

VIRTUAL LOCAL AREA NETWORKING AND ITS IMPLEMENTATION
AT UNC-CHAPEL HILL

by

William F. Bobzien, IV

A Master's paper submitted to the faculty
of the School of Information and Library Science
of the University of North Carolina at Chapel Hill
in partial fulfillment of the requirements
for the degree of Master of Science in
Information Science.

April, 1999

Advisor

Virtual Local Area Networking and Its Implementation at UNC-Chapel Hill

Introduction

In 1996, the University of North Carolina at Chapel Hill (UNC-CH) completed a major upgrade to their computer networking infrastructure. Previously, a broadband coaxial-cable network transmitting at 5 megabits per second (Mbps) had connected the various local area networks (LANs) on campus. While this backbone implementation had met the needs of the University as outlined in the Academic Computing Advisory Committee report ten years earlier (Academic Computing Advisory Committee, 1986), improved technology and the growing demand for network connectivity on the UNC-CH campus had outstripped the capacity of this system (Gogan, 1997). The upgraded system replaced the coax-based broadband backbone with fiber distributed data interface (FDDI) rings transmitting at 100 Mbps. In addition, the campus networking authority, Academic Technology and Networks (ATN) Networking and Communications, redesigned the distribution network to a flat, switched topology, first at 10 Mbps and later at 100 Mbps in core segments. Networking and Communications selected Cabletron Systems, Inc. as their preferred switch vendor, in part because Cabletron's switches included an implementation of virtual local area network (VLAN) technology (Gogan, 1997).

VLAN technology is, in concept, relatively simple. It enables end stations (typically microcomputers, but also printers, file servers, etc.) that are located on multiple LANs to operate as if they were all connected to the same LAN (Passmore & Freeman, 1996). Thus, following a classic example from VLAN literature (Passmore & Freeman, 1996; UC Davis, 1997), workers in the same department who work in different physical locations can "see" and use a single, departmental network (a VLAN), even though their workstations are in actuality connected to different LANs. Another popular use for VLAN technology is to contain broadcast traffic that can waste network resources (Passmore & Freeman, 1996).

Although this paper will look at what VLANs are and what functionality they add to a network (which are primarily end-user concerns) its main concentration will be answering the following questions, which address the issues of concern to most information technology professionals:

How does a VLAN work?

Many discussions of VLANs stay at the "what-it-does" level. While this perspective is a useful one, it is important to understand how a VLAN implementation works, especially because there currently are no standards for VLAN implementation. Each method is proprietary to the vendor providing the solution. As the specific implementation described in the paper uses the Cabletron Systems, Inc. VLAN technology, that system will be described in some detail. However, it should be possible to generalize from the Cabletron and other VLAN models a set of general methods or techniques that should be useful in evaluating and comparing the various proprietary standards.

What lies ahead for this technology?

VLAN technology is still fairly new, and its future is uncertain. Will VLAN implementations become standardized and interoperable, or will they remain vendor-centric? Will they catch on in the networks of many enterprises, or will they become specialized solutions to a narrow range of problems?

The UNC-CH Implementation

In addition, this paper will look at the implementation of VLAN technology at UNC-CH in particular, beyond illustrating general VLAN principles. It will focus on the following questions:

Why did ATN choose to implement VLAN technology?

This paper will examine the choice by ATN to implement VLAN technology in general and Cabletron VLANs in particular. What factor(s) or community need(s), either existing or projected, led ATN to VLANs?

Has this implementation worked as planned?

Looking back on the reasons for implementing VLANs, is ATN as an organization pleased or at least satisfied with the performance of their VLAN implementation? This paper will examine whether or not the system is successfully addressing the needs of the UNC-CH networking community, and whether the projected uses for VLANs (and projected needs for them) came to pass as planned. Have VLANs at UNC-CH lived up to their promise?

What plans does ATN have for VLANs at UNC-CH in the future?

Given the almost three years' experience the organization now has with VLANs, are there new ideas or applications for this technology that ATN sees in the future?

Literature Review

Since VLAN technology is relatively new, and is different from vendor to vendor, it is not surprising that there is sparse mention of the technology in the literature.

However, there is more extensive literature available on networking in general, the problems that VLANs are intended to solve, the networking issues, such as network security, that they are intended to address, and the network applications that VLANs can enhance. The different vendors of VLAN systems have also published information, in the form of white papers and other material regarding VLAN technology in general and their solutions in particular.

Additional literature includes background information on computing and networking at UNC-Chapel Hill. The UNC-CH literature will provide a comparison between both the old and new networking technology as well as between old and new roles and policies for the ATN department at UNC.

General Networking

A complete discussion of networking is beyond the scope of this paper. However, Derfler and Freed (1996) usefully define many of the terms used in discussing networks. A *local area network (LAN)* is "a group of computers typically connected by no more than 1,000 feet of cable, which interoperate and allow people to share resources." A *network interface card* (or *LAN adapter*) is the device which packages data for transmission and acts as "a gatekeeper to control access to the shared

network cable.” Network interface cards break data streams into *packets*, which are reassembled at the destination. *Bridges* segment LANs or join LANs together; they act to control traffic by learning the “station address” of each machine on the networks in question, and only send a packet across the bridge if the destination of the packet is a station on the other side. *Routers* function similarly to bridges, but look at the network address of packets and use *routing algorithms* to send the packet to its destination efficiently.

Henry and De Libero (1996) describe the use of *switching* to divide the network into smaller segments. Switching helps to reduce then number of nodes trying to use the same network segment, resulting in lower congestion on each segment. In switched hubs or bridges, each node can have its own network segment, and therefore have access to all of the network bandwidth of the segment. Switching bridges can look deep into a packet and use protocol information and the like to provide a level of filtering and prioritization (Henry and De Libero, 1996).

Networking Problems and Issues

Virtual local area networks address and attempt to solve many of the issues and problems facing network administrators, particularly on large, enterprise-wide networks. Some common issues include network utilization, particularly *collisions* and *broadcasts*, and network security. In addition, administrators want to reduce the amount of time and resources required to perform “moves, adds, and changes” to the workstations on a network; such activities often take up a disproportionate amount of an administrator’s time and resources.

Network Utilization: Collisions and Broadcasts

Of particular interest to network administrators is the area of network utilization.

Network utilization describes the percentage of available network resources that are being used by end stations on the network. Martin, Chapman, and Leben (1994) and Tittel and Robbins (1994) provide a great deal of information on general networking theory, including the issue of network utilization. The most common type of network, Ethernet, allows any station to transmit information on the network as long as no other station is currently transmitting. However, it is possible for two or more stations to simultaneously “sense” that the network is clear and transmit at the same time, causing a *collision*. While Ethernet and other network protocols include methods for dealing with collisions, the larger the network (i.e. the more users it supports), the higher the frequency of collisions (Martin et al., 1994). As network activity increases, the frequency of collisions can severely degrade network speeds, to the point that the network may seem to have stopped working.

Comer (1995) and Chappell and Hakes (1994) describe a feature of local area networks that is related to the problem of collision frequency and its impact on network utilization: the propagation of “useless” network traffic. All signals from a station on a given network are sent to all other stations on the network, regardless of whether they are intended for a station or whether that station can even interpret those signals (Comer, 1995). The designers of Ethernet had the foresight to place the destination address at the beginning of each Ethernet packet (Comer, 1995), and thus the network interface on a particular workstation can rapidly determine whether or not a packet is addressed to it. Packets addressed to other stations can be examined and discarded with minimal use of system resources. However, the Ethernet protocol itself (Comer, 1995) and several higher-level protocols, such as

NetWare's IPX/SPX (Chappell & Hakes, 1994) utilize packets that are designed to be received *and processed* by all interfaces on a network. These packets have a special "broadcast address" instead of the destination address of a single station. When a workstation's network interface receives such a packet, it does not discard the packet based on its destination address; it examines it further to determine what action should be taken. If the interface "speaks" the protocol for which the packet is used, it takes action on the packet's contents; otherwise, the packet is discarded. Determining whether or not a broadcast packet should be discarded requires that the receiver look many bytes deeper into the packet, with a correspondingly greater use of CPU cycles.

