NEAR EAST UNIVERSITY

Faculty of Engineering

Department of Computer Engineering

VIRTUAL LOCAL AREA NETWORK

Graduation Project
COM- 400

Student: Nael Attoun(20002109)

Supervisor: Assist. Prof. Dr. Firudin Muradov

Nicosia - 2005

ACKNOWLEDGEMENTS

Firstly I would like to present my special appreciation to my supervisor Prof. Dr. Firudin Muradov, without whom it is not possible for me to complete the project. His trust in my work and me and his priceless awareness for the project has made me do my work with full interest. His friendly behavior with me and his words of encouragement kept me doing my project.

Secondly I offer special thanks to my parents, who encouraged me in every field of life and try to help whenever I needed. They enhanced my confidence in myself to make me able to face every difficulty easily. I am also grateful to my mother whose prayers and my father whose words for me had made this day comes true. And because of them I am able to complete my work

I would also like to pay my special thanks to my all friends who helped me and encouraged me for doing my work. Their reluctance and friendly environment for me has helped me. I want to thank them as they contributed their time and provided very helpful suggestions to me.

ABSTRACT

A Local Area Network (LAN) was originally defined as a network of computers located within the same area. Today, Local Area Networks are single broadcast domains. This means that if a user broadcasts information on his/her LAN, the broadcast will be received by every other user on the LAN. Broadcasts are controlled within a LAN by using a router. The disadvantage of this method is that routers usually take more time to process incoming data compared to a bridge or a switch. Most importantly, the formation of broadcast domains depends on the physical connection of the devices in the network. Virtual Local Area Networks (VLAN's) were developed as an alternative solution to using routers to contain broadcast traffic.

In this project, we describe VLAN's, examine the difference between a LAN and a VLAN. And show the advantages of VLAN's introduce to a network. While bandwidth may be a reason big enough to go for switching, Virtual LAN (VLAN) Support may also be attractive. A VLAN is logical grouping of ports into workgroups. With VLAN support network managers can define workgroups independent of underlying network topology. VLANs are becoming popular because of the flexibility they offer.

TABLE OF CONTENETS

ACKNOWLEDGMENT i

ABSTRACT ii

TABLE OF CONTENTS iii

INTRODUCTION 1

CHAPTER 1: INTRODUCTION TO VIRTUAL LOCAL AREA NETWORK

1.1 Computer Networks 3

1.2 Needs of Networks 3

1.3 Goals of Computer Networks 5

1.4 Classification of Computer Networks 5

1.5 Local Area Networks 7

1.6 Ethernet Local Area Network 8

1.7 Major Components of LANs 10

1.8 Types of Local Area Networks 11

1.8.1 Peer-To-Peer 11

1.8.2 Client Server 11

1.9 Local Area Networks Connectivity Devices 11

1.9.1 Repeaters 11

1.9.2 Bridges 11

1.9.3 Routers 12

1.9.4 Brouters 12

1.9.5 Gateways 12

1.10 Local Area Networks in the work place and its advantages 12

1.11 Emerging Technology, Wireless Networks 13

CHAPTER 2: VIRTUAL LOCAL AREA NETWORK

2.1 Virtual Area Network 15

2.1.1 VLAN Benefits 20
2.1.1.1 Increased Performance 19
2.1.1.2 Improved Manageability 21
2.1.1.3 Network Tuning and Simplification of Software Configurations 22
2.1.1.4 Physical Topology Independence 23
2.1.1.5 Increased Security Options 23
2.1.1.6 Reduced Cost 24
2.1.2 VLAN Limitations 24
2.1.2.1 Broadcast Limitations 24
2.1.2.2 Device Limitations 24
2.1.2.3 Port Constraints 25
2.2 VLAN`s Working 25
2.2.1 Types Of VLAN`s 26
2.2.1.1 Layer 1 VLAN: Membership by Port 26
2.2.1.2 Layer 2 VLAN: Membership by MAC Address 27
2.2.1.3 Layer 3 VLAN: Membership by Protocol Type 28
2.2.1.4 Layer 4 VLAN: Membership by IP Subnet Address 28
2.2.1.5 Higher Layer VLAN`s 29
2.2.2 Types of Connections 29
2.2.2.1 Trunk Link 29
2.2.2.2 Access Link 30
2.2.2.3 Hybrid Link 30
2.2.3 Frame Processing 31
2.2.4 Filtering Database 31
2.2.4.1 Static Entries 31
2.2.4.2 Dynamic Entries 32
2.2.5 Tagging 34

CHAPTER 3: APPLICATION OF VIRTUAL AREA NETWORK

3.1 Membership by MAC Address 37
3.2 Layer 3-Based VLANs 38

3.3 Multicast Groups as VLAN`s 39
3.4 Combination VLAN Definitions 40
3.5 Automation of VLAN Configuration 40
3.6 Communicating VLAN Membership Information 41
3.7 Virtual LAN Management Protocol 42
3.8 VLAN Trunk Protocol 43
3.8.1 VTP Messages in Detail 43
3.8.2 Configuration Revision Number 44
3.8.3 Summary Advertisements 44
3.8.4 Subset Advertisements 46
3.9 Virtual LAN Security Best Practices 48
3.9.1 Basic Security 49
3.9.2 VLAN-Based Security 49
3.9.3 Control Plane 50
3.9.4 Precautions for the Use of VLAN 1 51
3.9.5 Secure environments of VLAN1 52
3.9.6 The Layer 2 Security 53
3.9.7 Attacks in a VLAN-Based Network 54
3.9.7.1 MAC Flooding Attack 54
3.9.7.2 802.1Q and ISL Tagging Attack 55
3.9.7.3 Double-Encapsulated 802.1Q/Nested VLAN Attack 56
3.9.7.4 ARP Attacks 58
3.9.7.5 Private VLAN Attack 59
3.9.7.6 Multicast Brute Force Attack 60
3.9.7.7 Spanning-Tree Attack 61
3.9.7.8 Random Frame Stress Attack 61

CHAPTER 4: Use of Virtual LANs (VLANs) in Networks

4.1 Use of virtual 63
4.2 VLANs in the DSL Access 63
4.3 VLANs in the Enterprise Network 64

4.4 Virtual LAN Configuration 65
4.5 Broadband Network Design 66
4.6 Requirements of the Service Providers Network 67
4.7 Internet service 68
4.8 Business Service 68
4.9 Speed 69
4.10 The Core Network 69
4.11 Access Components 70
4.12 VLAN Stacking 70
4.13 Net To Net Technologies Use of VLANs 72

CONCLUSION 73
REFERENCES 74

INTRODUCTION

Because of the fast development in the world of technology, and the big number of people -this is increasing every year and every day- who uses computers, and as those people are connected to each other all around the world by networks; known networks needs to be improved to suite this high speed developing area of technology, here we are talking about a side of this development in area networks; the Virtual Local Area Networks.

Virtual Area Networks are a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. Any switch port can belong to a VLAN, uni-cast broadcast, and multicast packets are forwarded and flooded only to stations in the VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router or bridge.

Virtual Local Area Networks can be viewed as a group of devices on different physical LAN segments which can communicate with each other as if they were all on the same physical LAN segment. Virtual Local Area Networks provide a number of benefits over the network, which we will discuss in the next section. In order to take advantage of the benefits of Virtual Local Area Networks, a different network topology is needed.

The objective of this project is to talk about the new generation of local area networks; the virtual local area networks the development and the uses of it, the project consists of introduction, four chapters and conclusion.

Chapter one describes computer networks, the need of it, the goals and the classifications of computer networks, the local area network and Ethernet local area network.

Chapter two presents the Virtual Area Network, the differences than the other local area networks the advantages and disadvantages.

Chapter three presents the applications of virtual local area network in life, how it works and its protocol.

Chapter four presents the way virtual local area networks effects networks and its configuration. The obtained results of Virtual Local Area Networks are analyzed finally; the conclusion section presents the important results obtained within the project.

1. INTRODUCTION TO VERTUAL LOCAL AREA NETWORK

1.1 Computer Networks

A network is a group of computers, printers, and other devices that are connected together with cables. Information travels over the cables, allowing network users to exchange documents & data with each other, print to the same printers, and generally share any hardware or software that is connected to the network. Each computer, printer, or other peripheral device that is connected to the network is called a node. Networks can have tens, thousands, or even millions of nodes. In the simplest terms, a network consists of two or more computers that are connected together to share information.

Principal components of a computer network:

- Computers (processing nodes or hosts)
- Data communication system (transmission media, communication processors, modems, routers, bridges, radio systems, satellites, switches, etc)

1.2 Needs of Networks

The concept of linking a large numbers of users to a single computer via remote terminal is developed at MIT in the late 50s and early 60s. In 1962, Paul Baran develops the idea of distributed, packet-switching networks. The first commercially available WAN of the Advances Research Project Agency APRANET in 1969. Bob Kahn and Vint Cerf develop the basic ideas of the Internet in 1973.

In early 1980s, when desktop computers began to proliferate in the business world, then intent of their designers was to create machines that would operate independently of each other. Desktop computers slowly became powerful when applications like spreadsheets, databases and word processors included. The market for desktop computers exploded, and dozens of hardware and software vendors joined in the fierce competition to exploit the open opportunity for vast profits. The competition spurred intense technological development, which led to increased power on the desktop and lower prices. Businesses soon discovered that information is useful only when it is

communicated between human beings. When large information being handled, it was impossible to pass along paper copies of information and ask each user to reenter it into their computer. Copying files onto floppy disks and passing them around was a little better, but still took too long, and was impractical when individuals were separated by great distances. And you could never know for sure that the copy you received on a floppy disk was the most current version of the information-the other person might have updated it on their computer after the floppy was made.

For all the speed and power of the desktop computing environment, it was sadly lacking in the most important element: communication among members of the business team. The obvious solution was to link the desktop computers together, and link the group to shared central repository of information. To solve this problem, Computer manufactures started to create additional components that users could attach to their desktop computers, which would allow them to share data among themselves and access centrally located sources of information. Unfortunately the early designs for these networks were slow and tended to breakdown at critical moments.

Still, the desktop computers continued to evolve. As it became more powerful, capable of accessing larger and larger amounts of information, communications between desktop computers became more and more reliable, and the idea of a Local Area Network (LAN) became practical reality for businesses. Today, computer networks, with all their promise and power, are more complicated and reliable than stand-alone machines. Figure 1.1 shows the network connectivity of the world.

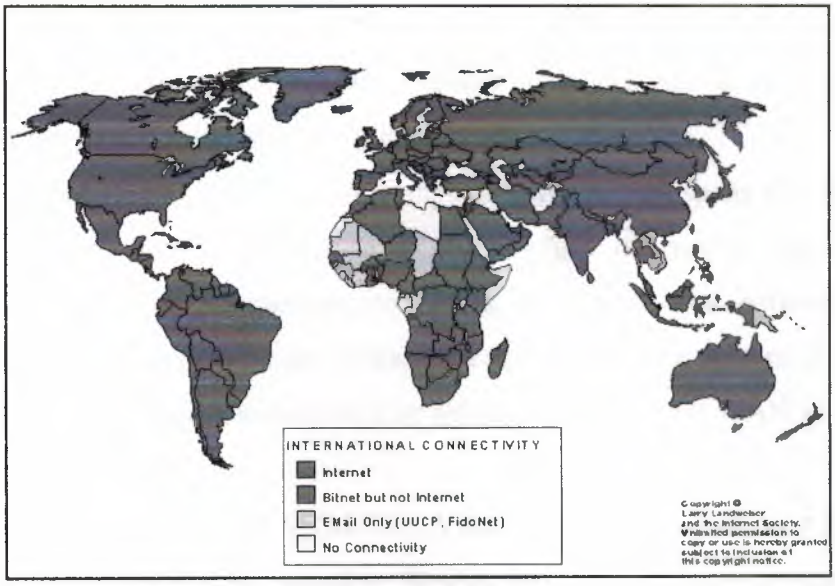

Figure 1.1 Computer Network Connectivity of the World

1.3 Goals of Computer Networks

.1. Resource sharing and accessing them independently of their location.

2. Providing a universal environment for transmission of all kinds of information: data, speech, video, etc.

3. Supporting high reliability of accessing resources.

4. Distribution of loads according to the requirements very fast main frames, minis, PCs, etc.

1.4 Classification of Computer Networks

Network Classification Like snowflakes, no two networks are ever alike. So, it helps to classify them by some general characteristics for discussion. A given network can be characterized by its:

- Size: The geographic size of the network
- Security and Access: Who can access the network? How is access controlled?
- Protocol: The rules of communication in use on it (ex. TCP/IP, NetBEUI, AppleTalk, etc.)
- Hardware: The types of physical links and hardware that connect the network

Computer experts generally classify computer network into following categories:

- Local Area Network (LAN): A computer network, with in a limited area, is known as local area network (e.g in the same building)
- Wide Area Network (WAN): A computer network that spans a relatively large geographical area. Typically, a WAN consists of two or more local-area networks (LANs). Computers connected to a wide-area network are often connected through public networks, such as the telephone system. They can also be connected through leased lines or satellites. The largest WAN in existence is the Internet.
- Metropolitan Area Network (MAN): A data network designed for a town or city. In terms of geographic breadth, MANs are larger than local-area networks (LANs), but smaller than wide-area networks (WANs). MANs are usually characterized by very high-speed connections using fiber optical cable or other digital media.
- Campus Area Network (CAN): The computer network within a limited geographic area is known as campus area network such as campus, military base etc.
- Home Area Network (HAN): A network contained within a user's home that connects a person's digital devices. It connects a person's digital devices, from multiple computers and their peripheral devices to telephones, VCRs, televisions, video games, home security systems, fax machines and other digital devices that are wired into the network.

Computer networks are used according to specified location and distance. In table 1.1 it is shown that which technology can be applied to the specific location and specific distance.

Table 1.1 Network Techonologies that Fit in Different Communication Spaces

NETWORK TYPE	DEFINITION	DISTANCE RANGE	COMMUNICATION SPACE
LAN	Local Area Network	0.1 to 1 Km	Building, floor, Room
WAN	Wide Area Network	100 to 10000+ Km	Region, Country
MAN	Metropolitan Area Network	10 to 100 Km	City
CAN	Campus Area Network	1 to 10 Km	Campus, Military base, Compnay site
HAN	Home Area Network	0.1 Km	Home

1.5 Local Area Networks

LANs are networks usually confined to a geographic area, such as a single building, office. LANs can be small, linking as few as three computers, but often link hundreds of computers used by thousands of people. The development of standard networking protocols and media has resulted in worldwide proliferation of LANs throughout business organizations. This means that many users can share expensive devices, such as laser printers, as well as data. Users can also use the LAN to communicate with each other, by sending e-mail or engaging in chat sessions. Most LANs are built with relatively inexpensive hardware such as Ethernet cable and network interface cards (although wireless and other options exist). Specialized operating system software is also often used to configure a LAN. For example, some flavors of Microsoft Windows -- including Windows 98 SE, Windows 2000, and Windows ME -- come with a package called Internet Connection Sharing (ICS) that support controlled access to resources on the network.

LANs are usually faster than WANs, ranging in speed from 230 Kbps up to and beyond 1 Gbps (billion bits per second) as shown in Figure 1.2. They have very small delays of less than 10 milliseconds.

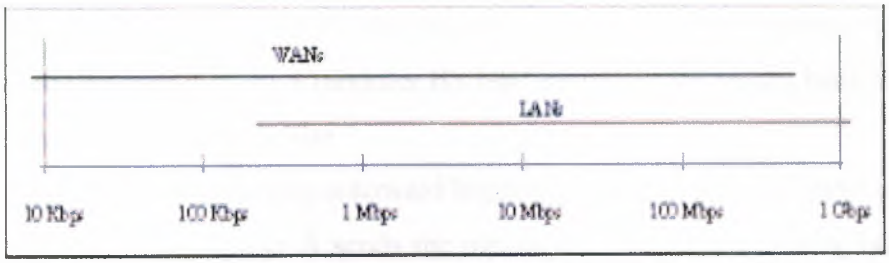

Figure 1.2 Data Speeds on LANs and WANs

How does one computer send information to another? It is actually rather simple. The figure 1.3 shows and explains a simple network.

