decisys

The Virtual LAN
Technology Report

The Virtual LAN Technology Report

Contents

The Virtual LAN Technology Report

by David Passmore and John Freeman

David Passmore is president and co-founder of Decisys, Inc., a Sterling, Virginia–based consulting firm specializing in network design, architecture, and management for end-user organizations, and network product marketing and strategic planning for vendors. Before founding Decisys, David was vice president of the Gartner Group and a partner in Ernst & Young's Center for Information Technology and Strategy in Boston, Massachusetts.

David received a B.S. in computer science and engineering and an M.S. in electrical engineering and computer science, both from the Massachusetts Institute of Technology.

Introduction

Virtual LANs (VLANs) have recently developed into an integral feature of switched LAN solutions from every major LAN equipment vendor. Although end-user enthusiasm for VLAN implementation has yet to take off, most organizations have begun to look for vendors that have a well-articulated VLAN strategy, as well as VLAN functionality built into products today. One of the reasons for the attention placed on VLAN functionality now is the rapid deployment of LAN switching that began in 1994/1995.

The shift toward LAN switching as a replacement for local/departmental routers—and now even shared media devices (hubs)—will only accelerate in the future. With the rapid decrease in Ethernet and Token Ring switch prices on a per-port basis, many more ambitious organizations are moving quickly toward networks featuring private port (single user/port) LAN switching architectures. Such a desktop switching architecture is ideally suited to VLAN implementation. To understand why private port LAN switching is so well suited to VLAN implementation, it is useful to review the evolution of segmentation and broadcast containment in the network over the past several years.

In the early 1990s, organizations began to replace two-port bridges with multiport, collapsed backbone routers in order to segment their networks at layer 3 and thus also contain broadcast traffic. In a network using only routers for segmentation, segments and broadcast domains correspond on a one-to-one basis. Each segment typically contained between 30 and 100 users.

With the introduction of switching, organizations were able to divide the network into smaller, layer 2–defined segments, enabling increased bandwidth per segment. Routers could now focus on providing broadcast containment, and broadcast domains could now span multiple switched segments, easily supporting 500 or more users per broadcast domain. However, the continued deployment of switches, dividing the network into more and more segments (with fewer and fewer users per segment) does not reduce the need for broadcast containment. Using routers, broadcast domains typically remain in the 100 to 500 user range.

VLANs represent an alternative solution to routers for broadcast containment, since VLANs allow switches to also contain broadcast traffic. With the implementation of switches in conjunction with VLANs, each network segment can contain as few as one user (approaching private port LAN switching), while broadcast domains can be as large as 1,000 users or perhaps even more. In addition, if implemented properly, VLANs can track workstation movements to new locations without requiring manual reconfiguration of IP addresses.

Why haven't more organizations deployed VLANs? For the vast majority of end-user organizations, switches have yet to be implemented on a large enough scale to necessitate VLANs. That situation will soon change. There are, however, other reasons for the lukewarm reception that VLANs have received from network users up to now:

- VLANs have been, and are still, proprietary, single-vendor solutions. As the networking industry has shown, proprietary solutions are anathema to the multivendor/open systems policies that have developed in the migration to local area networks and the client server model.
- Despite the frequently quoted numbers illuminating the hidden costs of networking, such as administration and moves/adds/changes, customers realize that VLANs have their own administrative costs, both straightforward and hidden.
- Although many analysts have suggested that VLANs enhance the ability to deploy centralized servers, customers may look at enterprise-wide VLAN implementation and see difficulties in enabling full, high-performance access to centralized servers.

This paper discusses these and other issues in greater detail, and attempts to determine the strategic implications that VLANs, present and future, pose for enterprise networks.

Defining VLANs

What is a 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.

However, at this point, issues such as the extent to which end-stations are not constrained by physical location, the way VLAN membership is defined, the relationship between VLANs and routing, and the relationship between VLANs and ATM have been left up to each vendor. To a certain extent these are tactical issues, but how they are resolved has important strategic implications.