Roesse (1998) discusses the particular problems associated with a "flat," switched network. Unlike large-scale networks consisting of subnetworks connected through a series of routers, a flat network is essentially one large broadcast domain. While this does have some advantages over traditional, routed networks, namely higher-speed connections between segments, lower cost of networking equipment, and lower administrative overhead, flat networks do have disadvantages as compared to routed networks. According to Roesse, connecting switches as routers are connected, with multiple possible paths from one point to another, can lead to "loops" in the network, wherein broadcast packets propagate infinitely, creating "broadcast storms" that can severely degrade network performance.

Network Security

Baker (1995) discusses a broad range of topics related to network security. He provides a good summary of the network security problem. "Good" networks should operate smoothly with other networks, be transparently to users, provide remote

access, and maintain peak performance. On the other hand, “secure” networks protect confidential information, keep network performance reliable, and emphasize data integrity. The two dimensions are often at odds (Baker, 1995). Most of what Baker terms “network security” would be more precisely called “server security.” He is more concerned with securing machines on a network than the network itself. Such a focus is appropriate, since common sense tells us that the targets of most malicious attacks are end stations and the data that reside in them, rather than the network itself. However, network abuses (and misuses) do occur, and in any event the means of accessing a server for purposes of breaching security is often a network (Baker, 1995). Network hardware, such as switches and routers, can implement some kinds of security, such as routing traffic in such a way that it travels by the most direct path, thereby minimizing the chance of interception. They can also implement security-oriented functions such as authentication and encryption (Baker, 1995).

VLAN Definitions and Capabilities

Most literature on VLANs available today comes from vendors who are supplying VLAN technologies. As mentioned earlier, no fully qualified standards exist for defining VLAN implementation; thus definitions are often different from vendor to vendor. One third-party source for VLAN information is the UC-Davis Network 21 initiative (1998). It defines much of the terminology involved in discussing VLANs, and includes a discussion of the uses of VLANs, especially with regard to an academic network.

Henry and De Libero (1996) define a VLAN as the construction of logical LANs across a switched network using “virtual circuits or connections (p. 75).” They further defines virtual circuits as

. . .the pathways created between two devices communicating with each other in a switched network or communications environment. These circuits are active only for the duration of the originating data packet. Even though an exclusive connection is established between two devices, it’s only temporary and is closed, or taken down, when the communications session is completed (p. 75).

The only VLAN-related standard currently under development comes from the Institute of Electrical and Electronics Engineers (IEEE). Their standard, "IEEE P802.1Q, IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Networks" (Institute of Electrical and Electronics Engineers, 1998), describes enhancements to the 802.x LAN/MAN standards for packet structure. The standard (from here on abbreviated to "802.1Q") complements the 802.1p standard for inter-bridge/switch communication, which includes the "Spanning Tree" algorithm used to eliminate network loops and broadcast storms. The packet structure of the major IEEE-defined network architectures (Ethernet, Token Ring, etc.) are redefined by 802.1Q to include "tags" that further describe the contents of the packet (IEEE, 1998.) The 802.1Q standard does not, however, define the actual content of these tags under the current draft; rather, it simply "makes room" for the tags in the existing packet structure.

A relatively unbiased overview of VLAN technology comes from Passmore and Freeman (1998), writing a white paper for 3Com, Inc. VLANs, they say, “represent an alternative solution to routers for broadcast containment, since VLANs allow switches to also contain broadcast traffic (p. 2).” While many enterprises have used switches to segment their networks, standard switches do not stop broadcast traffic.

VLAN technology allows broadcast containment without the high cost and speed penalty of routers. Passmore and White also discuss the typical reasons enterprises do not readily adopt VLAN technology:

- They are proprietary solutions, which are “anathema” (p. 2) to the networking market, which emphasizes open systems and interconnectivity
- VLANs add additional cost, both visible and hidden, to the administration of a network
- VLANs can impede full-speed access to centralized servers.

Passmore and White divide VLANs into four categories, based on the means by which they assign stations to a given VLAN: port grouping, MAC-layer grouping, network-layer grouping, and multicast grouping.

Cisco Systems (*Virtual LAN communications*, 1996) views VLAN technology (at least, their version of it) as providing flexibility in organization and greater segmentation of an enterprise's network. Cisco concentrates on port grouping, in which the port to which a user connects her or his workstation is grouped together with the ports of other users in her or his workgroup. Thus members of the same workgroup (the example in the text is the Accounting department) can work in different locations throughout the organizations, be it different floors, offices, buildings, or even campuses, and still connect to each other as if on the same physical network.

Finally, Cabletron Systems (1998), in a series of white papers and technical documentation, and Roese and Knapp (1997) describe Cabletron's proprietary VLAN system, SecureFast Virtual Networking. SecureFast implements a tagging system similar to that proposed by the IEEE's 802.1Q standard, but with some enhancements, such as utilizing both the source and the destination address in determining packet routing (Roese & Knapp, 1997).

Methodology

This project is a combination of a descriptive and a case study. As such, the research methodology concentrated primarily on fact gathering and interpretation rather than experimentation and data analysis. Background information was taken from the literature. This information including theory on the design and implementation of local area and wide area networks (WANs), concentrating on large, flat switched networks, common problems in the implementation of such networks, and a description of VLAN solutions in general and the Cabletron Systems, Inc. VLAN technologies in particular. This information formed the basis for the first portion of the paper. Information for the second portion, dealing with the UNC-CH experience with VLANs, came from two sources: ATN documentation and interviews with ATN personnel who were involved in the design and implementation of VLAN technology at UNC-CH.

ATN Documentation

Available ATN documentation regarding the design of the new network was examined to provide information on the design process. This documentation illuminated many of the factors leading to the implementation of VLAN technology in the design of the new network. Of particular interest was the decision to implement the Cabletron Systems VLAN technology over competing systems. Facts unearthed by analyzing the documentation also informed the design of interview questions.

Documentation from the actual implementation phase of the new network, and of subsequent maintenance, expansion, and problem resolution, will be crucial to an

accurate description of the UNC-CH experience with VLAN technology. A comparison of the design goals versus the actual, production system is a central point of the proposed project.

ATN Personnel Interviews

While an examination and interpretation of the available ATN documentation will provide a framework and fact base for analysis of the UNC-CH VLAN effort, first-hand accounts of the design, implementation, and maintenance process will complete the paper. Interviews concentrated on personnel involved in both the design process and the continued maintenance of the UNC network.

Networking Overview

As mentioned above, a comprehensive discussion of networking is beyond the scope of this paper. However, a brief review of key concepts, especially the particulars of Ethernet technology and the TCP/IP protocol, will be helpful in order to explain the workings and uses of VLANs.