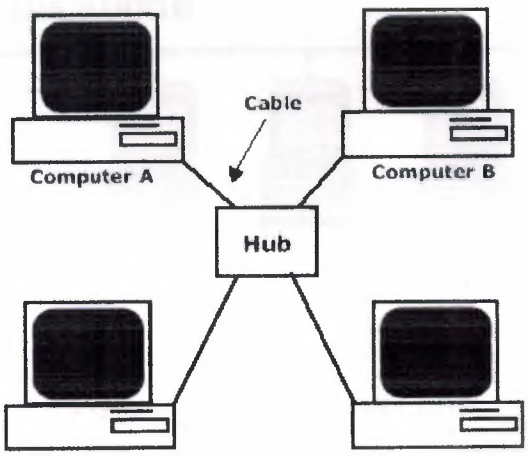

Figure 1.3 Simple Networks

1.6 Ethernet local area network

If Computer A wants to send a file to Computer B, the following would take place:

1. Based on a protocol that both computers use, the NIC in Computer A translates the file (which consists of binary data -- 1's and 0's) into pulses of electricity.

2. The pulses of electricity pass through the cable with a minimum (hopefully) of resistance.
3. The hub takes in the electric pulses and shoots them out to all of the other cables.
4. Computer B's NIC interprets the pulses and decides if the message is for it or not. In this case it is, so, Computer B's NIC translates the pulses back into the 1's and 0's that make up the file.

Sounds easy however, if anything untoward happen along the way, you have a problem, not a network. So, if Computer A sends the message to the network using NetBEUI, a Microsoft protocol, but Computer B only understands the TCP/IP protocol, it will not understand the message, no matter how many times Computer A sends it. Computer B also won't get the message if the cable is getting interference from the fluorescent lights etc. or if the network card has decided not to turn on today etc. Figure 1.4 shows small Ethernet local area network.

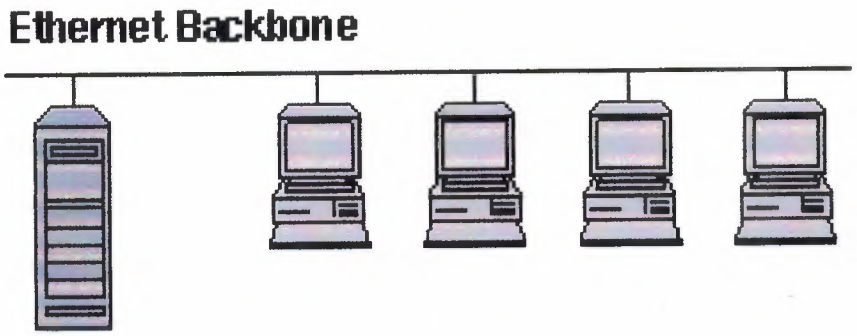

Figure 1.4 Small Ethernet LAN

The figure 1.5 shows briefly the interconnection of two LANs

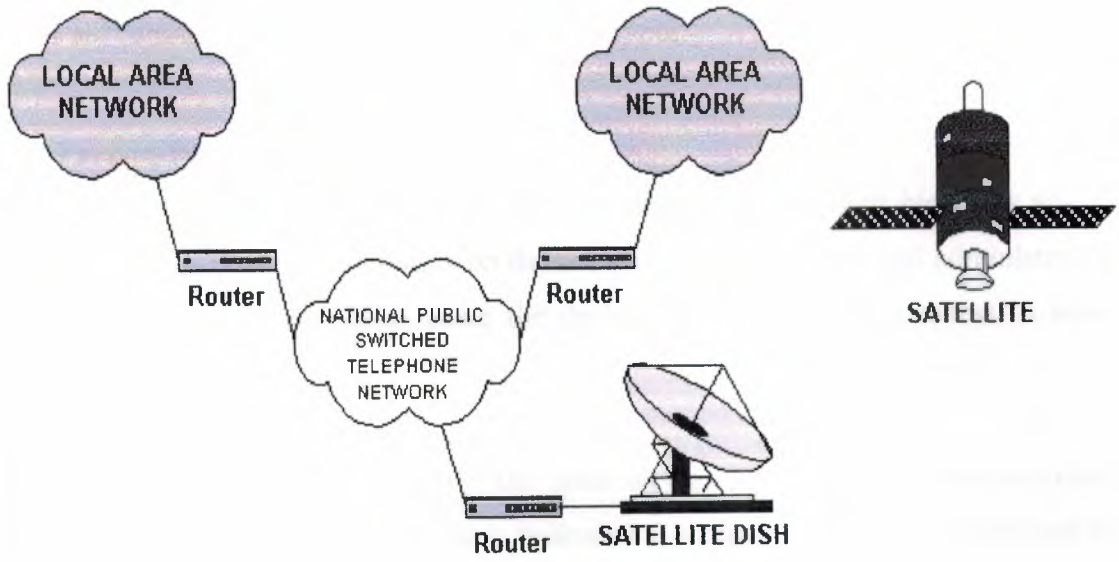

Figure 1.5 Interconnection of two LANs

1.7 Major Components of LANs

- Servers.
- Client / Workstation.
- Media.
- Shared Data.
- Shared Printers and other peripherals.
- Network Interface Card.
- Hubs / Concentrator.
- Repeaters, Bridges, Routers, Brouters, Gateways
- Physical connectors.
- Protocols.
- Network operating system (NOS).

1.8 Types of Local Area Networks

LANs are usually further divided into two major types:

1.8.1 Peer-to-Peer:

A peer-to-peer network doesn't have any dedicated servers or hierarchy among the computers. All of the computers on the network handle security and administration for themselves. The users must make the decisions about who gets access to what.

1.8.2 Client-Server:

A client-server network works the same way as a peer-to-peer network except that there is at least one computer that is dedicated as a server. The server stores files for sharing, controls access to the printer, and generally acts as the dictator of the network.

1.9 Local Area Networks Connectivity Devices

1.9.1 Repeaters

Boost signal in order to allow a signal to travel farther and prevent attenuation. Attenuation is the degradation of a signal as it travels farther from its origination. Repeaters do not filter packets and will forward broadcasts. Both segments must use the same access method, meaning that you can't connect a token ring segment to an Ethernet segment. Repeaters will connect different cable types.

1.9.2 Bridges

Functions the same as a repeater, but can also divide a network in order to reduce traffic problems. A bridge can also connect unlike network segments (i.e. token ring and Ethernet). Bridges create routing tables based on the source address. If the bridge can't find the source address it will forward the packets to all segments.

1.9.3 Routers

A router will do everything that a bridge will do and more. Routers are used in complex networks because they do not pass broadcast traffic. A router will determine the most efficient path for a packet to take and send packets around failed segments. Unroutable protocols can't be forwarded.

1.9.4 Brouters

A brouter has the best features of both routers and bridges in that it can be configured to pass the unroutable protocols by imitating a bridge, while not passing broadcast storms by acting as a router for other protocols.

1.9.5 Gateways

Often used as a connection to a mainframe or the internet. Gateways enable communications between different protocols, data types and environments. This is achieved via protocol conversion, whereby the gateway strips the protocol stack off of the packet and adds the appropriate stack for the other side.

1.10 Local Area Networks (LAN) in the work place and its advantages

Network allows more efficient management of resources. For example, multiple users can share a single top quality printer, rather than putting lesser quality printers on individual desktops. Also network software licenses can be less costly than separate, stand alone licenses for the same number of users.

Network helps keep information reliable and up-to-date. A well managed, centralized data storage system allows multiple users to access data from different locations, and limit access to data while it is being processed.

Network helps speeds up data sharing. Transferring files across a network is almost always faster than other, non-network means of sharing files.

Networks help business service their clients more effectively. Remote access to centralized data allows employees to service clients in the field, and clients to communicate directly to suppliers.

Speed: Networks provide a very rapid method for sharing and transferring files. Without a network, files are shared by copying them to floppy disks, then carrying or sending the disks from one computer to another. This method of transferring files is very time-consuming.

Security: Files and programs on a network can be designated as "copy inhibit," so that you do not have to worry about illegal copying of programs. Also, passwords can be established for specific directories to restrict access to authorized users.

Centralized Software Management: One of the greatest benefits of installing a local area network is the fact that all of the software can be loaded on one computer (the file server). This eliminates that need to spend time and energy installing updates and tracking files on independent computers throughout the building.

Electronic Mail: The presence of a network provides the hardware necessary to install an e-mail system. E-mail aids in personal and professional communication for all personnel, and it facilitates the dissemination of general information to the entire school staff. Electronic mail on a LAN can enable students to communicate with teachers and peers at their own school. If the LAN is connected to the Internet, people can communicate with others throughout the world. Network allows workgroups to communicate more effectively. Electronic mail and messaging is a staple of most network systems, in addition to scheduling systems, project monitoring, on-line conferencing and groupware, all of which help work teams be more productive.

Workgroup Computing: Workgroup software (such as Microsoft BackOffice) allows many users to work on a document or project concurrently. For example, educators located at various schools within a county could simultaneously contribute their ideas about new curriculum standards to the same document and spreadsheets.

1.11 Emerging Technology, Wireless Networks

Wireless networking refers to hardware and software combinations that enable two or more appliances to share data with each other without direct cable connections. Thus, in its widest sense, wireless networking includes cell and satellite phones, pagers, two-way radios, wireless LANs and modems, and Global Positioning Systems (GPS). Wireless LANs enable client computers and the server to communicate with one another without direct cable connections. Figure 1.6 shows the wireless network.

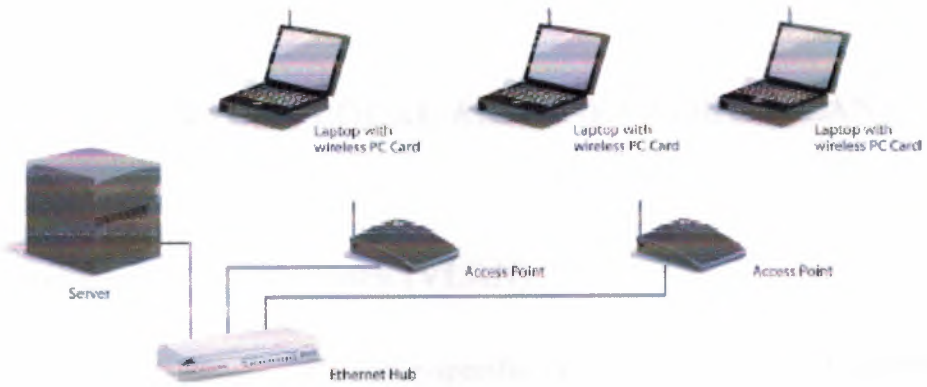

Figure 1.6 Wireless Network

2. VIRTUAL LOCAL AREA NETWORK (VLAN)

2.1 Virtual Local Area Network (VLAN)

With the multitude of vendor-specific VLAN solutions and implementation strategies, defining precisely what VLANs are has become a contentious issue. Nevertheless, most people would agree that a VLAN can be roughly equated to a broadcast domain. More specifically, VLANs can be seen as analogous to a group of end-stations, perhaps on multiple physical LAN segments, that are not constrained by their physical location and can communicate as if they were on a common LAN.

Virtual LANs (VLANs) can be viewed as a group of devices on different physical LAN segments which can communicate with each other as if they were all on the same physical LAN segment. VLANs provide a number of benefits over the network described in Figure 2.1, which we will discuss in the next section. In order to take advantage of the benefits of VLANs, a different network topology is needed.

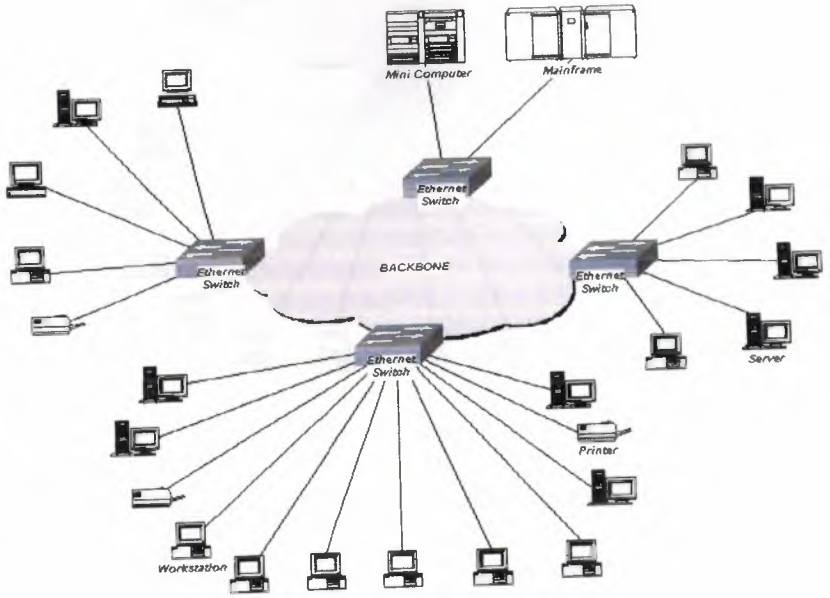

Figure 2.1 Typical Switched Networks

Using the same end nodes as in Figure 2.1, the switched network in Figure 2.2 provides the same connectivity as Figure 2.1. Although the network above has some distinct speed and latency advantages over the network in Figure 2.1, it also has some serious drawbacks. The most notable of these for the purposes of this discussion is that all hosts (end nodes) are now in the same broadcast domain. This adds a significant amount of traffic to the network that is seen by all hosts on the network. As this network grows, the broadcast traffic has the potential impact of flooding the network and making it essentially unusable.

Switches using VLANs create the same division of the network into separate broadcast domains but do not have the latency problems of a router. Switches are also a more cost-effective solution. Figure 2.2 shows a switched network topology using VLANs.

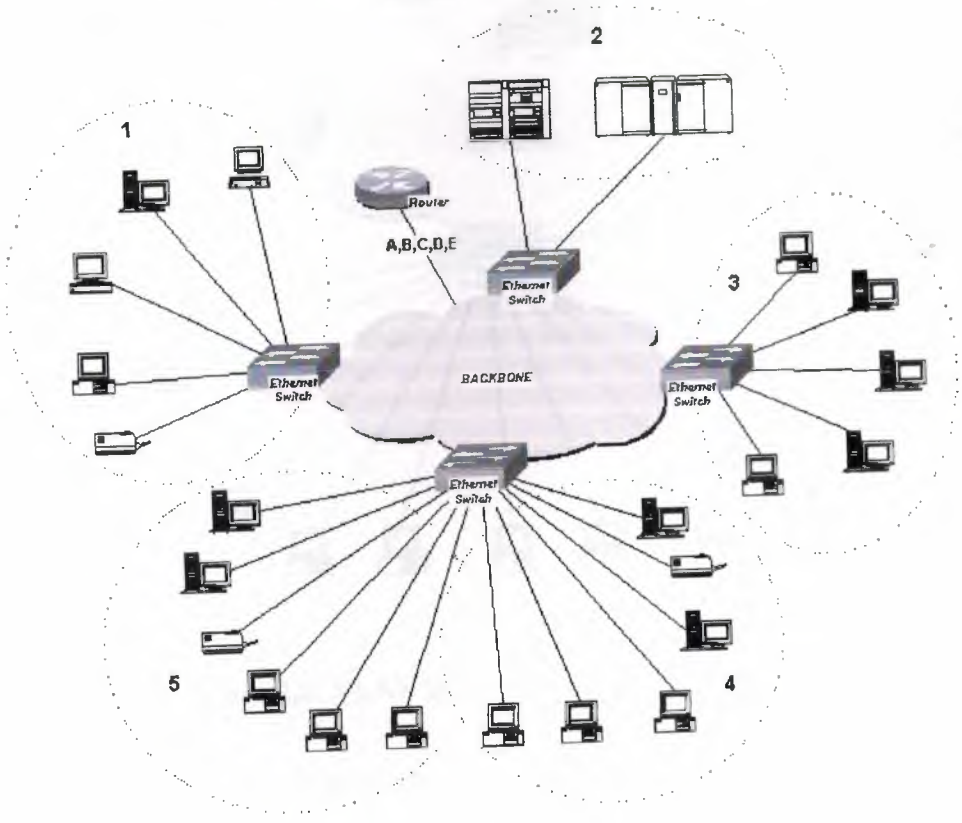

Figure 2.2 Switched Network with VLANs

Notice that the initial logical LAN topology from Figure 2.1 has been restored, with the major changes being the addition of Ethernet switches and the use of only one router. Notice also that the LAN identifiers appear on the single router interface. It is still necessary to use a router when moving between broadcast domains, and in this example, the router interface is a member of all of the VLANs. There are a number of ways to do this, and most are still proprietary and vendor-based.

