vehicle used by several upper-level protocols such as TCP/IP.

Ethernet is a very common method of networking computers in a local area network using copper cabling. Ethernet will handle about 100,000,000 (100 Mbps) bits per second and can be used with almost any kind of computer.

Firewall

A firewall is software running on a networked device that actively scans and restricts non-authorized layer 2 Ethernet frames (computers not part of a specific VLAN), and/or layer 3 data packets (suspect information coming from the outside network).

FTP

File transfer protocol is a very common method of moving files between two Internet sites. FTP is a special way to login to another Internet site for the purposes of retrieving and/or sending files. There are many Internet sites that have established publicly accessible repositories of material that users can obtain using FTP, by logging in using the account name anonymous; thus, these sites are called anonymous FTP servers *(2)*.

Hub

A hub joins multiple Ethernet connections and blindly repeats any incoming data from one connection (called a port) to all other connections. In most cases, a hub will not look at the Ethernet frames to make switching decisions; it just repeats what it hears across all of its ports. A small minority of hubs check the Ethernet frames to see if the frames need forwarding through the hub's uplink connection. A few hubs check incoming Ethernet frames for data integrity. Because of this small subset, hubs operate primarily on the Open Systems Interconnection (OSI) layer 1 (physical layer) with a minority operating at layer 1 and layer 2 (datalink layer). Since a hub does little more than repeat what it hears, it is cheap to construct and sell. The disadvantage of a hub occurs when the hosts connecting to it are causing a large volume of Ethernet traffic. Since every transmit from each host repeats across all ports, there is an increased chance of data collisions as two hosts try to transmit at the same time. A hub is cheaper than a switch and works fine when connecting a handful of hosts that generate light Ethernet activity.

An example of where a cheaper hub is preferred over a more expensive switch might be a traffic cabinet with a computer and two 9600 baud, RS-422 to Ethernet transceivers. Since the

two transceivers are sending infrequent data, they and the computer are unlikely to have frequent data collisions.

A hub is a poor choice in a situation where the hosts connected to it are generating heavy Ethernet activity, such as a bank of video to Ethernet encoders all flooding the hub with outbound data. Here the heavy transmissions would collide often and produce congested data flow out of the hub.

IGMP

Internet group management protocol is used by switches and routers to manage multicast group memberships and multicast data flow on a network (see multicast). IGMP allows the switches and routers on the network to deliver multicast packets only to the hosts that request the multicast data rather than wasting network bandwidth delivering the multicast data to all hosts whether they need the data or not.

Internet/Intranet

The Internet is the vast collection of interconnected LANs that all use TCP/IP. An "intranet" is a private network inside a company or organization that uses the same kinds of software that you would find on the public Internet, but that is only for internal use. An intranet may be on the Internet or may simply be a network *(3)*.

LAN

LAN stands for local area network, which is a computer network that spans only a small area. Most LANs are confined to a single building, group of buildings, or perhaps just a few miles. A LAN can be connected to other LANs over any distance via telecom providers. A system of LANs connected in this way is called a wide area network. Devices communicate with each other to share resources and information, such as disk storage and files, printers, and e-mail *(4)*.

Latency

Broadly, latency is a term used to describe a delay in transmitting data between two computers. During this time, the processes associated with the communication are hung up and

cannot continue. It is the time taken to transfer a packet of information from one point to another, and it is often cited as one of the reasons for the sluggish performance of Internet connections (*5*).

Multicast

A multicast is a targeted broadcast of frames or packets to a group of hosts that have requested to receive the data. A multicasting host addresses its packets to a special IP address. The IP address represents a channel of information rather than a specific host address. Individual network hosts that wish to receive this multicast data will listen for and process packets on the network that have that multicast address in their destination field. This is similar in concept to tuning a radio to a specific channel to receive the data presented on that channel.

Protocols

Protocols are a set of formal rules describing how to transmit data, especially across a network. Low-level protocols define the electrical and physical standards to be observed, bit- and byte-ordering, and the transmission and error detection and correction of the bit stream. High-level protocols deal with the data formatting, including the syntax of messages, the terminal-to-computer dialog, character sets, and sequencing of messages (*6*).