The OSI Model

Modern networks are built using the Open Systems Interconnection (OSI) Model as a baseline. Developed by the ISO in 1977 (Zimmermann, 1980), the OSI Model presents a generic description of the functions of a telecommunications network. The OSI Model is divided into seven layers. Each layer provides a specific type of function, and is independent of the other layers. A network following the OSI Model would thus be able to change the way a given layer is implemented without having to change the functioning of the other layers. It is important to note, however, that

no networking technology currently used fully implements the OSI model; some combine the functions of two or more layers, while others perform the functions of one layer using more than one protocol. The Model is a baseline or ideal rather than a prescription. It enables network designers, engineers, vendors, and the like to have a common framework for describing and integrating various products.

The OSI Model is as follows:

Application Layer – 7
Presentation Layer – 6
Session Layer – 5
Transport – 4
Network – 3
Data Link – 2
Physical – 1

Figure 1: The OSI Model

Communications flows "down" the model, from the application at the top to the actual transmitter at the bottom, physical layer, and then back up the model at the receiving end. As a message travels down, each layer adds its required information to the beginning of the message. At the receiver, each layer reads its information, acts on it as required, then removes it before passing the remainder of the message to the next layer up. Messages in reality go down the model and then up. However, because each layer receives the message with its required information at the beginning (since the layers below have already removed theirs), the result is the same as if the message had been sent directly between the layers at each end.

(Zimmermann, 1980)

As described above, each layer has a different function. Layer 1 (Physical) actually transmits data between hosts. Layer 2 (data link) assembles the data into packets for transmission. Layer 3 (network) controls the routing of data between different

networks. Layer 4 (transport) insures that data gets from sender to receiver successfully. The functions of the other layers, while important, do not have a great deal of impact on VLANs for the purposes of this discussion. Protocols and devices which are used to implement each layer are called “layer *X*” protocols and devices. For example, a bridge uses the data from layer 2 (data link) to operate, and is often referred to as a “layer 2 device.” The Internet Protocol (IP), which implements layer 3 (network) on TCP/IP networks such as the Internet, is similarly referred to as a “layer 3 protocol.”

Ethernet

Ethernet is a typical networking architecture. Ethernet was developed by Metcalfe and others in 1976 as a means of connecting office microcomputers (Metcalfe & Boggs, 1976). It has since grown to become one of the most widely used networking architectures today. Ethernet is a combination of hardware (layer 1) and a specific data structure (layer 2). While different “flavors” of Ethernet exist, both in hardware (10Base5, 100BaseT, etc.) and at layer 2 (Ethernet II, Ethernet 802.2, etc.) they are largely the same. Ethernet hardware consists of shared media (the network cable) connected to network interface cards (NICs) at each station. Stations on the shared media “listen” for it to be clear, and transmit once it is. If multiple stations hear a clear line and start transmitting at the same time, a “collision” occurs; the stations involved wait a random amount of time and transmit again (Metcalfe & Boggs, 1976). The structures of Ethernet frames, while different for each version, have general similarities; the first part of the frame is the address of the station for which the packet is intended (destination address or DA), the second part is the address of the originating station (source address or SA). Depending on the version of Ethernet

in use, following the addresses will be a field indicating type (Ethernet II) or length followed by additional information (802.2/802.3). The basic structure is as follows:

DA	SA	Type/Length	Data (possibly including additional Ethernet information)
----	----	-------------	---

Figure 2: Ethernet Packet Structure

The destination address field of Ethernet allows for a broadcast address to be used; as discussed above this is a special address that tells all receiving stations to process the frame as if it had their own particular destination address.

At the hardware level, several different kinds of devices can be attached to an Ethernet cable. Simplest is the *repeater*, which takes a signal and retransmits it at full strength down the wire. Repeaters allow signals to travel farther on an Ethernet network. They repeat everything they “hear,” including broadcasts and collisions. Repeaters can have multiple ports; such a device is called a *hub* or *concentrator*. Repeaters are Layer 1 devices; they do not examine the frames they receive, but simply pass them on.

A more sophisticated device is a *bridge*, also known as a *switch*. A bridge receives and retransmits signals on an Ethernet network, but is more selective than a repeater (Martin et al., 1994). A bridge “learns” by listening to the network traffic, which NICs are on each side of it (using the layer 2 information in each frame.) If a device on one side of the bridge is transmitting to a device on the same side of the bridge, the bridge does not retransmit the signal. However, if the source address is on one side and the destination on the other, the bridge passes the signal from one side to the other (Tittel & Robbins, 1994). Figure 3 illustrates this concept. If station A sends a frame intended for station B, it is sent to both station B and the bridge. The frame is received on port 1. The bridge examines its bridging table and sees

that station B is also on port 1. Thus it knows that it does not need to transmit the packet out port 2. However, if station D sends a frame addressed to station A, the frame is sent to station C and the bridge (port 2). Since the bridging table indicates that station A is on port 1, but the frame came into port 2, the bridge retransmits the frame out port 1. The frame is then received by both station A and station B. Thus all stations on both sides receive the frame.

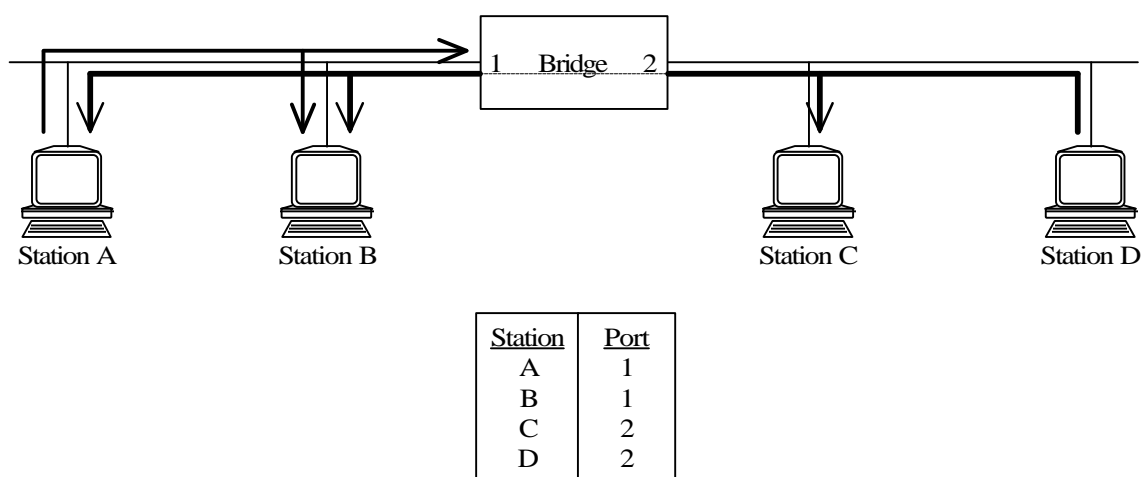


Figure 3: Bridging

If the bridge sees a DA that it has not “learned,” it will automatically retransmit the frame. Since a bridge “learns” from source addresses, and no station ever sends with a source address equal to the broadcast address, a bridge sends broadcast frames to all stations. However, since a collision creates a malformed packet that does not have a DA, a bridge will not retransmit a collision as a repeater would. Thus a bridge divides a network into separate *collision domains* while leaving the stations in the same *broadcast domain* (Martin et al., 1994). Bridges also often have filtering capability, so those frames of a certain type are not transmitted across the bridge. Alternatively, bridges can filter so that frames of a certain type or types are the *only* frames transmitted across the bridge. Filters can be useful in preventing packets

using a certain protocol from crossing the bridge, if the stations on the other side do not use that protocol and thus would not be able to process the packets (Derfler & Freed, 1996).