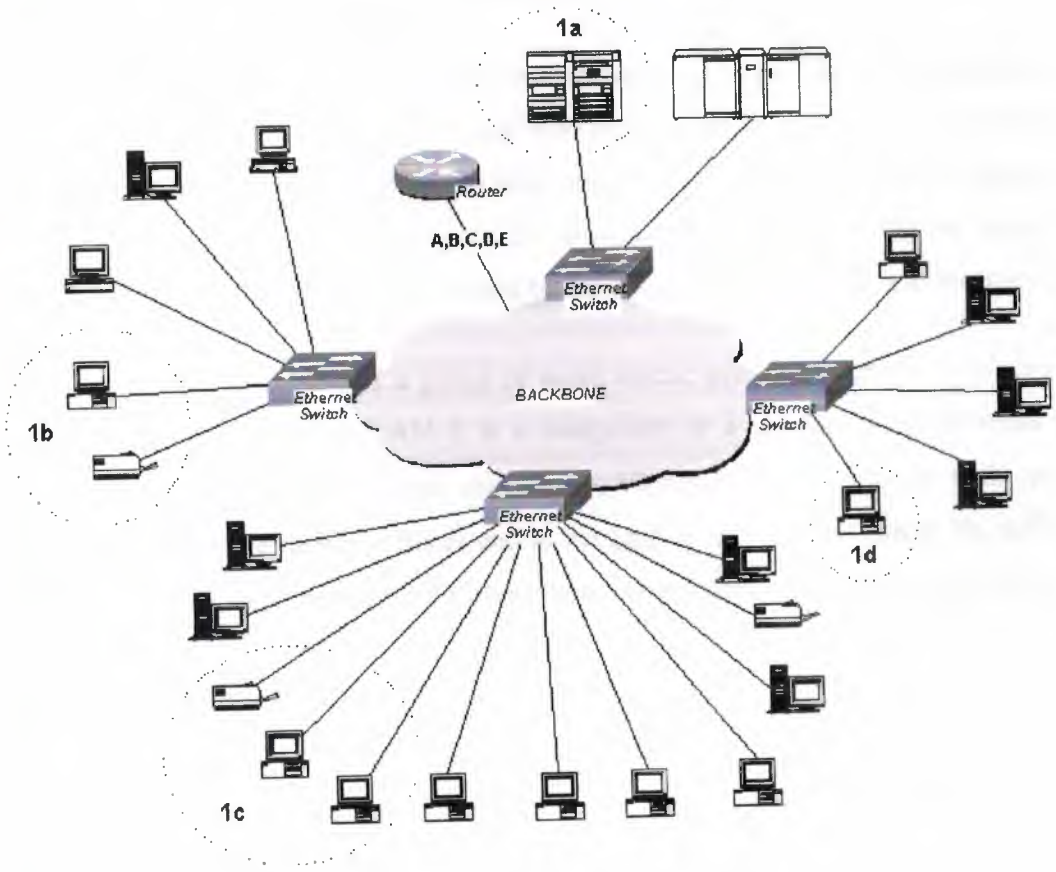

Figure 2.3 VLAN grouping using traffic patterns

By now you are probably wondering why someone would go to all this work to end up with what appears to be the same network (at least from a logical standpoint) as the original one. Consider Figure 2.4, where we begin to take advantage of some of the benefits of VLANs.

In the previous examples, LANs have been grouped with physical location being the primary concern. In Figure 2.4, VLAN 1 has been built with traffic patterns in mind. All of

the end devices in 1b, 1c, and 1d are primarily used for minicomputer access in 1a. Using VLANs, we are able to group these devices logically into a single broadcast domain. This allows us to confine broadcast traffic for this workgroup to just those devices that need to see it, and reduce traffic to the rest of the network. There is an increased connection speed due to the elimination of latency from router connections. An additional benefit of increased security could be realized if we made the decision to not allow access to the host from foreign networks, i.e., those that originate from another subnet beyond the router.

If we extend this thinking, we can now create a network that is independent of physical location and group users into logical workgroups. For instance, if a department has users in three different locations, they can now provide access to servers and printers as if they were all in the same building. Figure 2.5 illustrates this concept using the same end devices as in Figure2.1 and logically grouped by function, traffic patterns, and workgroups.

As in Figure2.4, VLAN 1 is a group of users whose primary function is to access a database on a minicomputer. VLAN 2 is a comprised of a similar group of users that require access to local servers and the mainframe. VLAN 3 is a department with servers and user workstations on different floors and in the case of the workstations in 3b, different buildings. VLANs 4 and 5 represent different departments with workstations and servers in single buildings.

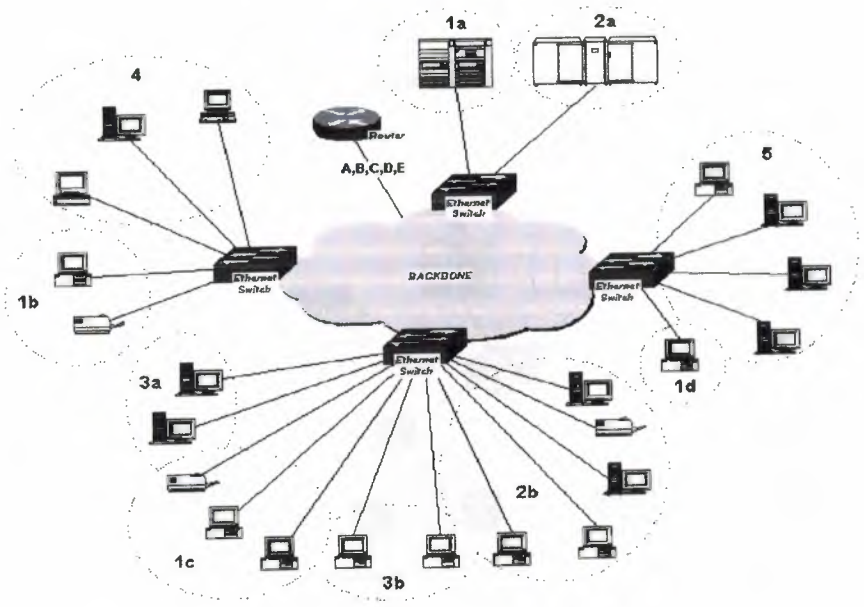

Figure 2.4 Logically grouped VLANs

One problem remains from the picture above. In a campus environment the size of UC Davis, it is difficult to scale the model above due to physical distances and sheer numbers.

Enter ATM and Network 21. The solution to these problems is to install ATM in the cloud and use something called LAN Emulation (LANE) to provide backbone services to the edge devices, or in this case, the Ethernet switches shown in Figure 2.3. Without going into detail, LAN Emulation over ATM provides the means to fully support existing LAN-based applications without changes. Advanced LAN Emulation software provides transparency to the underlying network's move to ATM. In addition, LANE provides the following benefits:

- Higher capacity
- Superior allocation and management of network capacity
- Easier management of the constantly changing LAN membership
- Access to multiple VLANs from the same physical interface
- Ease of evolution to new applications.

Figure2.5 gives us a look at VLANs in an ATM LANE environment. You'll notice that nothing has changed at the edges of the network, and a little more detail has been added at the core.

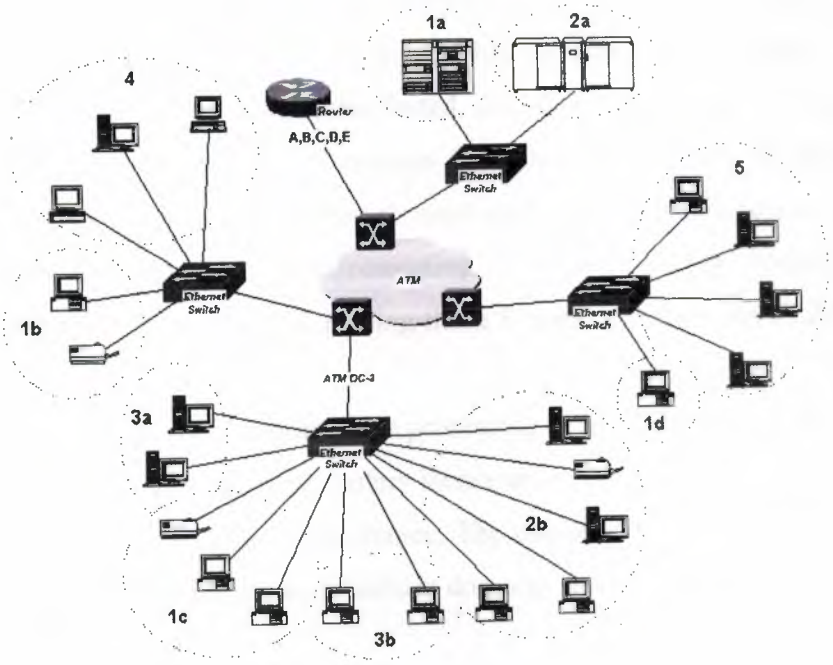

Figure 2.5 VLANs with ATM backbone

2.1.1 VLAN Benefits

As we have seen, there are several benefits to using VLANs. To summarize, VLAN architecture benefits include:

- Increased performance
- Improved manageability
- Network tuning and simplification of software configurations
- Physical topology independence
- Increased security options
- Reduced Cost

2.1.1.1 Increased performance

Switched networks by nature will increase performance over shared media devices in use today, primarily by reducing the size of collision domains.

Grouping users into logical networks will also increase performance by limiting broadcast traffic to users performing similar functions or within individual workgroups.

Additionally, less traffic will need to be routed, and the latency added by routers will be reduced, In networks where traffic consists of a high percentage of broadcasts and multicasts, VLAN's can reduce the need to send such traffic to unnecessary destinations. For example, in a broadcast domain consisting of 10 users, if the broadcast traffic is intended only for 5 of the users, then placing those 5 users on a separate VLAN can reduce traffic.

Compared to switches, routers require more processing of incoming traffic, As the volume of traffic passing through the routers increases, so does the latency in the routers, which results in reduced performance, The use of VLAN's reduces the number of routers needed, since VLAN's create broadcast domains using switches instead of routers.

2.1.1.2 Improved manageability

VLANs provide an easy, flexible, less costly way to modify logical groups in changing environments. VLANs make large networks more manageable by allowing centralized configuration of devices located in physically diverse locations.

Seventy percent of network costs are a result of adds, moves, and changes of users in the network [Buerger]. Every time a user is moved in a LAN, recabling, new station addressing, and reconfiguration of hubs and routers becomes necessary.

Some of these tasks can be simplified with the use of VLAN's. If a user is moved within a VLAN, reconfiguration of routers is unnecessary. In addition, depending on the type of VLAN, other administrative work can be reduced or eliminated [Cisco white paper]. However the full power of VLAN's will only really be felt when good management tools are created which can allow network managers to drag and drop users into different VLAN's or to set up aliases.

Despite this saving, VLAN's add a layer of administrative complexity, since it now becomes necessary to manage virtual workgroups

2.1.1.3 Network tuning and simplification of software configurations

VLANs will allow LAN administrators to "fine tune" their networks by logically grouping users. Software configurations can be made uniform across machines with the consolidation of a department's resources into a single subnet. IP addresses, subnet masks, and local network protocols will be more consistent across the entire VLAN. Fewer implementations of local server resources such as BOOTP and DHCP will be needed in this environment. These services can be more effectively deployed when they can span buildings within a VLAN.

Nowadays, it is common to find cross-functional product development teams with members from different departments such as marketing, sales, accounting, and research. These workgroups are usually formed for a short period of time. During this period, communication between members of the workgroup will be high. To contain broadcasts and multicasts within the workgroup, a VLAN can be set up for them. With VLAN's it is easier to place members of a workgroup together.

Without VLAN's, the only way this would be possible is to physically move all the members of the workgroup closer together. However, virtual workgroups do not come without problems. Consider the situation where one user of the workgroup is on the fourth floor of a building, and the other workgroup members are on the second floor. Resources such as a printer would be located on the second floor, which would be inconvenient for the lone fourth floor user.

Another problem with setting up virtual workgroups is the implementation of centralized server farms, which are essentially collections of servers and major resources for operating a network at a central location. The advantages here are numerous, since it is more efficient and cost-effective to provide better security, uninterrupted power supply, consolidated backup, and a proper operating environment in a single area than if the major resources were scattered in a building. Centralized server farms can cause problems when setting up virtual workgroups if servers cannot be placed on more than one VLAN. In such a case, the

server would be placed on a single VLAN and all other VLAN's trying to access the server would have to go through a router; this can reduce performance.

2.1.1.4 Physical topology independence

VLANs provide independence from the physical topology of the network by allowing physically diverse workgroups to be logically connected within a single broadcast domain. If the physical infrastructure is already in place, it now becomes a simple matter to add ports in new locations to existing VLANs if a department expands or relocates. These assignments can take place in advance of the move, and it is then a simple matter to move devices with their existing configurations from one location to another. The old ports can then be "decommissioned" for future use, or reused by the department for new users on the VLAN.

2.1.1.5 Increased security options

VLANs have the ability to provide additional security not available in a shared media network environment. By nature, a switched network delivers frames only to the intended recipients, and broadcast frames only to other members of the VLAN. This allows the network administrator to segment users requiring access to sensitive information into separate VLANs from the rest of the general user community regardless of physical location. In addition, monitoring of a port with a traffic analyzer will only view the traffic associated with that particular port, making discreet monitoring of network traffic more difficult. It should be noted that the enhanced security that is mentioned above is not to be considered an absolute safeguard against security infringements. What this provides is additional safeguards against "casual" but unwelcome attempts to view network traffic. Periodically, sensitive data may be broadcast on a network. In such cases, placing only those users who can have access to that data on a VLAN can reduce the chances of an outsider gaining access to the data. VLAN's can also be used to control broadcast domains, set up firewalls, restrict access, and inform the network manager of an intrusion.

2.1.1.6 Reduced Cost

VLAN's can be used to create broadcast domains, which eliminate the need for expensive routers.

2.1.2 VLAN Limitations

There are a few limitations to using VLANs, some of the more notable being:

- Broadcast limitations
- Device limitations
- Port constraints

2.1.2.1 Broadcast limitations

In order to handle broadcast traffic in an ATM VLAN environment it is necessary to have a special server that is an integrated part of the ATM infrastructure. This server has limitations in the number of broadcasts that may be forwarded. Some network protocols that will be running within individual VLANs, such as IPX and AppleTalk, make extensive use of broadcast traffic. This has the potential of impacting thresholds on the switches or broadcast servers and may require special consideration when determining VLAN size and configuration.

2.1.2.2 Device limitations

The number of Ethernet addresses than can be supported by each edge device is 500. This represents a distribution of about 20 devices per Network 21 port. These numbers are actual technical limitations that could be further reduced due to performance requirements of attached devices.

These limitations are above the recommended levels for high performance networking. From a pure performance standpoint, the ideal end-user device to Network 21 port ratio would be one device per port. From a practical point of view, a single Network 21 port could be shared by a number of devices that do not require a great deal of bandwidth and

belong to the same VLAN. An example of this would be a desktop computer, printer, and laptop computer for an individual user.