Because there are several ways in which VLAN membership can be defined, this paper divides VLAN solutions into four general types: port grouping, MAC-layer grouping, network-layer grouping, and IP multicast grouping. We will discuss the issue of manual vs. automatic VLAN configuration, and describe techniques by which VLANs may be extended across multiple switches in the network. Finally, the paper takes a look at the present state of VLAN standards.

Membership by Port Group

Many initial VLAN implementations defined VLAN membership by groups of switch ports (for example, ports 1, 2, 3, 7, and 8 on a switch make up VLAN A, while ports 4, 5, and 6 make up VLAN B). Furthermore, in most initial implementations, VLANs could only be supported on a single switch.

Second-generation implementations support VLANs that span multiple switches (for example, ports 1 and 2 of switch #1 and ports 4, 5, 6, and 7 of switch #2 make up VLAN A; while ports 3, 4, 5, 6, 7, and 8 of switch #1 combined with ports 1, 2, 3, and 8 of switch #2 make up VLAN B). This scenario is depicted in Figure 1.

Port grouping is still the most common method of defining VLAN membership, and configuration is fairly straightforward. Defining VLANs purely by port group does not allow multiple VLANs to include the same physical segment (or switch port). However, the primary limitation of defining VLANs by port is that the network manager must reconfigure VLAN membership when a user moves from one port to another.

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.

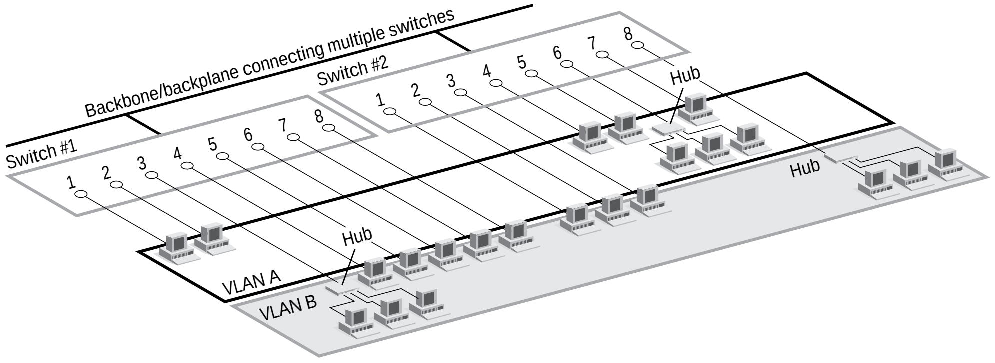

Figure 1. *VLANs Defined by Port Group*

John Freeman is a senior consultant at Decisys, Inc., where he specializes in the development of technology marketing and vendor strategies. John also works with end-user clients to help them understand and evaluate emerging technologies and vendor strategies. Before joining Decisys, John worked as a consultant in Japan in the areas of networking and systems integration. He is fluent in Japanese and is an expert in the Japanese networking market.

John holds a B.A. in East Asian Studies from Harvard University.

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 remain on the desktop, while the notebook travels with the user. When the user moves to a new desk and docking station, the MAC-layer address changes, making VLAN membership impossible to track. 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.

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

Acronyms and Abbreviations

AAL5
ATM Adaptation Layer Type 5

ASIC
Application-specific integrated circuit

ATM
Asynchronous Transfer Mode

DHCP
Dynamic Host Configuration Protocol

ELAN
Emulated LAN

FDDI
Fiber Distributed Data Interface

IPX
Internet Packet Exchange

LANE
LAN Emulation

LEC
LAN Emulation client

LES
LAN Emulation server

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 “unroutable” protocols such as NetBIOS. End-stations running unroutable protocols cannot be differentiated and thus cannot be defined as part of a network-layer VLAN.

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.

The dynamic nature of VLANs defined by IP multicast groups enables a very high degree of flexibility and application sensitivity.

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.

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:

- *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.
- *Semiautomated.* Semiautomated configuration refers to the option to automate either 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. Semiautomated configuration could also refer to

Acronyms and Abbreviations (Cont.)