Switches use the same technology as bridges but with more ports than a typical bridge; using a switch, the network designer can allocate a port for each station, rather than a port on each side for two sets of stations (Martin et al., 1994). Each frame that comes into the bridge is therefore sent only to the station that needs it. Each station has its own segment of the network (Henry & De Libero, 1996). Thus in the second example above, traffic that crosses the bridge (station D to station A) is seen by all stations (A, B, C, D). Using a switch, such traffic is sent only to station A. One downside to switches is that broadcast frames are “flooded” out of every port, just as a bridge always retransmits broadcasts.

A third device often found on networks is known as a *router*. Routers are designed to work at level 3 of the OSI model, and as their name implies they route traffic from one network to another. On Ethernet networks, routers are often used to break a large network into subnetworks. They stop collisions like bridges and switches, but can also prevent broadcasts from travelling from one subnetwork to another.

TCP/IP

The Transmission Control Protocol/Internet Protocol (TCP/IP) is the Layer 4/Layer 3 protocol combination that is used to send packets across the Internet. It has also gained a following in LANs as well.

IP, the layer 3 protocol, provides for connectionless delivery of packets in a non-guaranteed manner. Packets are addressed to their destination using its IP address, a 4-byte sequence. They are then sent out over the LAN or Internet. IP makes no

provision for confirming the receipt of the packet by the destination station, nor does it provide a set route for packets to follow between the source and destination. Each packet sent to a given destination is routed across the intervening network segments independently of the others (Comer, 1995). Packets sent out in a given sequence may arrive at the destination at different times, and therefore in a different sequence. While this may not seem like the best way to communicate between hosts, the connectionless, freeform nature of IP makes it highly resilient; packets sent between two hosts are largely immune to disruptions in the networking segments between them. If several packets are sent along a given network segment, and then that segment goes down, subsequent packets are sent over a different route automatically. IP's nature reflects the Internet's origin as a communications system for the US military that could survive the loss of network segments due to nuclear bombardment.

TCP, the layer 4 protocol, provides the reliability and packet-sequencing facilities not found in IP. TCP adds information to the packet consisting of packet numbers (at the source) and acknowledgments (from the destination). This system of transmission and acknowledgement allows for confirmation of delivery and for the re-transmission of lost packets. TCP also allows hosts to exchange information regarding their sending and receiving capacities (flow control) and implements mechanisms to alter transmission rates in response to network limitations (congestion control). TCP is based on the concept of *data streams* that are sent out over *virtual circuit* connections, in contrast to IP, which sends out data as discrete packets into a distributed, connectionless "cloud" (Comer, 1995).

A host communicates with another host using TCP/IP in the following manner.

First, the host determines the IP address of the destination. This may be specified in

advance, or determined by translating the “name” of the host (e.g. “ruby.ils.unc.edu”) into an IP address using the Domain Name Service (DNS). The host examines the IP address to determine whether or not the destination station is on the same network. If it is, the host uses a secondary protocol, the Address Resolution Protocol (ARP), to find the associated hardware address. ARP consists of a broadcast across the network of the host's hardware address and the IP address of interest; the host with that IP address responds with its hardware address. It also stores the hardware address of the sending host in memory, so that it does not need to ARP for the sender's address in turn. The sending host then passes the IP packet to the layer 2 protocol (Ethernet, Token Ring, etc.). Otherwise, it uses ARP to find the hardware address of the network's router, then sends the IP packet to the router. The router then forwards the packet to the network of the destination station. Encapsulated within the IP packet is the TCP information. The destination system sends back IP packets, as well as acknowledgments of the TCP information, in a similar manner.

Other Network Protocols

While the Ethernet-TCP/IP combination is extremely popular in networking today, other layer 2 and higher protocols exist. One layer 2 technology that has implications for VLANs is ATM. Many enterprise networks also must deal with the presence of the IPX/SPX, AppleTalk, and NetBEUI protocols on their networks in addition to TCP/IP.

ATM

Asynchronous Transfer Mode (ATM) is a layer 2 protocol like Ethernet. Unlike Ethernet, ATM is a connection-oriented protocol. ATM communications work in a

similar fashion to a telephone call. The sending host builds a connection across the network to the receiver, setting aside a certain amount of bandwidth and dictating several parameters such as an acceptable packet-loss rate. Once the network indicates that it can support such a connection (analogous to the other party picking up their telephone handset), the sender begins transmitting data. When the transmission is complete, the connection is "torn down" to allow other hosts to use that portion of the network's resources. If the network cannot support the requested connection, the sender gets the equivalent of a "busy signal" and must either wait and try again or else request different connection parameters. ATM connections are virtual connections similar to those used by TCP. ATM has the capability to build *permanent virtual circuits* (PVCs) that are not torn down when data transmission ends. PVCs save on the overhead required repeatedly setting up and tearing down connections between two frequently communicating hosts.

IPX/SPX, AppleTalk, and NetBEUI

IPX/SPX, AppleTalk, and NetBEUI are layer 3 and 4 protocol suites used by Novell NetWare clients and servers, Apple Macintoshes and associated printer and file servers, and Microsoft Windows 95/98/NT clients and servers, respectively. While they operate in a similar fashion to TCP/IP, they rely more heavily on broadcasts than TCP/IP in order to operate. AppleTalk in particular is designed to require little or no configuration on the part of the user. To achieve this goal AppleTalk servers, clients, printers, and the like broadcast identifying information across the network frequently so that new devices coming on line can configure themselves or discover the locations of network resources automatically.

Networking Issues

While the Ethernet-based LAN has proven to be a robust and scalable network architecture, it does have limitations. As the number of users grows, the number of protocols in use on the network increases, and/or users become more geographically dispersed, a simple Ethernet LAN cannot be counted on to perform adequately. To illustrate, let us start with a small LAN and then observe what happens as it grows. The initial LAN configuration is five users and a central file server. These machines are all in a small office, and can be connected using a simple Ethernet LAN as shown. For the sake of argument, let us assume they all use the same protocols for layer 3 and 4.




Figure 4: Initial LAN configuration

All is well until another group of users wants to be connected to the same LAN. They are relatively far away from the original workgroup. Lengthening the LAN cable may or may not work, as it can only be a certain length. If they are too far away, a solution is to connect the two sets of users through a repeater, which regenerates the signals on the cable and allows a greater distance between groups.

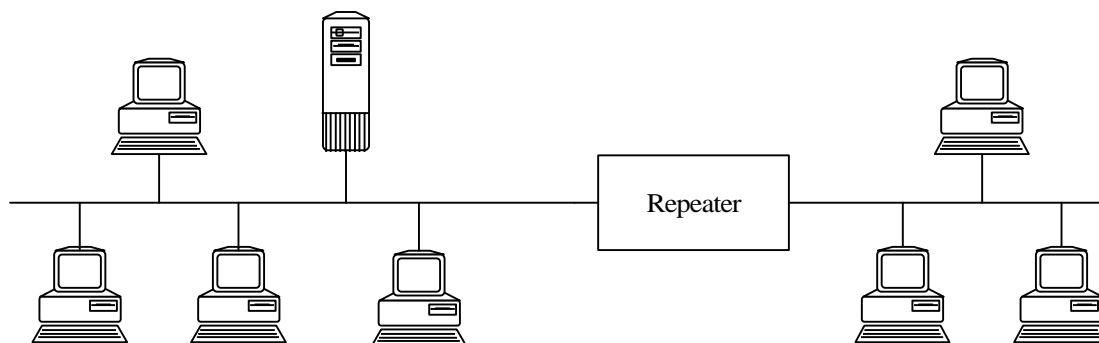


Figure 5: Connecting a Second Group Using a Repeater

Adding a repeater solves the distance problem, but the increased number of users on the network increases the potential for collisions. The network could become unusable at times because of the high collision rate. One solution: install a bridge to divide the network into two separate collision domains. Collisions on one side do not affect users on the other side.