2.1.2.3 Port Constraints

If a departmental hub or switch is connected to a Network 21 port, every port on that hub must belong to the same VLAN. Hubs do not have the capability to provide VLANs to individual ports, and VLANs can not be extended beyond the edge device ports even if a switch capable of supporting VLANs is attached.

2.2 VLAN's working

When a LAN bridge receives data from a workstation, it tags the data with a VLAN identifier indicating the VLAN from which the data came. This is called explicit tagging. It is also possible to determine to which VLAN the data received belongs using implicit tagging. In implicit tagging the data is not tagged, but the VLAN from which the data came is determined based on other information like the port on which the data arrived. Tagging can be based on the port from which it came, the source Media Access Control (MAC) field, the source network address, or some other field or combination of fields. VLAN's are classified based on the method used. To be able to do the tagging of data using any of the methods, the bridge would have to keep an updated database containing a mapping between VLAN's and whichever field is used for tagging. For example, if tagging is by port, the database should indicate which ports belong to which VLAN. This database is called a filtering database. Bridges would have to be able to maintain this database and also to make sure that all the bridges on the LAN have the same information in each of their databases. The bridge determines where the data is to go next based on normal LAN operations. Once the bridge determines where the data is to go, it now needs to determine whether the VLAN identifier should be added to the data and sent. If the data is to go to a device that knows about VLAN implementation (VLAN-aware), the VLAN identifier is added to the data. If it is to go to a device that has no knowledge of VLAN implementation (VLAN-unaware), the bridge sends the data without the VLAN identifier.

In order to understand how VLAN's work, we need to look at the types of VLAN's, the types of connections between devices on VLAN's, the filtering database which is used to send traffic to the correct VLAN, and tagging, a process used to identify the VLAN originating the data.

2.2.1 Types of VLAN's

VLAN membership can be classified by port, MAC address, and protocol type.

2.2.1.1 Layer 1 VLAN: Membership by Port

Membership in a VLAN can be defined based on the ports that belong to the VLAN. For example, in a bridge with four ports, ports 1, 2, and 4 belong to VLAN 1 and port 3 belongs to VLAN 2 (Figure2.6).

Port	VLAN
1	1
2	1
3	2
4	1

Figure 2.6 Assignment of ports to different VLAN's.

The main disadvantage of this method is that it does not allow for user mobility. If a user moves to a different location away from the assigned bridge, the network manager must reconfigure the VLAN.

2.2.1.2 Layer 2 VLAN: Membership by MAC Address

Here, membership in a VLAN is based on the MAC address of the workstation. The switch tracks the MAC addresses which belong to each VLAN (Figure2.7). Since MAC addresses form a part of the workstation's network interface card, when a workstation is moved, no reconfiguration is needed to allow the workstation to remain in the same VLAN. This is unlike Layer 1 VLAN's where membership tables must be reconfigured.

MAC Address	VLAN
1212354145121	1
2389234873743	2
3045834758445	2
5483573475843	1

Figure 2.7 Assignment of MAC addresses to different VLAN's.

The main problem with this method is that VLAN membership must be assigned initially. In networks with thousands of users, this is no easy task. Also, in environments where notebook PC's are used, the MAC address is associated with the docking station and not with the notebook PC. Consequently, when a notebook PC is moved to a different docking station, its VLAN membership must be reconfigured.

2.2.1.3 Layer 2 VLAN: Membership by Protocol Type

VLAN membership for Layer 2 VLAN's can also be based on the protocol type field found in the Layer 2 header (Figure2.8).

Protocol	VLAN
IP	1
IPX	2

Figure 2.8 Assignment of protocols to different VLAN's

2.2.1.4 Layer 3 VLAN: Membership by IP Subnet Address

Membership is based on the Layer 3 header. The network IP subnet address can be used to classify VLAN membership (Figure2.9).

IP Subnet	VLAN
23.2.24	1
26.21.35	2

Figure 2.9 Assignment of IP subnet addresses to different VLAN's.

Although VLAN membership is based on Layer 3 information, this has nothing to do with network routing and should not be confused with router functions. In this method, IP addresses are used only as a mapping to determine membership in VLAN's. No other processing of IP addresses is done.

In Layer 3 VLAN's, users can move their workstations without reconfiguring their network addresses. The only problem is that it generally takes longer to forward packets using Layer 3 information than using MAC addresses.

2.2.1.5 Higher Layer VLAN's

It is also possible to define VLAN membership based on applications or service, or any combination thereof. For example, file transfer protocol (FTP) applications can be executed on one VLAN and telnet applications on another VLAN. The 802.1Q draft standard defines Layer 1 and Layer 2 VLAN's only. Protocol type based VLAN's and higher layer VLAN's have been allowed for, but are not defined in this standard. As a result, these VLAN's will remain proprietary.

2.2.2 Types of Connections

Devices on a VLAN can be connected in three ways based on whether the connected devices are VLAN-aware or VLAN-unaware. Recall that a VLAN-aware device is one which understands VLAN memberships (i.e. which users belong to a VLAN) and VLAN formats.

2.2.2.1 Trunk Link

All the devices connected to a trunk link, including workstations, must be VLAN-aware. All frames on a trunk link must have a special header attached. These special frames are called tagged frames (figure 2.10).

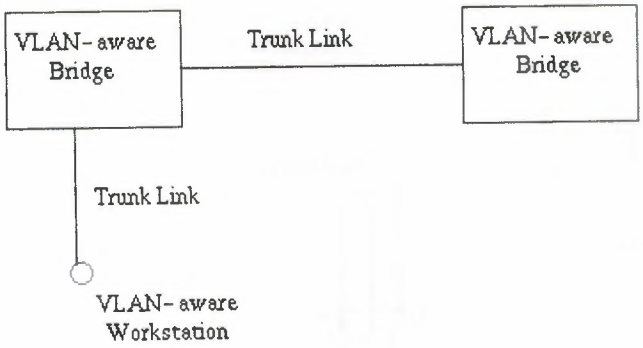

Figure 2.10 Trunk link between two VLAN-aware bridges.

2.2.2.2 Access Link

An access link connects a VLAN-unaware device to the port of a VLAN-aware bridge. All frames on access links must be implicitly tagged (untagged) (Figure2.11). The VLAN-unaware device can be a LAN segment with VLAN-unaware workstations or it can be a number of LAN segments containing VLAN-unaware devices (legacy LAN).

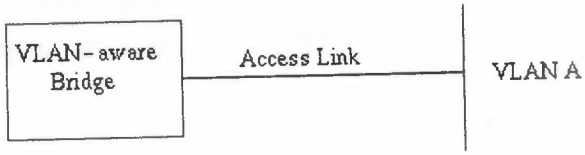

Figure 2.11 Access link between a VLAN-aware bridge and a VLAN-unaware device.

2.2.2.3 Hybrid Link

This is a combination of the previous two links. This is a link where both VLAN-aware and VLAN-unaware devices are attached (Figure2.12). A hybrid link can have both tagged and untagged frames, but all the frames for a specific VLAN must be either tagged or untagged.

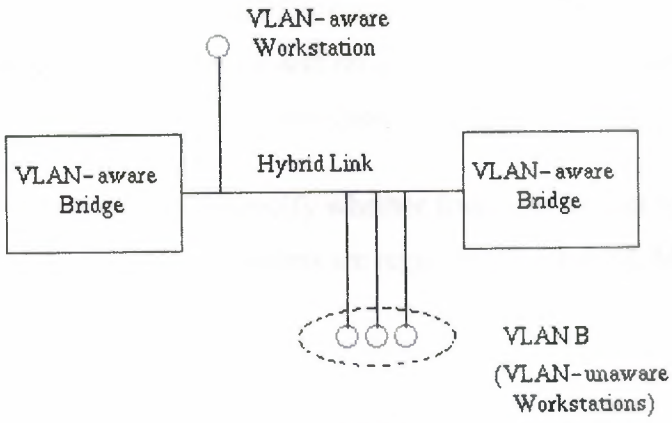

Figure 2.12 Hybrid link containing both VLAN-aware and VLAN-unaware devices.

It must also be noted that the network can have a combination of all three types of links.

2.2.3 Frame Processing

A bridge on receiving data determines to which VLAN the data belongs either by implicit or explicit tagging. In explicit tagging a tag header is added to the data. The bridge also keeps track of VLAN members in a filtering database which it uses to determine where the data is to be sent. Following is an explanation of the contents of the filtering database and the format and purpose of the tag header [802.1Q].

2.2.4 Filtering Database

Membership information for a VLAN is stored in a filtering database. The filtering database consists of the following types of entries:

2.2.4.1 Static Entries

Static information is added, modified, and deleted by management only. Entries are not automatically removed after some time (ageing), but must be explicitly removed by management. There are two types of static entries:

a) Static Filtering Entries: which specify for every port whether frames to be sent to a specific MAC address or group address and on a specific VLAN should be forwarded or discarded, or should follow the dynamic entry, and

b) Static Registration Entries: which specify whether frames to be sent to a specific VLAN are to be tagged or untagged and which ports are registered for that VLAN.

2.2.4.2 Dynamic Entries

Dynamic entries are learned by the bridge and cannot be created or updated by management. The learning process observes the port from which a frame, with a given source addresses and VLAN ID (VID), is received, and updates the filtering database. The entry is updated only if all the following three conditions are satisfied:

a) This port allows learning,

b) The source address is a workstation address and not a group address, and

c) There is space available in the database.

Entries are removed from the database by the ageing out process where, after a certain amount of time specified by management (10 sec --- 1000000 sec), entries allow automatic reconfiguration of the filtering database if the topology of the network changes. There are three types of dynamic entries:

a) Dynamic Filtering Entries: which specify whether frames to be sent to a specific MAC address and on a certain VLAN should be forwarded or discarded.

b) Group Registration Entries: which indicate for each port whether frames to be sent to a group MAC address and on a certain VLAN should be filtered or discarded. These entries are added and deleted using Group Multicast Registration Protocol (GMRP). This allows multicasts to be sent on a single VLAN without affecting other VLAN's.

c) Dynamic Registration Entries: which specify which ports are registered for a specific VLAN. Entries are added and deleted using GARP VLAN Registration Protocol (GVRP), where GARP is the Generic Attribute Registration Protocol.

GVRP is used not only to update dynamic registration entries, but also to communicate the information to other VLAN-aware bridges.

In order for VLAN's to forward information to the correct destination, all the bridges in the VLAN should contain the same information in their respective filtering databases. GVRP allows both VLAN-aware workstations and bridges to issue and revoke VLAN memberships. VLAN-aware bridges register and propagate VLAN membership to all ports that are a part of the active topology of the VLAN. The active topology of a network is determined when the bridges are turned on or when a change in the state of the current topology is perceived.

The active topology is determined using a spanning tree algorithm which prevents the formation of loops in the network by disabling ports. Once an active topology for the network (which may contain several VLAN's) is obtained, the bridges determine an active topology for each VLAN. This may result in a different topology for each VLAN or a common one for several VLAN's. In either case, the VLAN topology will be a subset of the active topology of the network (Figure2.13).

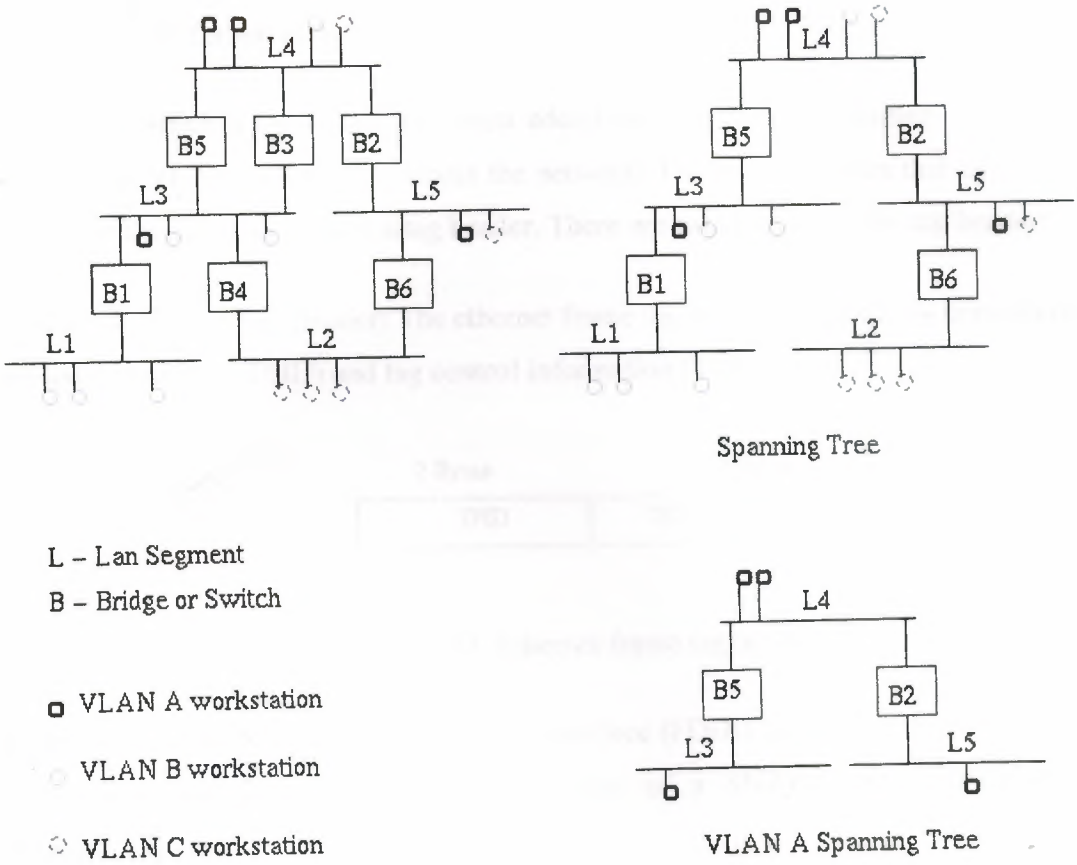

Figure 2.13 Active topology of network and VLAN A using spanning tree algorithm.

2.2.5 Tagging

When frames are sent across the network, there needs to be a way of indicating to which VLAN the frame belongs, so that the bridge will forward the frames only to those ports that belong to that VLAN, instead of to all output ports as would normally have been done. This information is added to the frame in the form of a tag header. In addition, the tag header:

a. allows user priority information to be specified.

b. allows source routing control information to be specified.

c. indicates the format of MAC addresses.

Frames in which a tag header has been added are called tagged frames. Tagged frames convey the VLAN information across the network. The tagged frames that are sent across hybrid and trunk links contain a tag header. There are two formats of the tag header:

a. Ethernet Frame Tag Header: The ethernet frame tag header (Figure2.14) consists of a tag protocol identifier (TPID) and tag control information (TCI).

Figure 2.14 Ethernet frame tag header.

b. Token Ring and Fiber Distributed Data Interface (FDDI) tag header: The tag headers for both token ring and FDDI networks consist of a SNAP-encoded TPID and TCI. (Figure2.15)

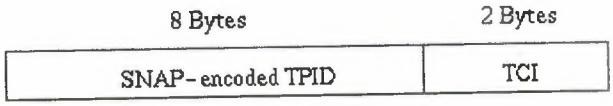

Figure 2.15 Token ring and FDDI tag header.

TPID is the tag protocol identifier which indicates that a tag header is following and TCI (Figure2.16) contains the user priority, canonical format indicator (CFI), and the VLAN ID.

Figure 2.16 Tag control information (TCI).

User priority is a 3 bit field which allows priority information to be encoded in the frame. Eight levels of priority are allowed, where zero is the lowest priority and seven is the highest priority. How this field is used is described in the supplement 802.1p.