MAC
Media access control

MPOA
Multiprotocol over ATM

NIC
Network interface card

OSPF
Open Shortest Path First

PVC
Permanent virtual circuit

RIP
Routing Information Protocol

SVC
Switched virtual circuit

TCP/IP
Transmission Control Protocol/Internet Protocol

TDM
Time-division multiplexing

situations where VLANs are initially configured manually, with all subsequent moves being tracked automatically. Combining both initial and subsequent configuration automation would still imply semiautomated configuration, because the network administrator always has the option of manual configuration.

- *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 paper.

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, communication of VLAN membership may also be implicit in the case of layer 3–based VLANs in a multiprotocol environment.

To date, outside of implementing an ATM backbone, three methods have been implemented for interswitch communication of VLAN information across a backbone: table maintenance via signaling, frame tagging, and time-division multiplexing (TDM).

- *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.
- *Frame Tagging.* In the frame-tagging approach, a header is typically inserted into each frame on interswitch 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.
- *TDM.* The third, and least utilized method, is time-division multiplexing. TDM works the same way on the interswitch 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 paper. 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.

Standards and the Proprietary Nature of VLANs

Given the variety of types of VLAN definitions and the variety of ways that switches can communicate VLAN information, it should not be surprising that each vendor has developed its own unique and proprietary VLAN solutions and products. The fact that switches from one vendor will not interoperate entirely with VLANs from other vendors may force customers to buy from a single vendor for VLAN deployment across the enterprise. An exception to this rule arises when VLANs are implemented in conjunction with an ATM backbone and LANE. This is discussed further in “VLANs and ATM,” later in this paper.

The fact that single-vendor VLAN solutions in the LAN backbone will be the rule for the foreseeable future contributes to the recommendation that VLANs should not be deployed indiscriminately throughout the enterprise. It also implies that purchase decisions should be more highly centralized or coordinated than they may traditionally have been. Thus, from both a procurement and a technological perspective, VLANs should be considered as elements of a strategic approach.

> The standardized format for frame tagging, known as 802.1Q, represents a major milestone in enabling VLANs to be implemented using equipment from several vendors.

The following two VLAN standards have been proposed:

- *802.10 “VLAN Standard.”* In 1995, Cisco Systems proposed the use of IEEE 802.10, which was originally established to address LAN security for VLANs. Cisco attempted to take the optional 802.10 frame header format and “reuse” it to convey VLAN frame tagging instead of security information. Although this can be made to work technically, most members of the 802 committee have been strongly opposed to using one standard for two discrete purposes. In addition, this solution would be based on variable-length fields, which make implementation of ASIC-based frame processing more difficult and thus slower and/or more expensive.
- *802.1 Internetworking Subcommittee.* In March, 1996, the IEEE 802.1 Internetworking Subcommittee completed the initial phase of investigation for developing a VLAN standard, and passed resolutions concerning three issues: the architectural approach to VLANs; a standardized format for frame tagging to communicate VLAN membership information across multiple, multivendor devices; and the future direction of VLAN standardization. The standardized format for frame tagging, in particular, known as 802.1Q, represents a major milestone in enabling VLANs to be implemented using equipment from several vendors, and will be key in encouraging more rapid deployment of VLANs. Furthermore, establishment of a frame format specification will allow vendors to immediately begin incorporating this standard into their switches. All major switch vendors, including 3Com, Alantec/FORE, Bay Networks, Cisco, and IBM voted in favor of this proposal.

However, due to the lag time necessary for some vendors to incorporate the frame format specification and the desire on the part of most organizations to have a unified VLAN management platform, VLANs will, in practice, continue to retain characteristics of a single-vendor solution for some time. This has significant ramifications for deployment and procurement of VLANs. Department-level procurement for LAN equipment, particularly in the backbone, is not practical for organizations deploying VLANs. Purchasing decisions and standardization on a particular vendor’s solution throughout the enterprise will become the norm, and price-based product competition will decrease. The structure of the industry itself may also shift in favor of the larger networking vendors that can furnish a complete solution across a wide range of components.

VLAN Implementation Benefits

Why are vendors paying so much attention to VLAN implementation? Will VLANs solve all of the network manager’s problems with respect to moves, changes, broadcasts, and performance?