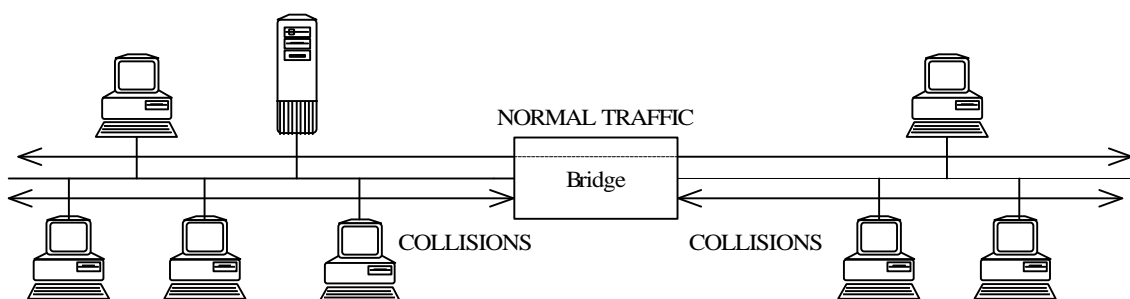


Figure 6: Installing a Bridge to Reduce Collisions

This solution works to reduce the number of collisions. Unfortunately, network performance is still degraded, since the users on the other side of the bridge frequently send traffic across the bridge to the server, which then uses up bandwidth and causes collisions on both segments. As an alternative to a bridge, a switch can be installed, with at least one port reserved for the server and one for each workgroup. Ideally, each station gets its own port on the switch, so that traffic intended for one station is never seen by the others.




Figure 7: Switching Gives Each Station Its Own Network Segment

This solution still leaves the problem of broadcasts. Broadcasts in and of themselves are not a problem; as discussed above they are necessary for the operation of most protocols (Martin et al., 1994). However, broadcasts used by one protocol are a waste of bandwidth and processing time for stations using other protocols. Each station must process broadcast packets to determine whether or not they should take action. Bridges and switches always pass broadcast packets to all attached stations. As an example, suppose we have the two workgroups as before, separated by a bridge with a port for each workgroup. Let us further suppose that two protocols are in use: TCP/IP and IPX/SPX (which is a much "chattier" or broadcast dependent protocol). If one group uses TCP/IP exclusively (the workgroup with the server) and the other group uses both TCP/IP and IPX/SPX, then the bridge can be programmed to only allow TCP/IP traffic to cross. However, if the second group uses IPX/SPX exclusively, then they will be unable to communicate with the server, even if it is configured to speak both IPX/SPX and TCP/IP. If both groups have a mixture of

TCP/IP and IPX/SPX machines, the NetWare machines would be unable to communicate across the bridge.

If a switch were in use, the situation would be even worse. If the switch filters IPX/SPX traffic, then none of the machines using IPX/SPX would be able to communicate, *even with each other*, since all traffic would go through the switch.

In such a case, the only solution may be to separate the networks using a router instead of a switch. This solution will stop the general broadcasts from going through, but can still enable broadcast-dependent protocols like IPX/SPX to operate across the router boundary, given the proper configuration. However, routers are much more expensive than switches, and are slower in their processing. Also, some protocols such as NetBEUI are not routable, so devices using these protocols cannot communicate across many routers.

The examples thus far have examined a relatively small network. The problems discussed are compounded greatly when the network in question encompasses dozens, hundreds, or even thousands of devices, as is more and more common in business and on college campuses. For maximum flexibility a network should be "flat," with all stations ideally on a single LAN, or at least separated only by bridges/switches (Roese, 1998). However, as the examples above demonstrate, such a network would most likely operate poorly, unless large amounts of resources, both money and time, are spent on its design and configuration. Large, flat networks have their own set of problems, as well; the algorithms used by switches to prevent "loops" in the network, such as Spanning Tree, can lead to inefficiencies in packet transmission. Spanning Tree requires that the switches on a flat network determine the one best loop-free route for packets to travel, rather than taking advantage of the multiple possible routes available. Without Spanning Tree, however, loops in the

network can lead to "broadcast storms," where a single broadcast is repeated over and over, wasting bandwidth at best and causing network failures at worst.

Ambiguous traffic on looped topologies can also cause individual switches to "misunderstand" the actual layout of the network (Roese, 1998). The solution in today's networks of separating groups of users by bridges, switches, and routers solves one set of problems but introduces another: users who for logical reasons, such as organizational affiliation, protocol use, or resource needs, should be on a single LAN, are instead separated into different networks. Roese (1998) identifies making "moves, adds, and changes" as the single biggest problem in such networks, since such operations are often quite costly in terms of the networking authority's resources, to say nothing of the lost productivity while the end user waits for the time-consuming process to be completed. Virtual Local Area Networks provide the capability to deploy a flat, switched network that allows users to be grouped together logically, while preventing the broadcast, collision, security, and other problems associated with large, flat networks.

Virtual Local Area Networks

A General Definition of VLANs

By Henry and De Libero's (1996) definition, a VLAN is a system that uses virtual circuits for transmission. This definition is too broad for our purposes, however. One of the foundation protocols of the Internet, the Transmission Control Protocol, utilizes virtual circuits (Comer, 1995). The TCP "stacks" of both the sender and receiver of data communicate with each other before any data are sent, negotiating a system for monitoring the transmission and reception of data packets sent using

the Internet Protocol (IP). The reliable transport mechanism thus established appears, at least to the communicating applications, to be a dedicated hardware connection between them. The X.25 Wide Area Network (WAN) protocol makes use of a similar feature (Comer, 1995). Thus by Henry and De Libero's definition any group of two or more hosts communicating over the Internet, or even the seldom-used X.25 protocol, are on a VLAN. While a case could be made in either instance that in some sense they are on a "virtual LAN," this kind of VLAN is not really applicable to the kinds of problems discussed above. The use of virtual circuits does not, therefore, distinguish a VLAN from any other communications using TCP/IP; limiting the definition to virtual-circuit use means that any time a user opens a telnet session (which uses TCP) to a remote host, he or she is creating a VLAN! A more narrow definition is required.

By moving our focus to the LAN rather than WANs or the Internet, we can craft a more appropriate definition for our exploration of VLANs. Recall that the TCP and IP protocols used by the Internet are layer 3 and 4 protocols, referring to their place in the OSI Reference Model. Hosts at each end of a communication, and the Internet routers in between, understand these protocols. These machines use TCP to set up virtual circuits for communication. However, while end stations on a LAN (which would be hosts on the Internet) still understand these protocols, the network hardware in between, notably switches, typically only understands layer 2 protocols. Bridges and switches originally used software to perform their tasks, which was slow; consequently, their designers relied on lower-level data, which is located at or close to the beginning of the packet and thus requires less time to access (Metcalfe & Boggs, 1976). Modern switches use dedicated embedded chips for processing, and are much faster, but still rely on layer 2 information. Most layer 2 protocols such as