Repeater

A repeater ties two separate network segments together and repeats traffic from one cable segment onto the other cable segment. This simple device operates on the OSI layer 1 (physical layer). It does not concern itself with what information is passing through it. It simply repeats the electrical signals it receives from the connected cable segments.

A repeater might be used if Ethernet cable running under the street from a traffic cabinet needed to be connected to another run of cable to further down the road. A repeater might be inserted between the two cable runs to amplify the signals being sent and extend the range of the cable run. A common rule of thumb when using Ethernet is to have no more than four repeaters joining a segment of cable. Extending a cable an excessive distance causes problems with data collision detection.

Response Time

Response time is defined as the elapsed time between the end of an inquiry or demand on a computer system and the beginning of the response; for example, the length of time between an indication of the end of an inquiry and the display of the first character of the response at a user terminal. For response time monitoring, it is the time from the activation of a transaction until a response is received. To make it even simpler; it is how long after you click on a button before you see something happen (*7*).

SMTP

Simple mail transport protocol is a subset of TCP/IP. Its primary purpose is the delivery of e-mail between e-mail servers. Many network devices use SMTP to send alert flags to users when something goes wrong. For example, if a networked power switch from a remote site supports SMTP, it might be programmed to send the operator an e-mail every time the power is cycled at the remote site. Since e-mail can easily route over protected networks, SMTP notifications require no "inside the LAN" tools to keep the operator informed about network dynamics.

SNMP

Simple network management protocol is a layer 7 or application-layer protocol that allows the exchange of network management messages between connected hosts. Network administrators troubleshoot, analyze, and monitor network traffic by using SNMP information.

Switch

The usage of the word "switch" is ambiguous. One definition of a switch attaches the label to anything that moves data. With this definition, a layer 3 router or a layer 2 bridge are both considered switches.

A secondary usage of the word "switch" concerns layer 2 bridges with added functionality. Like a bridge, a switch looks at the destination address of incoming Ethernet frames and decides to which connected ports it should forward the inbound frames. Unlike a bridge, a switch has multiple forwarding modes for inbound Ethernet frames. One mode waits for an entire Ethernet frame to enter the port input buffer, performs a checksum, and then if all

things are valid it makes a forwarding decision and moves the frame to its appropriate destination. This is called "store-and-forward" switching. A second mode that a switch can operate in is "cut-through" mode. Here the switch waits only during the first part of the Ethernet frame that contains destination information before making a forwarding decision. This allows for faster routing of the Ethernet frames at the expense of no error checking.

T1

Trunk level 1 lines are digital transmission links able to operate at 1.544 Mbps. T1 is a common standard for digital transmission in the United States. Traditionally, T1s were used for digitized voice data. They arrived to the subscriber via four twisted pair cables, one pair for transmits and one pair for reception. Today the same T1 performance and specifications are normally delivered via fiber optic systems from the local exchange carrier to the subscriber. T1s can be broken up into separate channels that take up portions of the 1.544 Mbps bandwidth (called Partial T1s), or the entire channel can be used to provide 1.536 Mbps of throughput. The traditional delivery of a T1 is split into 24 channels, each able to carry 64 Kbps per channel. These channels can be configured to carry a mix of data, voice, or digital voice traffic. Data T1s are able to route various protocols such as frame relay, ATM, or Internet traffic. T1 lines are not dependent on a specific media. T1 channels can be carried by optical, wireless, satellite, or wired media. The included table compares the T1 channel to similar channel designations. Table 5 displays some common trunk line types.

Table 5. Common Trunk Line Types.

Type	Bandwidth	# Voice Channels
T1 (DS1)	1.544 Mbps	24 VC
T1 C	2.152 Mbps	48 VC
T2 (DS2)	6.312 Mbps	96 VC
T3 (DS3)	44.736 Mbps	672 VC
T4 (DS4)	274.176 Mbps	4032 VC

TCP

Transmission control protocol is a subset of TCP/IP. TCP allows for data transmissions that require data integrity. Data designated for transmission is broken down into manageable chunks called packets. As these packets are sent over to the receiver, TCP has features built into it to help manage and guarantee that the packets are reconstructed successfully back into one solid block of data. During transmission, packets are occasionally dropped, or duplicated. TCP tracks this. The receiver and transmitter talk to each other to compensate for transmission errors. The transmitter sequentially numbers each outgoing data packet. If the receiver detects a gap in received packets, it requests a second transmission for the missing packet. Once reconstructed, TCP performs error checking to verify the transmitted information was not changed or accidentally corrupted.