The CFI bit is used to indicate that all MAC addresses present in the MAC data field are in canonical format. This field is interpreted differently depending on whether it is an ethernet-encoded tag header or a SNAP-encoded tag header. In SNAP-encoded TPID the field indicates the presence or absence of the canonical format of addresses. In ethernet-encoded TPID, it indicates the presence of the Source-Routing Information (RIF) field after the length field. The RIF field indicates routing on ethernet frames.

The VID field is used to uniquely identify the VLAN to which the frame belongs. There can be a maximum of $(2^{12} - 1)$ VLAN's. Zero is used to indicate no VLAN ID, but that user priority information is present. This allows priority to be encoded in non-priority LAN's

3. APPLICATION OF VIRTUAL LOCAL AREA NETWORK

3.1 Membership by MAC Address

VLAN membership based on MAC-layer address has a different set of advantages and disadvantages. Since MAC-layer addresses are hard-wired into the workstation's network interface card (NIC), VLANs based on MAC addresses enable network managers to move a workstation to a different physical location on the network and have that workstation automatically retain its VLAN membership.

In this way, a VLAN defined by MAC address can be thought of as a user based VLAN. One of the drawbacks of MAC address–based VLAN solutions is the requirement that all users must initially be configured to be in at least one VLAN. After that initial manual configuration, automatic tracking of users is possible, depending on the specific vendor solution. However, the disadvantage of having to initially configure VLANs becomes clear in very large networks where thousands of users must each be explicitly assigned to a particular VLAN. Some vendors have mitigated the onerous task of initially configuring MAC based VLANs by using tools that create VLANs based on the current state of the network—that is, a MAC address–based VLAN is created for each subnet. MAC address–based VLANs that are implemented in shared media environments will run into serious performance degradation as members of different VLANs coexist on a single switch port. In addition, the primary method of communicating VLAN membership information between switches in a MAC address–defined VLAN also runs into performance degradation with larger-scale implementations.

This is explained in "Communicating VLAN Membership Information," later in this paper. Another, but minor, drawback to VLANs based only on MAC-layer addresses emerges in environments that use significant numbers of notebook PCs with some docking stations.

The problem is that the docking station and integrated network adapter (with its hard-wired MAC-layer address) usually remains on the desktop, while the notebook travels with the user. Making VLAN membership impossible to track when the user moves to a new desk and docking station, the MAC-layer address changes. In such an environment, VLAN

membership must be updated constantly as users move around and use different docking stations. While this problem may not be particularly common, it does illustrate some of the limitations of MAC address–based VLANs.

3.2 Layer 3–Based VLANs

VLANs based on layer-3 information take into account protocol type (if multiple protocols are supported) or network-layer address (for example, subnet address for TCP/IP networks) in determining VLAN membership. Although these VLANs are based on layer 3 information, this does not constitute a "routing" function and should not be confused with network-layer routing.

Even though a switch inspects a packet's IP address to determine VLAN membership, no route calculation is undertaken, RIP or OSPF protocols are not employed, and frames traversing the switch are usually bridged according to implementation of the Spanning Tree Algorithm. Therefore, from the point of view of a switch employing layer 3–based VLANs, connectivity within any given VLAN is still seen as a flat, bridged topology.

Having made the distinction between VLANs based on layer-3 information and routing, it should be noted that some vendors are incorporating varying amounts of layer 3 intelligence into their switches, enabling functions normally associated with routing. Furthermore, "layer 3 aware" or "multi-layer" switches often have the packet-forwarding function of routing built into ASIC chip sets, greatly improving performance over CPU-based routers. Nevertheless, a key point remains: no matter where it is located in a VLAN solution, routing is necessary to provide connectivity between distinct VLANs.

There are several advantages to defining VLANs at layer 3. First, it enables partitioning by protocol type. This may be an attractive option for network managers who are dedicated to a service- or application-based VLAN strategy. Second, users can physically move their workstations without having to reconfigure each workstation's network address—a benefit primarily for TCP/IP users. Third, defining VLANs at layer 3 can eliminate the need for frame tagging in order to communicate VLAN membership between switches, reducing transport overhead.

One of the disadvantages of defining VLANs at layer 3 (vs. MAC- or port-based VLANs) can be performance. Inspecting layer 3 addresses in packets is more time consuming than looking at MAC addresses in frames. For this reason, switches that use layer-3 information for VLAN definition are generally slower than those that use layer-2 information. It should be noted that this performance difference is true for most, but not all, vendor implementations.

VLANs defined at layer 3 are particularly effective in dealing with TCP/IP, but less effective with protocols such as IPX™, DECnet®, or AppleTalk®, which do not involve manual configuration at the desktop. Furthermore, layer 3–defined VLANs have particular difficulty in dealing with "un-routable" protocols such as NetBIOS. End stations running un-routable protocols cannot be differentiated and thus cannot be defined as part of a network-layer VLAN.

3.3 IP Multicast Groups as VLANs

IP multicast groups represent a somewhat different approach to VLAN definition, although the fundamental concept of VLANs as broadcast domains still applies. When an IP packet is sent via multicast, it is sent to an address that is a proxy for an explicitly defined group of IP addresses that is established dynamically.

Each workstation is given the opportunity to join a particular IP multicast group by responding affirmatively to a broadcast notification, which signals that group's existence. All workstations that join an IP multicast group can be seen as members of the same virtual LAN. However, they are only members of a particular multicast group for a certain period of time. Therefore, the dynamic nature of VLANs defined by IP multicast groups enables a very high degree of flexibility and application sensitivity.

In addition, VLANs defined by IP multicast groups would inherently be able to span routers and thus WAN connections.

3.4 Combination VLAN Definitions

Due to the trade-offs between various types of VLANs, many vendors are planning to include multiple methods of VLAN definition. Such a flexible definition of VLAN membership enables network managers to configure their VLANs to best suit their particular network environment. For example, by using a combination of methods, an organization that utilizes both IP and NetBIOS protocols could define IP VLANs corresponding to preexisting IP subnets (convenient for smooth migration), and then define VLANs for NetBIOS end stations by dividing them by groups of MAC-layer addresses.

3.5 Automation of VLAN Configuration

Another issue central to VLAN deployment is the degree to which VLAN configuration is automated. To a certain extent, this degree of automation is correlated to how VLANs are defined; but in the end, the specific vendor solution will determine this level of automation.

There are three primary levels of automation in VLAN configuration:

1. Manual. With purely manual VLAN configuration, both the initial setup and all subsequent moves and changes are controlled by the network administrator.

Of course, purely manual configuration enables a high degree of control. However, in larger enterprise networks, manual configuration is often not practical. Furthermore, it defeats one of the primary benefits of VLANs: elimination of the time it takes to administer moves and changes—although moving users manually with VLANs may actually be easier than moving users across router subnets, depending on the specific vendor's VLAN management interface.

2. Semi automated configuration refers to the option to automate initial configuration, subsequent reconfigurations (moves/changes), or both. Initial configuration automation is normally accomplished with a set of tools that map VLANs to existing subnets or other

criteria. Semi automated configuration could also refer to situations where VLANs are initially configured manually, with all subsequent moves being tracked automatically. Combining both initial and subsequent configuration automation would still imply semi-automated configuration, because the network administrator always has the option of manual configuration.

3. Fully Automatic. A system that fully automates VLAN configuration implies that workstations automatically and dynamically join VLANs depending on application user ID or other criteria or policies that are preset by the administrator. This type of VLAN configuration is discussed in greater detail toward the end of this chapter.

3.6 Communicating VLAN Membership Information

Switches must have a way of understanding VLAN membership (that is, which stations belong to which VLAN) when network traffic arrives from other switches; otherwise, VLANs would be limited to a single switch. In general, layer 2–based VLANs (defined by port or MAC address) must communicate VLAN membership explicitly, while VLAN membership in IP-based VLANs is implicitly communicated by the IP address. Depending on the particular vendor's solution, communications of VLAN membership may also be implicit in the case of layer 3–based VLANs in a multi-protocol environment.
To date, outside of implementing an ATM backbone, three methods have been implemented for inter-switch communication of VLAN information across a backbone:

1. Table Maintenance via Signaling. This method operates as follows: When an end-station broadcasts its first frame, the switch resolves the end-station's MAC address or attached port with its VLAN membership in cached address tables. This information is then broadcast continuously to all other switches. As VLAN membership changes, these address tables are manually updated by a system administrator at a management console. As the network expands and switches are added, the constant signaling necessary to update the

cached address tables of each switch can cause substantial congestion of the backbone. For this reason, this method does not scale particularly well.

2. Frame Tagging. In the frame-tagging approach, a header is typically inserted into each frame on inter-switch trunks to uniquely identify which VLAN a particular MAC-layer frame belongs to. Vendors differ in the way they solve the problem of occasionally exceeding the maximum length of MAC-layer frames as these headers are inserted. These headers also add overhead to network traffic.

3. TDM. The third, and least utilized method, is time-division multiplexing. TDM works the same way on the inter-switch backbone to support VLANs as it does in the WAN environment to support multiple traffic types here, channels are reserved for each VLAN. This approach cuts out some of the overhead problems inherent in signaling and frame tagging, but it also wastes bandwidth, because a time slot dedicated to one VLAN cannot be used by another VLAN, even if that channel is not carrying traffic. Deploying an ATM backbone also enables the communication of VLAN information between switches, but it introduces a new set of issues with regard to LAN Emulation (LANE). ATM is discussed in detail in a separate section of this chapter. However, for the time being, it should be remembered that with port group–defined VLANs, the LANE standard provides for a nonproprietary method of communicating VLAN membership across a backbone.

3.7 Virtual LAN Management Protocol (VLMP)

The Virtual LAN Management Protocol (VLMP) is a protocol for communicating virtual LAN information between switches in a vendor-independent form. It allows switches from different vendors to interoperate even if the end-station or port membership in virtual LANs is specified in different ways on each switch.

The basic VLMP model is as follows. A switch identifies each virtual LAN by character string name. For each virtual LAN, there is a group of switches identified by the name of

the virtual LAN whose members consist of those switches that need to forward packets for that virtual LAN.

VLMP includes a group membership protocol portion that is used by a switch to join the group corresponding to a virtual LAN for to each virtual LAN the switch handles. VLMP also specifies a multicast call for disseminating virtual LAN membership information at the MAC address level. Optionally, a switch can use VLMP to query another switch to determine the virtual LANs to which a specified MAC address belongs. Finally, VLMP allows a switch to invalidate the virtual LAN information in other switches. VLMP is specified as an RPC protocol, first as a procedure interface and then as a packet format using ONC RPC and ONC XDR. This approach follows that taken with EGMP and has similar advantages. That is, it allows the protocol to be generated by an RPC stub generator1. It uses a standard familiar packet format to support versions and data representations, and it imposes a procedural structuring on the protocol.

3.8 VLAN Trunk Protocol (VTP)

Virtual Local Area Network (VLAN) Trunk Protocol (VTP) reduces administration in a switched network. When you configure a new VLAN on one VTP server, the VLAN is distributed through all switches in the domain. This reduces the need to configure the same VLAN everywhere. VTP is a Cisco-proprietary protocol that is available on most of the Cisco Catalyst Family.

3.8.1 VTP Messages in Detail

VTP packets are sent in either ISL frames or in dot1q frames. These packets are sent to the destination MAC address 01-00-0c-cc-cc-cc with a Logical Link Control (LLC) code of Sub-network Access Protocol (SNAP) (AAAA) and a type of 2003 (in the SNAP header). Below is the format of a VTP packet encapsulated in ISL frames (figure 3.1):

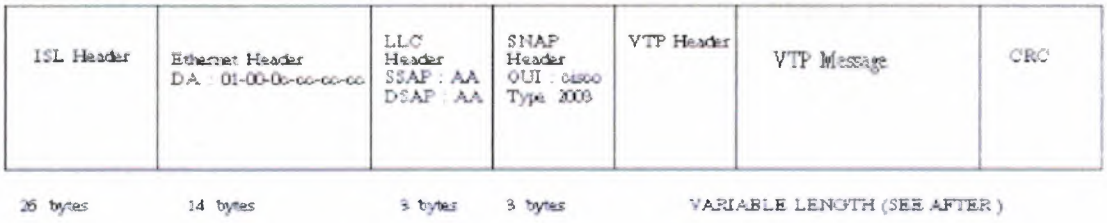

Figure3.1 Format of a VTP Packet in ISL Frames

You can, of course, have a VTP packet inside 802.1Q frames. In that case, the ISL header and Cyclic Redundancy Check (CRC) would be replaced by dot1q tagging. The format of the VTP header can vary depending on the type of VTP message. However, they all contain the following fields in the header:

1. VTP protocol version: 1 or 2
2. VTP message types:
    a. Summary advertisements
    b. Subset advertisement
    c. Advertisement requests
    d. VTP join messages
3. Management domain length
4. Management domain name

3.8.2 Configuration Revision Number

The configuration revision number is a 32-bit number that indicates the level of revision for a VTP packet. Each VTP device tracks the VTP configuration revision number assigned to it, and most of the VTP packets contain the VTP configuration revision number of the sender. This information is used to determine whether the received information is more recent than the current version. Each time you make a VLAN change in a VTP device, the configuration revision is incremented by one. In order to reset the configuration

revision of a switch, change the VTP domain name and then change it back to the original name.

3.8.3 Summary Advertisements

By default, Catalyst switches issue summary advertisements in five-minute increments. Summary advertisements inform adjacent Catalysts of the current VTP domain name and the configuration revision number.

When the switch receives a summary advertisement packet, it compares the VTP domain name to its own VTP domain name. If the name is different, the switch simply ignores the packet. If the name is the same, the switch then compares the configuration revision to its own revision. If its own configuration revision is higher or equal, the packet is ignored. If it is lower, an advertisement request is sent (figure 3.2).

Summary Advert Packet Format:

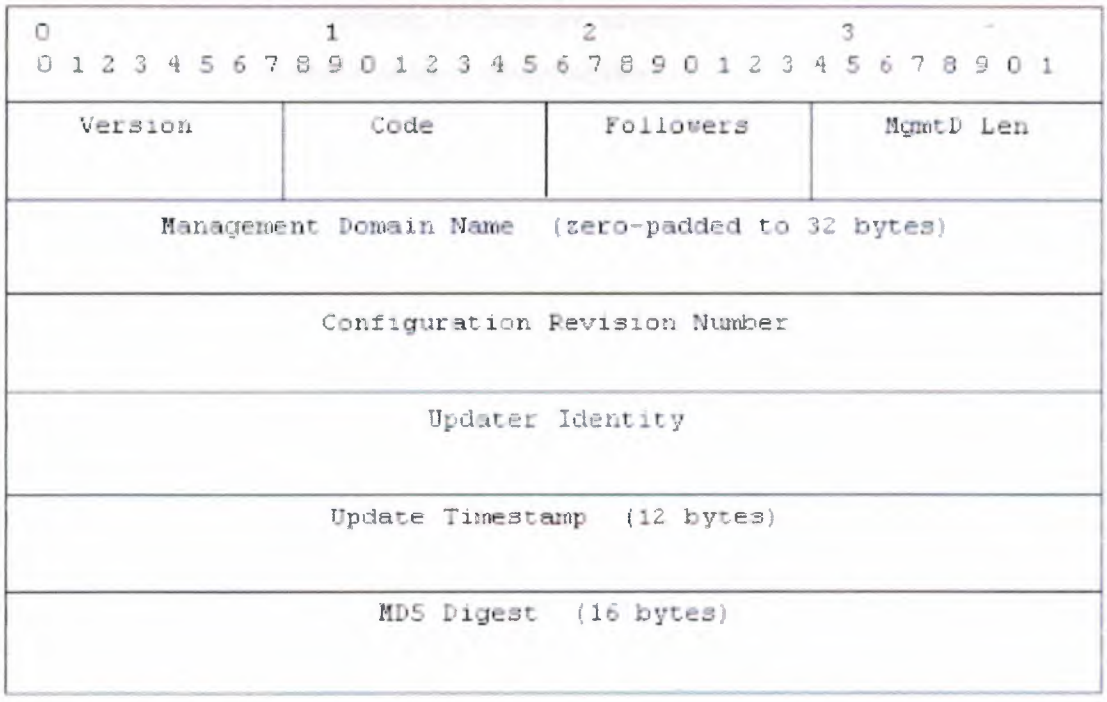

Figure3.2 Advert Packet Format

The following list clarifies the meaning of these fields in the summary advert packet:

1. Followers indicate that this packet is followed by a Subset Advertisement packet.
2. The updater identity is the IP address of the switch that is the last to have incremented the Configuration revision.
3. Update timestamps are the date and time of the last increment of the configuration revision.
4. Message Digest 5 (MD5) carries the VTP password if it is configured and used to authenticate the validation of a VTP update.