Ethernet, the most popular, do not provide facilities for virtual circuits as the layer 4 TCP does, thus switches using these protocols cannot implement virtual circuits. The goal, therefore, should be to have a VLAN system that can create virtual connections at layer 2. To do so would enable the use of any other higher-layer protocols on the network, such as TCP/IP, IPX/SPX, etc. Three possible alternatives for reaching this goal exist. The first is to replace an Ethernet switched network with an ATM network to take advantage of ATM's virtual circuit technology. This option would be expensive. Also, ATM has some scalability issues and is challenging to integrate with some popular higher-level protocols like TCP/IP (Comer, 1995). A second option is to modify existing layer 2 protocols to allow switches to "route" at layer 2. One standard for Ethernet that can be used to accomplish this feat is the IEEE 802.1Q standard, which provides a structure for appending additional information to the standard Ethernet (802.x) protocol. While 802.1Q is touted as a VLAN standard, it actually retains the "connectionless" quality of traditional Ethernet. It does, however, enable groups of end stations to be grouped together based on layer 2 information. This solution would not require as large an infrastructure change as ATM, but would require that switches be programmed to recognize the 802.1Q "tags" (or other, proprietary tagging system) and act on them. However, a modification of the existing Ethernet structure means that, in theory, the network retains the scalability and clean integration with higher-level protocols of Ethernet. A third route is to modify the network switches to become "layer 3/layer 4 aware." Such switches take advantage of the faster processing speeds of modern switches to look deeper in the packet than traditional switches, and bridge or switch based on information other than just the layer 2 Ethernet information. This scheme has the advantage that only the switches on the network need to be involved in the

implementation, and use the standard Ethernet packets already in use. However, this approach does not allow for many desirable features of a VLAN, such as inter-switch communication.

Working from the commonality of these three methods, we can define VLANs more narrowly. A VLAN is a network made up of virtual circuits that are defined and maintained at the layer 2 level, using specialized switch hardware and possibly new layer 2 protocols or modified versions of existing protocols. The VLAN should handle broadcast traffic needed for end-station protocols in such a way that the broadcasts are received by all stations on a particular VLAN but are not sent to stations on other VLANs.

Passmore and White (1996) divide VLANs into four categories, based on the means by which they assign stations to a given VLAN: port grouping, MAC-layer grouping, network-layer grouping, and multicast grouping. A port-grouping VLAN maintains a list of all the switch ports on the network and the VLAN or VLANs to which each port "belongs." Any machine connected to a given port "sees" the other machines on that port's VLAN or VLANs. While this type of VLAN is relatively simple to implement (*Virtual LAN communications*, 1996), its utility is somewhat limited. The system's administrator must manually construct Port groupings, and if a user moves, the port grouping must be changed to accommodate the move. While this process is easier on a flat, port-grouping VLAN network than on a traditional routed one (Roese, 1998), it is still time-consuming. Port grouping also has some security problems. Depending on the implementation, a user with malicious intent and knowledge of the VLAN layout in an enterprise can simply attach a machine to a port and be "on" the VLAN. Should such a user be detected, her or his port could be manually disconnected from the VLAN, but again, this is time-consuming. Also, the

user would then be free to connect to another port and continue her or his activities. Finally, port grouping can lack the granularity required for broadcast control. If a user belongs to two VLANs, one for IP and one for IPX, for example, and the network utilizes port-grouped VLANs, then all traffic that user sends to her or his port is sent to both VLANs. If he or she sends an IPX broadcast, it is flooded to both the IPX *and* the IP VLANs. Some vendors "solve" this problem by allowing only one VLAN membership per port. While this may work in some settings, in others it can make a port-grouped VLAN a liability rather than an asset.

Some of these shortcomings are addressed by MAC-layer grouping, which groups machines (actually, network interface cards) rather than ports. Each NIC has an address that is unique. By grouping machines by MAC address rather than port, the MAC-layer VLAN enables users to move around at will; when they connect to the network at any given point traffic for "their" VLAN is routed to them, without any intervention by system administrators. MAC-layer VLANs do not suffer from the granularity problem described for port-grouped VLANs (Roese, 1998). MAC-layer grouping VLANs are better from a security standpoint as well, since the MAC address information can be used to deny access at any point automatically as easily as it allows such access for legitimate users. A malicious user can move from port to port as much as he or she wishes, but will be "locked out" no matter where he or she connects. However, MAC-layer VLANs must still be manually created initially (*Designing with Smart Network Services*, 1998).

VLANs can also be based on network-layer information. VLAN grouping is often automatic in this case. Any machine that the VLAN implementation detects "speaking" a given protocol is automatically added to that protocol-based VLAN. Thus all machines that use IP are on one VLAN, all those using NetWare's IPX are

on another, and so on. Network-layer VLANs make sure that broadcasts for a certain protocol go to the machines that can process it, and not to others.

Passmore and Freeman's last VLAN category is multicast-grouped VLANs.

Multicasting is similar to broadcasting in that a defined address is used that is not an actual end-station address. End-station NICs can be programmed to accept packets with that address as if the packets were using the actual NIC MAC address.

Multicasting is distinguished from broadcasting in that NICs always recognize broadcast packets, but can either accept or reject multicast packets depending on whether or not it is a member of a multicast "group." While multicasting is a more elegant way of distributing packets to multiple end stations than broadcasting or, even worse, multiple unicasts, all end stations must still analyze multicast packets. A multicast-grouping VLAN, by comparison, would only send the multicast to end stations that are in the multicast group. Other stations would not see the multicast packets at all. Multicast-grouping VLANs have a great deal of potential as more and more networks deliver multicast-based applications such as video-over-IP (Passmore & Freeman, 1996).

Of course, the above categories are not mutually exclusive. A particular VLAN implementation could implement two or more types of VLAN grouping. The different types complement each other; a system with port-grouping VLANs could benefit from the security features of MAC-layer VLANs and the automatic configuration of network-layer VLANs, for example.

802.1Q: A VLAN "Standard"

The IEEE has been the standards-defining body for much of the layer 2 communications technology used today, including Ethernet. Currently, the IEEE

has a draft standard for VLAN support on bridges and switches (IEEE, 1998). The 802.1Q standard, as it is known, describes a means of including information on VLAN membership in a standard Ethernet or other layer 2 protocol frame. Bridges and switches that are "802.1Q-aware" can make forwarding decisions based on the VLAN and other information stored in these tags. The bridges and switches in an 802.1Q-compliant network are also responsible for adding the appropriate tags to untagged frames they receive for forwarding. While 802.1Q-aware switches must replace older switches on a network in order to construct VLANs, the standard does not require that end stations be modified in order to be used on such a network. The untagged frames sent out by end stations would be tagged and VLAN-routed by the switches on the network. In contrast, to make full use of ATM in enterprise VLANs, ATM must be implemented all the way to the end stations. 802.1Q does not actually create virtual circuits. However, ATM's virtual-circuit technology is required for the protocol to work; 802.1Q provides similar performance without virtual circuits and thus without the associated end-station upgrades and training of networking staff on a new layer 2 technology. 802.1Q is not yet a fixed standard, and is currently subject to revisions. It also would define VLAN membership by port only. While this is a useful VLAN-grouping method, it is not as flexible as MAC- or protocol-based VLANs. However, the standard does not prevent the use of other, MAC- or protocol-based implementations running "alongside" 802.1Q, so long as these implementations do not interfere with the operation of 802.1Q (IEEE, 1998). One other limitation on the usefulness of 802.1Q, at least in its present form, is that the standard only addresses frame structure and the position of 802.1Q tags in that structure; it does not define the nature of the tags' content. Thus 802.1Q provides a

useful structure, but does not in and of itself provide an full-fledged VLAN implementation.