Industry jargon labels TCP as a "connection oriented" protocol. TCP is useful when it is important to have a set of data transfer correctly and in sequence. One example might be a daily download of traffic information from a remote site. The data for that day must transfer correctly and be free of transmission errors if it is to be trusted. TCP would be a solid choice for this kind of transfer. Another application well suited for TCP would be the updating of a remote station's software. If the update file corrupts due to transmission, it adversely affects the operation of the remote device. TCP's error checking and sequence tracking reduces the chances of file corruption. The disadvantage of TCP is that it increases the overhead bytes in the packets to track data delivery. This means that TCP provides less throughput than other simpler protocols.

TCP/IP

Transmission control protocol/Internet protocol is a combined set of protocols that performs the transfer of data between two computers. TCP monitors and ensures correct transfer of data. IP receives the data from TCP, breaks it up into packets, and ships it off to a network within the Internet. This is the suite of protocols that defines the Internet. TCP and UDP are OSI layer 4 protocols, and IP is an OSI layer 3 protocol *(3)*.

UDP

User datagram protocol is a subset of the TCP/IP protocol. UDP allows for data transmissions that do not have guaranteed delivery or packet acknowledgements. It does not use

any flow control or have any method for recovering dropped packets. It tosses the data down the pipe and leaves it up to the receiver's software to catch and organize the data into a reconstructed message. The benefit of this protocol is the lack of functionality allows for the information that packages the data to be very small. On a channel without loss, it requires less bandwidth to move data using UDP than other protocols like TCP.

UDP is a good fit for data in which the youngest message overrides any previous message. Consider traffic video. The observer may only care to see the most recent image. The observer does not care if a few frames are dropped in transit because the most recent image nullifies the content of previous messages. Another good fit for UDP data transport are status messages. An observer is interested in knowing if a railroad gate arm is up or down. The observer is interested only in the most recent status message. UDP allows for the reduction of unnecessary network packets. Industry jargon labels UDP as a "connectionless" protocol because of its lack of transmission error checking.

APPENDICES REFERENCES

1. *Glossary*, DSL Center for Small Business, http://dslcenter.netopia.com/netopia/glossary.html. Accessed August 25, 2004.

2. *Glossary of Internet Terms*, Matisse Enzer, http://www.matisse.net/files/glossary.html. Accessed August 25, 2004.

3. *Glossary of Internet & Web Jargon*, UC Berkeley – Teaching Library Internet Workshops, http://www.lib.berkeley.edu/TeachingLib/Guides/Internet/Glossary.html. Accessed August 25, 2004.

4. *Glossary*, Texas State Library and Archives Commission, http://www.tsl.state.tx.us/ld/pubs/compsecurity/glossary.html. Accessed August 25, 2004.

5. *Glossary*, Western Carolina University, http://online.wcu.edu/orient/Glossary.htm. Accessed August 25, 2004.

6. *Geek Speak Glossary*, Geekazoid and friends, http://www.geekazoid.com/geekspeak/. Accessed August 25, 2004.

7. *Glossary*, Glasgow Caledonian University, http://www.gcal.ac.uk/cit/helpdesk/useful_definitions.htm. Accessed August 25, 2004.

APPENDICES BIBLIOGRAPHY

1. *An Introduction to Computer Networking,* Kenneth C. Mansfield, Jr., and James L. Antonakos (Prentice Hall, 2002).

2. *Newton's Telecom Dictionary 16th edition*, Harry Newton (Publisher's Group West, 2000).

3. *Interconnections, Second Edition, Bridges, Routers, Switches, and Internetworking Protocols*, Radia Perlman (Addison Wesley Longman, 2000).