3.8.4 Subset Advertisements

When you add, delete, or change a VLAN in a Catalyst, the server Catalyst where the changes were made increments the configuration revision and issues a summary advertisement, followed by one or several subset advertisements. A subset advertisement contains a list of VLAN information. If there are several VLANS, more than one subset advertisement may be required in order to advertise them all (figure 3.3).

Subset Advert Packet Format:

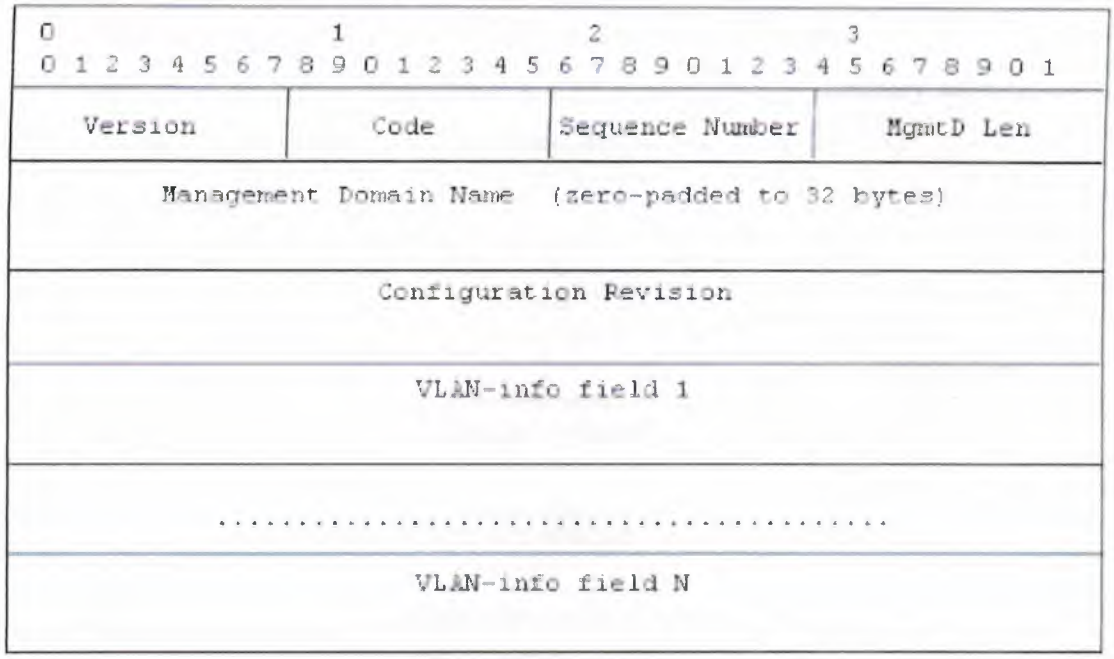

Figure3.3 Subset Advert Packet Format

The following formatted example shows that each VLAN information field contains information for a different VLAN (ordered with lowered-valued ISL VLAN IDs occurring first) (figure 3.4):

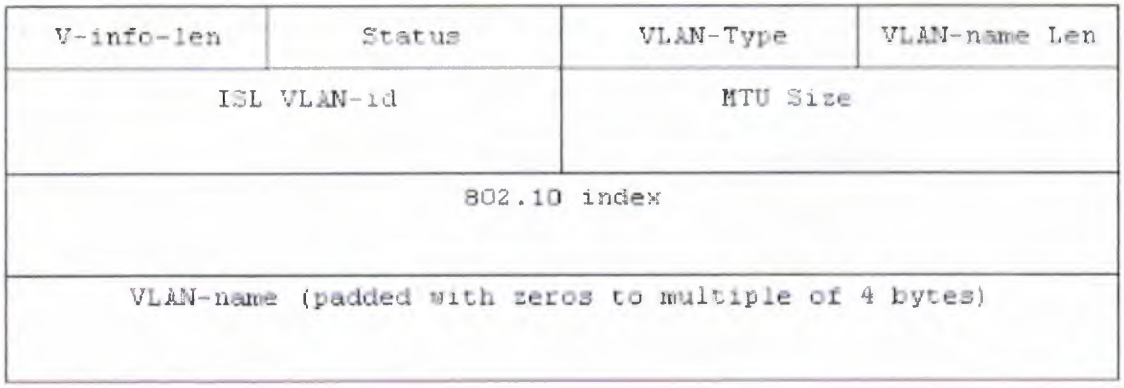

Figure3.4 Example 1

1. Code The format for this is 0x02 for subset advertisement.
2. Sequence number this is the sequence of the packet in the stream of packets following a summary advertisement. The sequence starts with 1.

Upon receipt of an advertisement request, a VTP device sends a summary advertisement, followed by one or more subset advertisements. Below is an example (3.5).

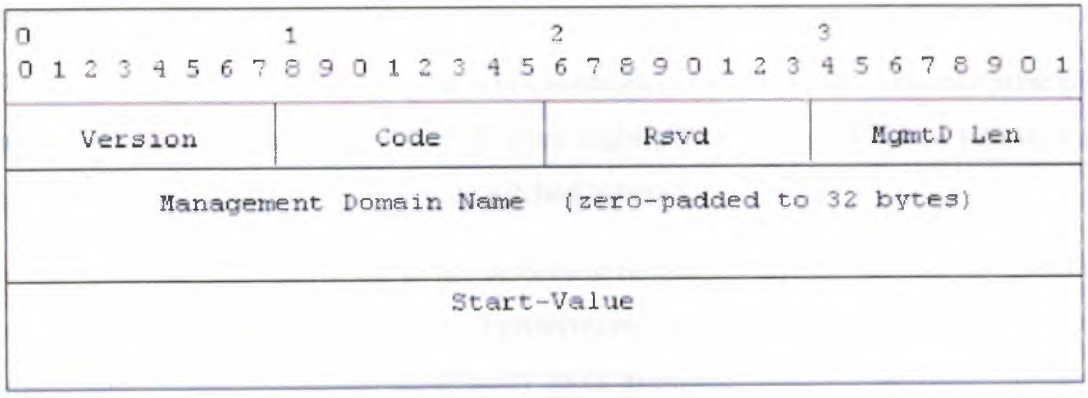

Figure3.5 Example 2

3. Code The format for this is 0x03 for an advertisement request.
4. Start Value this is used in cases where there are several subset advertisements. If the first (N) subset advertisement has been received and the subsequent one (N+1) has not, the Catalyst only requests advertisements from the (N+1) the one.

3.9 Virtual LAN Security Best Practices

The security of VLAN technology has proven to be far more reliable than its detractors had hoped for and only user miss configuration or improper use of features have been pointed out as ways to undermine its robustness.

The most serious mistake that a user can make is to underestimate the importance of the Data Link layer and of VLANs in particular, in the sophisticated architecture of switched networks. It should not be forgotten that the OSI stack is only as robust as its weakest link,

and that therefore an equal amount of attention should be paid to any of its layers so as to make sure that its entire structure is sound.

Any good networking design based on Cisco Catalyst switches should incorporate the best practice guidelines described here as an effective way to protect a network's L2 security architecture from dangerous vulnerabilities.

3.9.1 Basic Security

Any attempt to create a secure switched network starts from basic security principles. And in particular, basic rules such as the ones highlighted in the SAFE best practices are the cornerstone of any design of secure switched networks.

If a user does not want one of his or her devices to be tampered with, physical access to the device must be strictly controlled. Furthermore, it is important for any network administrator to use all the proven security tools available on Cisco platforms: from the very basic configuration of system passwords, the use of IP permit filters, and login banners, all the way to more advanced tools such as RADIUS, TACACS+, Kerberos, SSH, SNMPv3, IDS, and so forth.

3.9.2 VLAN-based security

A Layer 2 (L2) switch is a device capable of grouping subsets of its ports into virtual broadcast domains isolated from each other. These domains are commonly known as virtual LANs (VLANs).The concept of VLAN is akin to other concepts in the networking world where traffic is identified by the use of a tag or label. Identification is crucial for a L2 device to be able to isolate ports and properly forward the traffic received. As we will see later, lack of identification is sometimes a cause of insecurity and needs to be avoided.

If any packet in a device is tightly coupled to an appropriate VLAN tag, it is always possible to reliably discriminate traffic into separate and independent domains. This is the basic premise of VLAN-based switching architectures.

In particular, Cisco devices work in accordance with popular VLAN tagging technologies like ISL or 802.1Q across physical links (sometimes referred to as trunks) and employ

advanced tagging techniques to preserve the VLAN information internally and use it for the purpose of traffic forwarding.

The simple observation that can be made at this point is that if a packet's VLAN identification cannot be altered after transmission from its source and is consistently preserved from end to end, then VLAN-based security is no less reliable than physical security.

3.9.3 Control Plane

Malicious users often seek to gain access to the management console of a networking device, because if they are successful they can easily alter the network configuration to their advantage.

In a VLAN-based switch, in addition to having a direct connection to an out-of-band port, the management CPU can use one or more VLANs for in-band management purposes. It also uses one or more VLANs to exchange protocol traffic with other networking devices.

As basic physical security guidelines require networking equipment to be in a controlled (locked) space, VLAN-based security's primary rule is to confine in-band management and protocol traffic into a controlled environment.

This can be achieved with the following tools and best practices:

Traffic and protocol ACLs or filters.

1. Quos marking and prioritization (control protocols are differentiated by means of appropriate class-of-service or DSCP values).
2. Selective deactivation of L2 protocols on entrusted ports (for example, disabling DTP on access ports).
3. Configuration of in band management port(s) only in dedicated VLAN(s).
4. Abstention from using VLAN 1 to carry any data traffic.

3.9.4 Precautions for the Use of VLAN 1

The reason VLAN 1 became a special VLAN is that L2 devices needed to have a default VLAN to assign to their ports, including their management port(s). In addition to that, many L2 protocols such as CDP, PAGP, and VTP needed to be sent on a specific VLAN on trunk links. For all these purposes VLAN 1 was chosen.

As a consequence, VLAN 1 may sometimes end up unwisely spanning the entire network if not appropriately pruned and, if its diameter is large enough, the risk of instability can increase significantly. Besides the practice of using a potentially omnipresent VLAN for management purposes puts trusted devices to higher risk of security attacks from entrusted devices that by miss configuration or pure accident gain access to VLAN 1 and try to exploit this unexpected security hole.

To redeem VLAN 1 from its bad reputation, a simple common-sense security principle can be used: as a generic security rule the network administrator should prune any VLAN, and in particular VLAN 1, from all the ports where that VLAN is not strictly needed.

Therefore, with regard to VLAN 1, the above rule simply translates into the recommendations to:

1. Not use VLAN 1 for in band management traffic and pick a different, especially dedicated VLAN that keeps management traffic separate from user data and protocol traffic.
2. Prune VLAN 1 from all the trunks and from all the access ports that don't require it (including not connected and shutdown ports).

Similarly, the above rule applied to the management VLAN reads:

1. Don't configure the management VLAN on any trunk or access port that doesn't require it (including not connected and shutdown ports).
2. For foolproof security, when feasible, prefer out-of-band management to in band management.

As a general design rule it is desirable to "prune" unnecessary traffic from particular VLANs. For example, it is often desirable to apply VLAN ACLs and/or IP filters to the traffic carried in the management VLAN to prevent all telnet connections and allow only SSH sessions. Or it may be desirable to apply QOS ACLs to rate limit the maximum amount of ping traffic allowed.

If VLANs other than VLAN 1 or the management VLAN represent a security concern, then automatic or manual pruning should be applied as well. In particular, configuring VTP in transparent or off mode and doing manual pruning of VLANs is commonly considered the most effective method to exert a more strict level of control over a VLAN-based network.

3.9.5 Secure environments of VLAN 1

After proper handling of VLAN 1 has been decided upon and implemented, the next logical step is to turn one's attention to other equally important best practices commonly used in secure environments. The generic security principle applied here is: connect entrusted devices to entrusted ports, trusted devices to trusted ports, and disable all the remaining ports. What this means can be easily expanded into this list of common recommendations:

1. If a port is connected to a "foreign" device, don't try to speak any language with it: it could be turned to somebody else's advantage and used against you. So on that port makes sure to disable CDP, DTP, PAGP, UDLD, and any other unnecessary protocol, and to enable port fast/BPDU guard on it. After all, why risk a potentially dangerous communication with an untrustworthy neighbor?
2. Enable the root guard feature to prevent a directly or indirectly connected STP-capable device to affect the location of the root bridge.
3. Configure the VTP domains appropriately or turn off VTP altogether if you want to limit or prevent possible undesirable protocol interactions with regard to network-wide VLAN configuration. This precaution can limit or prevent the risk of an administrator error propagating to the entire network and the risk of a new switch with a higher VTP revision overwriting by accident the entire domain's VLAN configuration.

4. By default only those ports which are known to be `trusted' should be treated as such and all other ports should be configured as `entrusted'. This prevents attached devices from manipulating QOS values inappropriately.
5. Disable unused ports and put them in an unused VLAN. By not granting connectivity or by placing a device into a VLAN not in use, unauthorized access can be thwarted through fundamental physical and logical barriers.

3.9.6 The Layer 2 Security

The OSI stack was conceived so that different layers are able to function independently (with only the knowledge of their mutual interfaces). This allows for flexibility in that developments for a given layer of the protocol stack do not impact other layers so long as the standard interface between the layers is maintained.

Unfortunately this also means that if one layer is hacked, communication may be compromised without the other layers being aware of the problem (as shown in Figure 3.6).

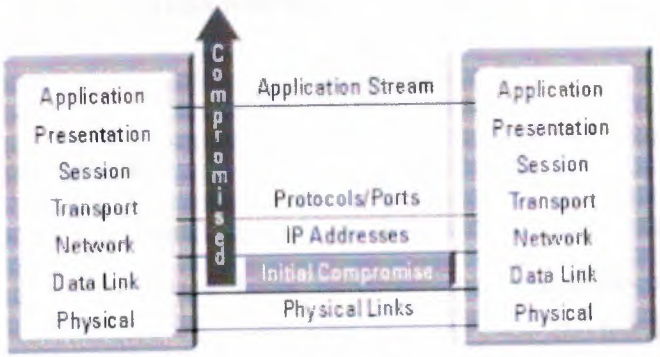

Figure3.6 OSI Stack Structure

In this architecture, security is only as strong as its weakest link.

The Data Link layer is as vulnerable as any other layer and can be subjected to a variety of attacks which the switch must be configured to protect against.

3.9.7 Attacks in a VLAN-Based Network

The majority of attacks at L2 exploit the inability of a device to track the attacker who can therefore perform undetected malicious actions on the forwarding path to alter it and then exploit the change.

These are the most talked-about L2 attacks and incidentally:

1. MAC Flooding Attack
2. 802.1Q and ISL Tagging Attack
3. Double-Encapsulated 802.1Q/Nested VLAN Attack
4. ARP Attacks
5. Private VLAN Attack
6. Multicast Brute Force Attack
7. Spanning-Tree Attack
8. Random Frame Stress Attack

A description of each of these threats follows.

3.9.7.1 MAC Flooding Attack

This is not properly a network "attack" but more a limitation of the way all switches and bridges work. They possess a finite hardware learning table to store the source addresses of all received packets: when this table becomes full, the traffic that is directed to addresses that cannot be learned anymore will be permanently flooded. Packet flooding however is constrained within the VLAN of origin; therefore no VLAN hopping is permitted.