Cabletron, Inc. and VLAN Technology

Cabletron implements its proprietary VLAN solution, called SecureFast Virtual Networking, through a combination of the above methods (*IP Host Communication*, 1997). Switches on the network, called SecureFast Packet Switches (SFPS) communicate with each other and with a central server called a SecureFast Virtual Network Server (SFVNS). While the SFVNS acts somewhat like a router, keeping track of addresses on the network, packets are not sent to it for routing. Instead, the SFPSs on the network send requests for connections to the SFVNS, which then determines the switch route between hosts and sets up a virtual circuit on those switches. In a sense, SecureFast implements a modification of the normal Ethernet protocol to allow SFPSs to implement virtual connections (*IP Host Communication*, 1997). Machines on a SecureFast network have their default gateway set to their own IP address. This tells the NIC that it is able to communicate directly with all machines on the network, even though it actually cannot (Roese, 1998). Because the NIC assumes it is on a flat network, it ARPs for every address it wishes to contact. The SFPS hears the ARP and blocks it from travelling onto the network. The reception of the ARP request triggers a connection request to the SFVNS from the SFPS, which then sets up the required connection. The SFVNS resolves the IP address to a MAC address and sends that information to the originating SFPS, which replies to the originating host as if the destination host had responded to the ARP. It also “seeds” the destination host’s IP-to-MAC cache by unicasting an ARP packet to it with the originating host’s MAC address. Thus from the point of view of

the originating and destination host it appears that the ARP process is the same as if they were on the same network segment, without requiring an actual ARP broadcast (*IP Host Communication*, 1997). While routers can also perform such ARP broadcast containment by using "proxy ARP," they are slower and more expensive than SecureFast switches (*IP Host Communication*, 1997).

Cabletron switches implement a "tagging" system similar to the 802.1Q system suggested by the IEEE. Unlike 802.1Q, however, Cabletron's InterSwitch Messaging Protocol (ISMP) is connection-oriented. 802.1Q tags and forwarding logic are based on the characteristics of the destination address only; ISMP tags include information about both the source and destination addresses. Cabletron calls its ISMP virtual connections "flows" (Roese & Knapp, 1998). The distinction is that an 802.1Q implementation treats communications between two hosts as discrete. Given a stream of packets transmitted from station A to station B, 802.1Q switches treat the returning stream from B to A no differently than they treat the data transmitted from, say, station Y to station Z. A network using SFPSs speaking ISMP, on the other hand, "understands" and acts on the A-B and B-A streams as a logically linked, two-way communication. ISMP, as the name implies, allows switches to communicate with each other about the source and destination stations in a flow. Also, switches can send messages to each other and to the SFVNS regarding the topology of the network, much as the Open Shortest Path First (OSPF) protocol, after which ISMP was modeled, allows routers to send similar messages to each other (*SecureFast VLAN InterSwitch Message Protocol*, 1998).

The Future of VLAN Technology

It is difficult to project the future of VLANs, or say the least. While the technology does have great potential, the fractured nature of the current landscape in terms of multiple implementations with limited or no interoperability severely limits the growth rate of VLAN implementation. The single proposed standard currently available only partially addresses this problem. However, should the standard be improved or, as with other technologies, a single vendor or consortium of vendors develops an implementation which becomes the *de facto* standard, VLAN technology should find its way into most if not all enterprise networks. With their capability to resolve the problems that hamper the otherwise excellent flat network topology, VLANs can enhance network performance and security while reducing management overhead.

Implementing VLANs at UNC-CH

In order to gain insight into the design of UNC-CH's campus network and the role of VLAN technology in it, an employee of the Networking and Communications division of ATN was interviewed. While a single interview may seem like a small sample size, the Networking and Communications division is itself rather small, and preliminary investigation showed that other employees would have little to add to the information gained from the single interview. The employee in question was familiar with the design process leading up to the deployment of Cabletron switches and SecureFast VLANs, and also had a role in the ongoing maintenance of the network as well as input into future planning.

The redesign of the campus network at the University of North Carolina at Chapel Hill was in response to a variety of factors. The requirements of the faculty, staff, and students of the University were no longer adequately met by the existing network, which consisted of a 5 Mb/sec broadband coaxial-cable backbone connecting a heterogeneous mix of departmental and business-unit networks (Gogan, 1997). Also, the University's mandate to the ATN Networking and Communications department had expanded considerably since the installation of the broadband system. In 1986, when ATN (then called the Office of Information Technology, or OIT) had led a task force to analyze the computing needs of the University, OIT's area of responsibility regarding networking stopped at the boundary of the backbone. Individual units provided their own connectivity to the backbone, with design decisions such as shared bandwidth versus switched bandwidth made at the department level (Academic Computing Advisory Committee, 1986). Networking and Communications, however, had been charged with providing connectivity for University affiliates "from the wall plate to the Internet", meaning from the end station all the way through the campus distribution network to the Internet router that connected the University with the rest of the Internet community. While this mandate technically applies to all units of the university, for a variety of reasons some units still maintain their own networking authority. Regardless, the scope of ATN's responsibility expanded dramatically between 1986 and 1995, and a new network infrastructure was clearly in order.

ATN Design Philosophy

According to the ATN employee, the overall philosophy of Networking and Communications regarding campus networking is "switched rather than shared, if

possible." Under this philosophy, a UNC LAN would look more like that depicted in Figure 7 than that in Figure 4 or the intervening diagrams. Such a network gives each end station a dedicated network segment with full bandwidth access. ATN encourages the use of switched rather than bus- or hub-based networks for several reasons.

First, switched networks are more "scalable" than shared networks. It is easier for ATN to add bandwidth to switched networks than to shared networks, a sentiment that is echoed by Roesse (1998). The ability to add bandwidth easily will become important to UNC-CH as newer, high-bandwidth applications proliferate. One University affiliate should not see her or his network access degraded because a neighboring user is using a high-bandwidth multimedia application on the same, shared segment.

Second, shared networks often require extensive subnetting, with the accompanying need for routers. End stations on such networks require manual configuration; the associated overhead from "moves, adds, and changes" can be crippling, as Roesse (1998) points out. By moving to a switched network, configuration becomes much more "plug and play," to use the popular phrase. The end station does not require as much configuration; rather, the network in a sense configures itself around the end station. ATN viewed this capability as a vital feature in its redesign, having projected a great increase in the number of users with mobile computers, such as laptops. Manual configuration of such machines as they moved about the campus would be extremely difficult to execute efficiently, and configuration errors could create serious problems on the network.

The Role of VLANs

Because of the problems described above that are inherent in deploying large, flat networks, ATN made an implementation of virtual networking a requirement for vendors' bids when preparing the new design in 1995 (Gogan, 1997). While the product lines of several vendors were examined and evaluated, the University selected Cabletron Systems, Inc. and their SecureFast technology as being the best solution for ATN's purposes (Gogan, 1997). While several reasons for the selection were listed, Cabletron's VLAN implementation set it apart from competing vendors. SecureFast allowed port, MAC layer, protocol, and multicast grouping in its VLANs, while other vendors implemented port-based VLANs only, based on IEEE 802.1Q or similar, proprietary systems. While VLANs of that type would be generally useful, they would not address the mobile-computing needs of University affiliates.

However, the ATN employee admitted that having multiple membership mechanisms created its own set of policy issues. For example, if a user who is a member of a MAC-based VLAN connects her or his machine to a port that is a member of a separate, port-based VLAN, which takes precedence?

VLAN Use at UNC-CH

ATN has applied its VLAN technology to a variety of issues on the campus network.

Primarily, VLANs are used for broadcast containment, through protocol-based VLANs and Cabletron's Automatic Membership Registration (AMR) mechanism.

However, the technology is used in several other ways as well.