This corner case behavior can be exploited by a malicious user that wants to turn the switch he or she is connected to into a dumb pseudo-hub and sniff all the flooded traffic. Several programs are available to perform this task: for example macof, part of the dsniff suite [4]. This weakness can then be exploited to perform an actual attack, like the ARP poisoning attack (see ARP Attacks for more details on the subject).

On non intelligent switches this problem arises because a sender's L2 identity is not checked, therefore the sender is allowed to impersonate an unlimited number of devices simply by counterfeiting packets.

Cisco's switches support a variety of features whose only goal is to identify and control the identities of connected devices. The security principle on which they are based is very simple: authentication and accountability are critical for all entrusted devices.

In particular, Port Security, 802.1x, and Dynamic VLANs are three features that can be used to constrain the connectivity of a device based on its user's login ID and based on the device's own MAC layer identification.

With Port Security, for instance, preventing any MAC flooding attack becomes as simple as limiting the number of MAC addresses that can be used by a single port: the identification of the traffic of a device is thereby directly tied to its port of origin.

3.9.7.2 802.1Q and ISL Tagging Attack

Tagging attacks are malicious schemes that allow a user on a VLAN to get unauthorized access to another VLAN. For example, if switch ports were configured as DTP auto and were to receive a fake DTP packet, it might become a trunk port and it might start accepting traffic destined for any VLAN. Therefore, a malicious user could start communicating with other VLANs through that compromised port.

Sometimes, even when simply receiving regular packets, a switch port may behave like a full-fledged trunk port (for example, accept packets for VLANs different from the native), even if it is not supposed to. This is commonly referred to as "VLAN leaking"

While the first attack can be prevented very easily by setting DTP to off on all non trusted ports (again the principle of trust at work), the second attack can usually be addressed by following simple configuration guidelines (such as the one suggested in the next section) or with software upgrades. Fortunately, Cisco Catalyst 2950, Catalyst 3550, Catalyst 4000, and Catalyst 6000 series switches don't need any such upgrade, since their software and

hardware have been designed to always enforce proper traffic classification and isolation on all their ports.

3.9.7.3 Double-Encapsulated 802.1Q/Nested VLAN Attack

While internal to a switch, VLAN numbers and identification are carried in a special extended format that allows the forwarding path to maintain VLAN isolation from end to end without any loss of information. Instead, outside of a switch, the tagging rules are dictated by standards such as ISL or 802.1Q.

ISL is a Cisco proprietary technology and is in a sense a compact form of the extended packet header used inside the device: since every packet always gets a tag, there is no risk of identity loss and therefore of security weaknesses.

On the other hand, the IEEE committee that defined 802.1Q decided that because of backward compatibility it was desirable to support the so-called native VLAN, that is to say, a VLAN that is not associated explicitly to any tag on an 802.1Q link. This VLAN is implicitly used for all the untagged traffic received on an 802.1Q capable port.

This capability is desirable because it allows 802.1Q capable ports to talk to old 802.3 ports directly by sending and receiving untagged traffic. However, in all other cases, it may be very detrimental because packets associated with the native VLAN lose their tags, for example, their identity enforcement, as well as their Class of Service (802.1p bits) when transmitted over an 802.1Q link.

For these sole reasons—loss of means of identification and loss of classification—the use of the native VLAN should be avoided. There is a more subtle reason, though. Figure 3.7 shows why.

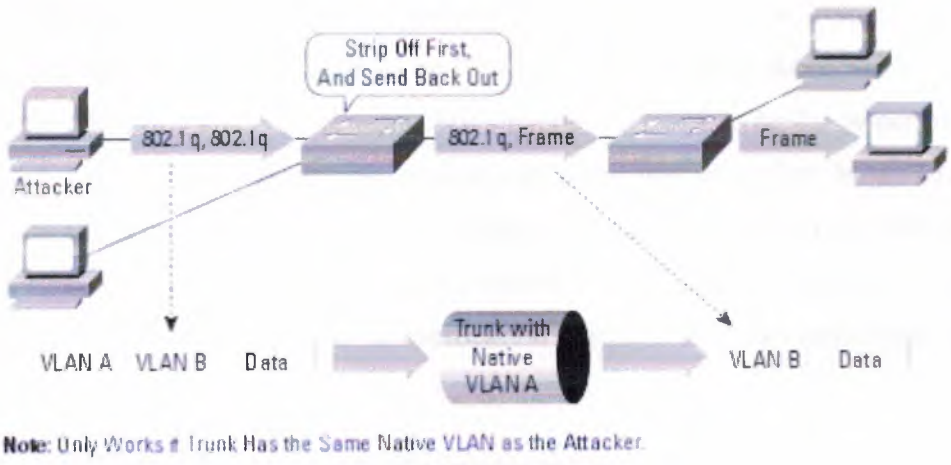

Figure3.7 Double Encapsulation Attack

When double-encapsulated 802.1Q packets are injected into the network from a device who's VLAN happens to be the native VLAN of a trunk, the VLAN identification of those packets cannot be preserved from end to end since the 802.1Q trunk would always modify the packets by stripping their outer tag. After the external tag is removed, the internal tag permanently becomes the packet's only VLAN identifier. Therefore, by double-encapsulating packets with two different tags, traffic can be made to hop across VLANs.

This scenario is to be considered a miss configuration, since the 802.1Q standard does not necessarily force the users to use the native VLAN in these cases. As a matter of fact, the proper configuration that should always be used is to clear the native VLAN from all 802.1Q trunks (alternatively, setting them to 802.1q-all-tagged mode achieves the exact same result). In cases where the native VLAN cannot be cleared, then always pick an unused VLAN as native VLAN of all the trunks; don't use this VLAN for any other purpose. Protocols like STP, DTP, and UDLD (check out [3]) should be the only rightful users of the native VLAN and their traffic should be completely isolated from any data packets.

3.9.7.4 ARP Attacks

The ARP protocol is quite an old technology. The ARP RFC is from a time when everyone in a network was supposed to be "friendly" and therefore there was no security built into the ARP function. As a consequence, anyone can claim to be the owner of any IP address they like. To be more precise, anyone can claim that his or her MAC address is associated to any IP address within a specific subnet. This is possible because ARP requests or replies carry the information about the L2 identity (MAC address) and the L3 identity (IP address) of a device and there is no verification mechanism of the correctness of these identities.

Again, this is another case where lack of a precise and reliable means of identification of a device leads to a serious security vulnerability. Also, this is a perfect example of why by compromising a lower level in the OSI stack it's possible to directly affect an upper level without the upper layer being aware of the problem. (ARP is a unique specimen of protocol living and breathing in the L2 world but logically residing at the boundary between the Data Link and the Network layer in the OSI stack.)

The ARP attacks that @stake performed were targeted to fool a switch into forwarding packets to a device in a different VLAN by sending ARP packets containing appropriately forged identities. However, in all Cisco devices VLANs are orthogonal to and therefore independent from MAC addresses: so by changing a device's identity in an ARP packet, it's not possible to affect the way it communicates with other devices across VLANs. As a matter of fact, as the report states, any VLAN hopping attempt was thwarted.

On the other hand, within the same VLAN, the so-called ARP poisoning or ARP spoofing attacks [7] are a very effective way to fool end stations or routers into learning counterfeited device identities: this can allow a malicious user to pose as intermediary and perform a Man-In-the-Middle (MiM) attack.

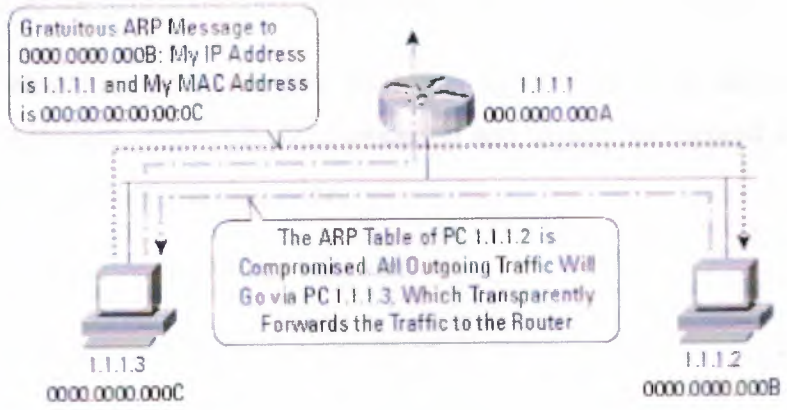

Figure3.8 ARP Poisoning Attack

In this case, a picture is worth more than a thousand words of explanation (see Figure 3.8).

The MiM attack is performed by impersonating another device (for example, the default gateway) in the ARP packets sent to the attacked device: these packets are not verified by the receiver and therefore they "poison" its ARP table with forged information.

This type of attack can be prevented either by blocking the direct communication at L2 between the attacker and the attacked device or by embedding more intelligence into the network so that it can check the forwarded ARP packets for identity correctness. The former countermeasure can be achieved with Cisco Catalyst Private VLANs or Private VLAN Edge features. The latter can be achieved by using a new feature called ARP Inspection, available first in Cat OS 7.5 on the Cisco Catalyst 6500 Supervisor Engine II and a little later also in the Cisco IOS Software for the Cisco Catalyst switches.

3.9.7.5 Private VLAN Attack

Private VLAN attack is actually a misnomer because it corresponds not to vulnerability but rather to the expected behavior of the feature. Private VLANs is a L2 feature and therefore it is supposed to isolate traffic only at L2. On the other hand, a router is a Layer 3 (L3) device and when it's attached to a Private VLAN promiscuous port it is supposed to

forward L3 traffic received on that port to whatever destination it is meant to, even if it's in the same subnet as the source (at stake refers to this behavior as Layer-2 Proxy).

Therefore, it is absolutely normal for two hosts in an Isolated VLAN to fail to communicate with each other through direct L2 communication and instead to succeed to talk to each other by using the router as a packet relay.

Figure 3.9 depicts the aforementioned behavior.

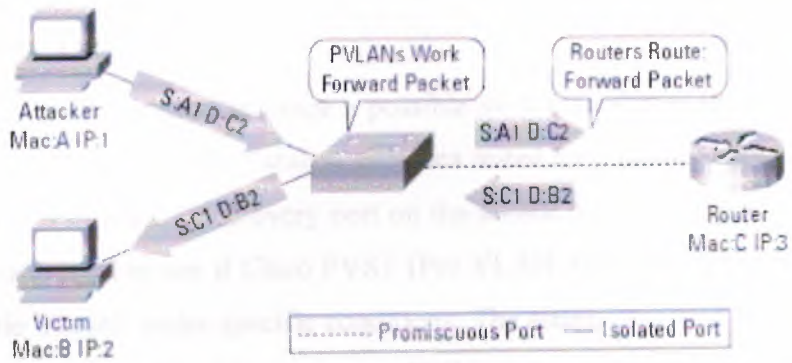

Figure3.9 L2 Proxy

As with regular routed traffic, packets relayed through L2 Proxy can be filtered, if desired, through the configuration of an appropriate ACL on the forwarding device.

Here is a simple example of output Cisco IOS ACL to block the relayed traffic:

Deny subnet/mask subnet/mask

Permit any subnet/mask

Deny any any

3.9.7.6 Multicast Brute Force Attack

This attack tries to exploit switches' potential vulnerabilities (read: bugs) against a storm of L2 multicast frames. @stake's test was designed to ascertain what happens when a

L2 switch receives lots of L2 multicast frames in rapid succession. The correct behavior should be to constrain the traffic to its VLAN of origin, the failure behavior would be to leak frames to other VLANs.

In @stake's results, this type of attack proved ineffective against Cisco Catalyst switches because they correctly contained all the frames within their appropriate broadcast domain (no surprise here: after all, in all Catalyst switches broadcasts are just special cases of multicasts).

3.9.7.7 Spanning-Tree Attack

Another attack that tries to leverage a possible switch weakness (for example, bug) is the STP attack. All of the Cisco Catalyst switches tested by @stake support this protocol. By default, STP is turned on and every port on the switch both speaks and listens for STP messages. @stake tried to see if Cisco PVST (Per VLAN Spanning Tree) would fail open across[1] multiple VLANs under specific conditions. The attack consisted in sniffing for STP frames on the wire to get the ID of the port STP was transmitting on. Next, the attacker would begin sending out STP Configuration/Topology Change Acknowledgement BPDUs announcing that he was the new root bridge with a much lower priority.

During this procedure broadcast traffic was injected by the testers to discover any possible VLAN leaks, but none were found. This is an indication of the robustness of STPs implementations on Cisco switches.

3.9.7.8 Random Frame Stress Attack

This last test can have many incarnations but in general it consists in a brute force attack that randomly varies several fields of a packet while keeping only the source and destination addresses constant. After repetitive testing by @stake's engineers, no packets were found to have successfully hopped VLANs.

Private VLANs can be used in this context to better isolate hosts at L2 and shield them from unwanted malicious traffic from untrustworthy devices. Communities of mutually-trusting hosts can be created so as to partition a L2 network into sub domains where only friendly devices are allowed to communicate with each other.

4. USE OF VIRTUAL LANs (VLANs) IN NETWORKS

4.1 use of virtual

The use of virtual LANs has already been established for a considerable amount of time in Enterprise networks, virtual LANs (VLANs) are now increasingly being implemented in the backbone networks of many service providers. The provider has a need to plan and divide his network to enable him to offer value added services on top of pure Internet connections. Virtual LANs supply the service provider with unlimited possibilities to segment his preferred network according to functionality, workgroups or service requirements. The service provider is then in a position to deliver a comprehensive service offering with no further investment. When the customer applies for DSL access they select their services. If desired the service may be upgraded at a later date with no visits to the customer premises (truck roll) being required.

4.2 VLANs in The DSL Access

Ethernet based networks are increasingly be deployed by carriers as they have been proven to be easier to configure and maintain than traditional core networks based on TDM or ATM technologies. This reduces costs in installation, maintenance and technical training. However an Ethernet network is by default a single network, which is not ideal for traffic control and management. A VLAN is a logically segmented (Ethernet) network. This divides the Ethernet network into various broadcast domains. This enables independent networks to be run in parallel over the same physical structure.

The key advantage of VLANs is that the network is divided into logical areas or groups. This maintains a high level of network security as only the members of a particular VLAN have access to one another and the server, i.e. a connection between one VLAN

and another is only possible through the higher levels of the protocol stack. This is at the discretion of the network administrator. A further and often under-estimated advantage of VLANs is the division of the network into different broadcast domains. Important resources and network areas are protected from unnecessary traffic and performance is dramatically increased.

4.3 VLANS in the Enterprise Network

Diagram 1 shows the structure of a business network divided by VLANs. The VLANs are categorised by department. The network designer took the approach that the majority of the traffic would be inter-departmental. For example, the purchasing department requires access to some common resources on the server, which colleagues in other departments only need on an occasional basis. The same situation is true of the purchasing and development departments. All employees use the existing Internet access. A similar architecture can be used by the service provider to provide access to gaming servers, VoIP servers, Portals, or even simply to direct traffic down preferred sections of the network based on Quality and conjestion requirements for the level of service being offered.

A corporate LAN could be divided up as below (figure 4.1):

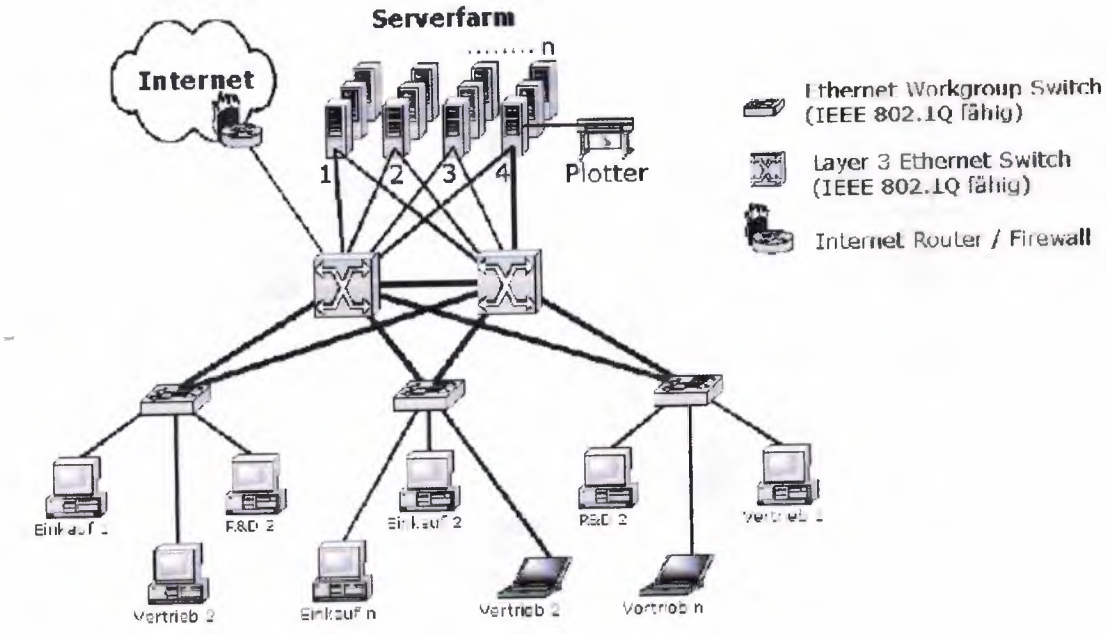

Figare4.1 Physical Structure of an Ethernet network by VLANs

4.4 Virtual LAN Configuration

VLAN IDs are assigned to various switch ports at the network administrator's discretion. This means a PC connected to port 2 on the switch may be a member of the same VLAN as a PC connected to ports 4.7 of that router and port 1 of another router. This is sometimes referred to as Port-based VLANs. Communication between different VLANs is only possible over a Layer 3 switch/router and consequently a high level protocol, such as IP. Every user port on the DSLAM could be assigned to a particular VLAN or number of VLANs according to services booked.

Diagram 2 shows the logical structure in a LAN environment. The VLANs depict three separate services (VLAN 11 – red, VLAN 12 – blue and VLAN 20 – green). A range of

sample IP addresses is also provided. This is necessary to allow communication between the configured VLANs at minimum expense and re-structuring (figure 4.2):

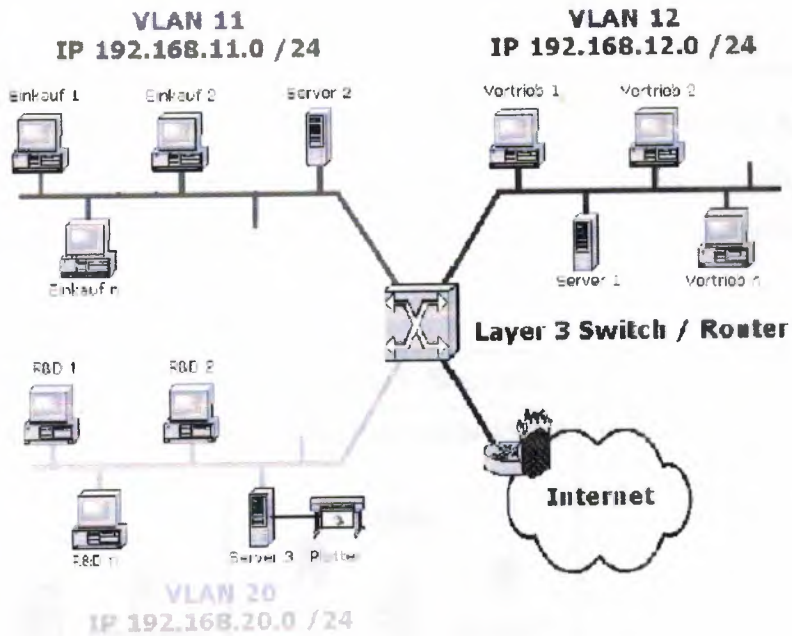

Figure4.2. Logical depiction of an Ethernet network divided by VLANs

4.5 Broadband Network Design

The following is an example configuration of a broadband access network. This network supports various service features and offers the customer a choice of services. It is assumed that the backbone network composes of switched or routed Ethernet and the Layer 3 protocol, IP. These components all support the VLAN standard IEEE 802.1Q and are able to route IP traffic over the so-called virtual routers. This example does not detail either the configuration of the switch/router or the dynamic routing protocols such as RIP, OSPF or BGP.

4.6 Requirements Of The Service Providers Network

The network should be able to serve both business and private customers within the same topology. Business customer should also have the possibility to connect up remote offices via a dedicated LAN extension service. Figure 4.3 shows network architecture already divided into VLANs with different services available in different VLANs.

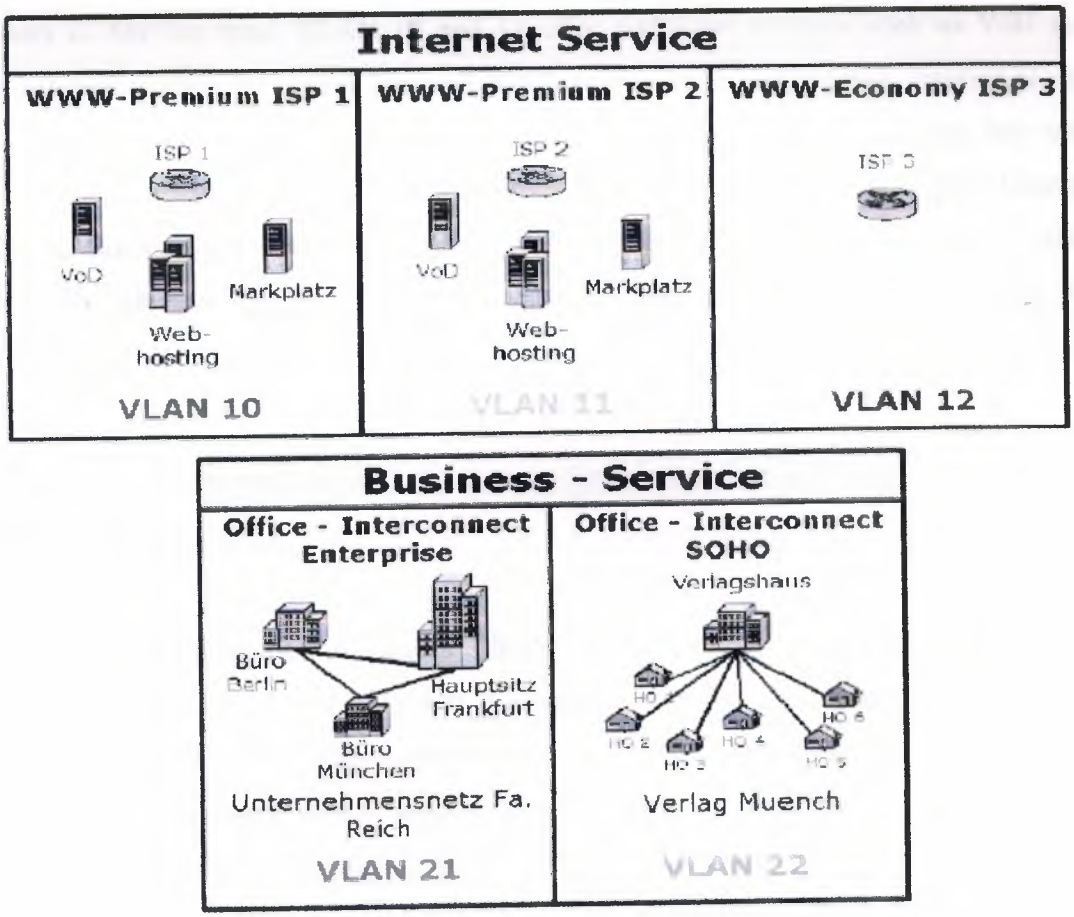

Figure4.3. Logical arrangement of the network

Note the relationship between service and VLAN ID, i.e. each service corresponds to one of altogether 4095 possible VLANs.

4.7 Internet service

The upper group shows public services, i.e. a new customer (in this example of an Internet service) is added to the created VLAN. Three different services are available, which are defined under the name WWW Premium ISP1, WWW Premium ISP2 and WWW Economy ISP3. For these services applicable technologies are ADSL and SDSL. If enabled the customer could not only have the ability to choose between different Internet Service Providers (ISPs), but also to choose how important to them the Internet access is. On one hand VLAN 10 and 11 offer additional services such as VoD and various web portals for on-line business and have improved performance. Internet access is available from providers ISP1 and ISP2. VLAN 12 also provides access but with reduced performance from the Economy provider. The service offered is pure Internet access, which may not have high performance unlike the connections to the two other ISPs. The services offered will have to be taken into account when determining the pricing structure.

4.8 Business service

The second group shown in Figure 3 offers private services. Here a direct relationship between VLAN ID and customer ID exists. This means that for each customer who rents one of these services, a dedicated VLAN is created. This is therefore a private network into which nobody else can transmit or receive data. As already mentioned, the VLAN technology offers a certain basic security that is sufficient for many customers. This is achieved through the 802.1Q protocol, which provides separate broadcast domains for each VLAN. If a higher safety class is required, then there is the possibility to encode the transferred data. From a technical viewpoint a similar approach is taken when connecting up Small Offices or Home Offices, by the service Office

Interconnect SOHO. Home-workers or small external offices have the possibility of a fast data link to the enterprise network. Here also, a single VLAN ID per customer is used.

4.9 Speed

Until now we have not mentioned the speed of each access method. This is because the speed is not affected by the use of VLANs, on net to net Technologies DSLAM all the access link ports can be VLAN enabled ADSL, ADSL2, ADSL2+, SHDSL and E1/T1 ports single or bonded. However the VLANs could be used to select different contention rations on the uplink ports by putting different access ports in different uplink VLANs to give quality of service to key customers.

4.10 The Core Network

Central to the network is the IP Router that supports the IEEE 802.1Q standard, this serves as the core of the network and if required can provide connections between different VLANs.

Interconnecting takes place over conventional Layer 3 Switching or Routing.

CAUTION: When connecting private and public services a Firewall should be connected between the private and public networks. It protects the enterprise network against unauthorized access.

Example: If the enterprise network of the company realm is to be tied up to the Internet connection of ISP 1, then the private enterprise network (VLAN 21) is connected via the central IP Router to the VLAN 10 (ISP 1). Without merging a Firewall the enterprise network will be visible from the public network (figure 4.4):

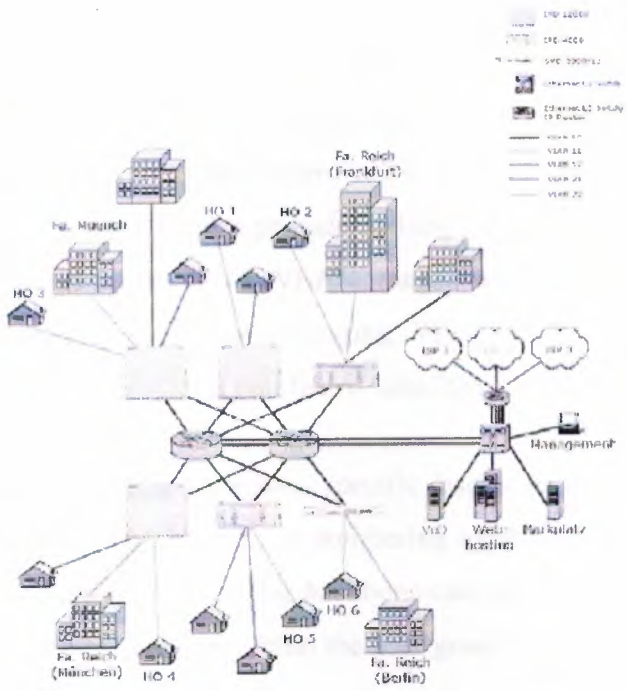

Figure 4.4. Physical Structure of the Network

4.11 Access Components

The IP DSLAMs are placed within the area of co-location. In fig. 4 these are the carrier class solutions based on the Net to Net Technologies IPD12000E (14 Slot modular DSLAM) and IPD 4000E (5 Slot modular DSLAM) but it could equally well be the Mini or Micro fixed port DSLAM products. Also in the application shown is an SMD2000-12, which is a standalone 12 port SDSL MiniDSLAM.

4.12 VLAN Stacking

The VLAN Stacking feature provides a scalability factor to the 802.1Q VLAN capability in the Net to Net Technologies IP DSLAM. The 802.1Q standard for VLANS supports 4096 VLANs per network. This low number limitation means that a scalability

issue is created in large IP based DSL networks, in particular where there is a need to allocate groups of VLAN numbers to service providers. With VLAN stacking Net to Net Technologies increases that number to VLANs able to be created in a backbone VLAN network to a total of 16,777,216 (4096 x 4096) still within the implementation of 802.1Q. VLAN stacking allows a DSL service provider selling access on their network to other providers to allocate an entire 4096 VLAN range, with maximum flexibility and expandability and avoiding issues of managing VLAN number ranges. A service provider can use VLANs to provide a number of benefits to the customer.

1. They can be used to create company specific backbone VLANs. This allows a company to use their own 802.1Q VLAN numbering system within the enterprise and still traverse the provider's backbone. This backbone can link between multiple users and sites and provides additional security within the user group and can include access to the Internet through a secured firewall facility.

2. A VLAN can be allocated to create dedicated lease line emulation.

3. They can be used to create groups or communities of interest for a particular service. For example a games server could sit on a particular VLAN and when a user subscribes to that service the service provider could subscribe their DSL port into the relevant VLAN to give them access to the service. This means that providers can very quickly and easily add users to new service groups.

4. VLAN allocations can include the DSLAM backhaul to create specific levels of backhaul service availability to users with Service Level Agreements offering specific level of bandwidth availability.

4.13 Net to Net Technologies use of VLANs

The VLAN technology implemented in Net to Net Technologies products provides the ability to provide customers with a full spectrum of access services beyond simple Internet access. This is done without any investment in additional equipment such as a Service Selection Gateway. The simple configuration and maintenance characteristic of the components reduce not only the installation and operating cost, but provides the Service Provider with the ability to react quickly to customer requirements. The Service Provider will thus always be able to remain one step ahead of its competitors.

CONCLUSION

This Project presents inclusive information for implementing the Virtual Area Networks (VLAN's). Virtual Area Networks are today's need for improving the known networks to suite this high speed developing area of technology. Virtual Local Area Networks (VLAN's) were developed as an alternative solution to using routers to contain broadcast traffic.

Virtual Local Area Networks can be viewed as a group of devices on different physical LAN segments which can communicate with each other as if they were all on the same physical LAN segment. However, In order to take advantage of the benefits of Virtual Local Area Networks, a different network topology is needed.

In this project I have defined VLAN's and examined the difference between a LAN and a VLAN. This is followed by a discussion on the advantages VLAN's introduce to a network. Finally, I explained how VLAN's work based on the current draft standards. While bandwidth may be a reason big enough to go for switching, Virtual LAN (VLAN) Support may also be attractive. A VLAN is logical grouping of ports into workgroups. With VLAN support network managers can define workgroups independent of underlying network topology, VLANs are becoming popular because of the flexibility they offer, Users can physically move but stay on the same VLAN.

REFERENCES

[1] Tanebaum Andrew S., *Computer Networks*, 1996

[2]Martin Michael J., *Understanding the Network: A Practical Guide to Internetworking*, Macmillan Computer publishing, USA, 2000

[3] Microsoft, *Networking Essentials*, Microsoft Corporation, Washington, 1996

[4] Mahler, Kevin. *CCNA Training Guide*. Indianapolis: New Riders, 1999.

[5] *Cisco IOS Wide Area Networking Solutions*. Indianapolis: Cisco Press, 1999.

[6] http://www.cisco.com/warp.

[7] http://www.sans.org/newlook/resources/IDFAQ/vlan.htm.