Port-based VLANs are used on campus to provide the dedicated connectivity and bandwidth demanded by some units. The University's Physical Plant unit, for example, maintains a network of monitoring devices in the various buildings on

campus. These devices are static, and thus are good candidates for a port-grouped VLANs. Without VLANs, the Physical Plant would have required separate, physical fiber links to each building from their central office, which is located north of the campus proper.

A novel application of the MAC-layer grouping mechanism is the "penalty box." ATN developed this VLAN to address the problem of "misbehaving" end stations on campus. If an end station violates University network or service-use policy, either as a result of incorrect configuration or maliciousness on the part of a user, it can be placed in this VLAN, either manually or in some cases automatically. The penalty box takes advantage of a VLAN feature that allows network administrators to define a VLAN as requiring a router to connect to other VLANs or the main network (*Designing with Smart Network Services*, 1998). The penalty box is defined in this manner, but no campus router is defined as part of the VLAN. Thus the "router required" restriction is never satisfied, and MAC-layer devices placed in this VLAN cannot communicate with any other host, unless that host (e.g. another policy violator) is also in the penalty box.

The AMR mechanism of SecureFast's VLAN manager also helps the network manage itself automatically. In the case of Novell NetWare servers, multiple Ethernet frame types can be specified. If a new server is brought up on the network, using a different frame type than the other servers, broadcasts from that server can cause clients to reconfigure themselves incorrectly or cause other problems. The AMR mechanism is sufficiently granular that it can create VLANs based not only on the protocol used but the frame type as well ("SecureFast Services Overview," 1998). Thus a misconfigured NetWare server would find itself in its own VLAN, unable to

communicate with the other NetWare servers and clients on campus and therefore unable to cause problems.

The Relationship between ATN and Cabletron Systems

Although the ATN employee stated that there has been "blood on the floor, some ours, some theirs" during the almost four-year relationship with Cabletron Systems, overall the relationship has been advantageous for both parties. ATN enjoys direct access to support by the vendor, described as a "strategic development relationship with the engineering group." Whereas most vendors require a customer, even one as large as UNC, to go through their sales department before a problem is relayed to engineering, Cabletron engineers can be reached with a telephone call and have become almost as familiar with the UNC network as ATN employees. Such reliance on a single vendor is not typical in networking environments, particularly in academic settings; it makes it difficult for UNC to change networking directions as quickly as may be required. On the whole, however, ATN feels the relationship is beneficial to the University. For their part, Cabletron can learn about the performance of SecureFast in a large-scale, real-world setting. Some of the forthcoming improvements in the VLAN Manager software will be the direct result of lessons learned through its use at UNC-CH. Cabletron has also learned about the scalability limitations of its system from its deployment at the University.

Future Trends at UNC-CH

ATN expects the use of VLANs to become more and more important as time goes by. They see more applications requiring "isolated" networks or dedicated bandwidth on the horizon, especially in multimedia. A collaborative effort is starting between the

University and Cabletron to set up "100 megabit to the desktop," meaning that end users would enjoy a 100BaseT 100 Mbps connection from their workstations, rather than the 10 Mbps connection they have now. ATN faces a quandary in this effort, in that 100BaseT SecureFast switches are still too expensive to deploy in sufficient numbers for this project. While 100BaseT hubs could be used, this would go against the ATN preference for switched over shared connections. One solution is the use of Cabletron's SmartStack switches, which are less expensive but do not support SecureFast ("SecureFast Services Overview," 1998). SmartStacks do support 802.1Q, and ATN along with Cabletron is investigating interoperability between the two technologies. In the future, data gathered at UNC may enable Cabletron to implement SecureFast using 802.1Q tagging, or at least more tightly integrate the two in its switches.

Conclusion

Network design has been and will be a process based on compromise. No one design offers a complete, perfect solution. While the deployment of flat, switched networks has definite advantages, such networks have problems not found in other architectures. The use of VLAN technology in a switched network can minimize or even eliminate these problems, allowing enterprises to make use of this topology without losing as much or more than they gain.

The University of North Carolina at Chapel Hill has made a commitment to using VLAN technology over a flat network fabric. While this decision has caused problems, and requires a much closer relationship with a single vendor than has traditionally been required in an academic IT environment, overall the system has done everything the University has wanted it to do. It has also increased

performance, improved network security and reduced the resources needed to manage the network.

References

Academic Computing Advisory Committee (1986). *A plan for campus-wide academic computing and digital communications*. Chapel Hill, NC: University of North Carolina

Baker, R. H. (1995). *Network security: How to plan for it and achieve it*. New York: McGraw-Hill, Inc.

Chappell, L. A. and Hakes, D. E. (1994). *Novell's guide to NetWare LAN analysis*. 2nd ed. Alameda, CA: SYBEX Inc.

Comer, D. E. (1995). *Internetworking with TCP/IP, volume I: Principles, protocols, and architecture*. 3rd ed. Upper Saddle River, NJ: Prentice Hall, Inc.

Derfler, F. J. and Freed, L. (1996). *How Networks Work*. 2nd ed. Emeryville, CA: Ziff-Davis Press, Inc.

Designing with Smart Networking Services: A Smart Network VLAN design guide from Systems Engineering. (1998). Rochester, NH: Cabletron Systems, Inc.

Gogan, J. (1997) *Why Cabletron as a preferred network vendor*. [online]. Available: <http://www.unc.edu/~gogan/whycatron.html>. (February 18, 1999)

Henry, P. D. and De Libero, G. *Strategic networking: From LAN and WAN to information superhighways*. London: International Thomson Computer Press

Institute of Electrical and Electronics Engineers (1998). *IEEE P8021.Q, IEEE standards for local and metropolitan area networks: Virtual bridged networks*. New York: Institute of Electrical and Electronics Engineers

IP host communication in bridged, routed and SecureFast virtual networks. 1997. [online]. Available: <http://www.cabletron.com/white-papers/cullerottoc/index.html>. (March 1, 1999)

Martin, J., Chapman, K. K., & Leben, J. (1994). *Local area networks: Architectures and implementations*. 2nd ed. New Jersey: P T R Prentice Hall, Inc.

Metcalfe, R. M. & Boggs, D. R. (1976) Ethernet: Distributed packet switching for local computer networks. *Communications of the ACM* 19, 395 - 404

Passmore, D. and Freeman, J. (1998). *The virtual LAN technology report*. [online]. Available: <http://www.3com.com/nsc/200374.html>. (October 2, 1998)

Roese, J. (1998) *Switched LANs: Implementation, operation, maintenance*. New York: McGraw-Hill, Inc.

References, continued

Roesse, J. and Knapp, E. (1997). *SecureFast: A comparative analysis of SecureFast and 802.1Q*. Rochester, NH: Cabletron Systems, Inc.

SecureFast services overview. (1998). *Cabletron Systems product marketing white papers*. Rochester, NH: Cabletron Systems, Inc.

SecureFast VLAN InterSwitch Message Protocol. (1998). Rochester, NH: Cabletron Systems, Inc.

Tittel, E. and Robbins, M. (1994). *Network design essentials*. Cambridge, MA: AP Professional

Virtual LAN communications. (1996). [online]. Available: <http://cio.cisco.com/warp/public/614/13.html>. (October 22, 1998)

VLAN information. *UCDavis Network 21*. (1998). [online]. Available: <http://net21.ucdavis.edu/newvlan.htm>. (October 21, 1998)

Zimmermann, H. (1980) OSI Reference Model-The ISO model of architecture for open systems interconnection. In C. Partridge (Ed.) *Innovations and Internetworking* (pp. 2-9). Norwood, MA: Artech House