Reducing the Cost of Moves and Changes

The reason most often given for VLAN implementation is a reduction in the cost of handling user moves and changes. Since these costs are quite substantial, this argument for VLAN implementation can be compelling.

Many venders are promising that VLAN implementation will result in a vastly increased ability to manage dynamic networks and realize substantial cost savings. This value proposition is most valid for IP networks. Normally, when a user moves to a different subnet, IP addresses must be manually updated in the workstation. This updating process can consume a substantial amount of time that could be used for more productive endeavors such as developing new network services. VLANs eliminate that hassle, because VLAN membership is not tied to a workstation's location in the network, allowing moved workstations to retain their original IP addresses and subnet membership.

It is certainly true that the phenomenon of increasingly dynamic networks absorbs a substantial portion of the budgets of most IS departments. However, not just any VLAN implementation will reduce these costs. VLANs themselves add another layer of virtual connectivity that must be managed in conjunction with physical connectivity. This is not to say that VLANs cannot reduce the costs of moves, and changes—if properly implemented, they will. However, organizations must be careful not to simply throw VLANs at the network, and they must make sure that the solution does not generate more network administration than it saves.

Virtual Workgroups

One of the more ambitious VLAN objectives is the establishment of the virtual workgroup model. The concept is that, with full VLAN implementation across the campus network environment, members of the same department or section can all appear to share the same "LAN," with most of the network traffic staying within the same VLAN broadcast domain. Someone moving to a new physical location but remaining in the same department could move without having workstations reconfigured. Conversely, a user would not have to change his or her physical location when changing departments—the network manager would simply change the user's VLAN membership.

This functionality promises to enable a more dynamic organizational environment, enhancing the recent trend toward cross-functional teams. The logic of the virtual workgroup model goes like this: teams formed on a temporary, project basis could be virtually connected to the same LAN without requiring people to physically move in order to minimize traffic across a collapsed backbone. Additionally, these workgroups would be dynamic: VLANs corresponding to these cross-functional project teams could be set up for the duration of the project and torn down when the project was completed, all the while allowing users to remain in the same physical locations.

Although this scenario seems attractive, the reality is that VLANs alone cannot pave the way for full utilization of the virtual workgroup model. There are several managerial and architectural issues that, at this point, pose problems for the virtual workgroup model:

- *Managing Virtual Workgroups.* From a network management perspective, the transitory nature of these virtual workgroups may grow to the point where updating VLAN membership becomes as onerous as updating routing tables to keep up with adds, moves, and changes today (although it may save on the time and effort involved in physically moving the user's workstation). Moreover, there are still cultural hurdles to overcome in the virtual workgroup model: people usually move to be physically close to those with whom they work, rather than to reduce traffic across a collapsed backbone.
- *Maintaining the 80/20 Rule.* Virtual LAN support for virtual workgroups is often tied to support of the "80/20 rule," that is, 80 percent of the traffic is "local" to the workgroup while 20 percent is remote or outside of the workgroup. In theory, by properly configuring VLANs to match workgroups, only the 20 percent of the traffic that is nonlocal will need to pass through a router and out of the workgroup, improving performance for the 80 percent of the traffic that is within the workgroup.

However, many believe that the applicability of the 80/20 rule is waning due to the deployment of servers and/or network applications such as e-mail and Lotus Notes® that users throughout the enterprise access on an equal basis.

- *Access to Local Network Resources.* The virtual workgroup concept may run into the simple problem that users must sometimes be physically close to certain resources such as printers. For example, a user is in the Accounting VLAN, but is physically located in an area populated by members of the Sales VLAN. The local network printer is also in the Sales VLAN. Every time this Accounting VLAN member prints to the local printer, his print file must traverse a router connecting the two VLANs. This problem can be avoided by making that printer a member of both VLANs. This clearly favors VLAN solutions that enable overlapping VLANs, discussed later. If overlapping VLANs are not possible, this scenario would require that routing functionality be built into the backbone switch. Then, the example print file would be routed by the switch rather than having to go through an external router.
- *Centralized Server Farms.* Server farms refer to the placement of departmental servers in a data center, where they can be provided with consolidated backup, uninterrupted power supply, and a proper operating environment. The trend toward server farm architecture has accelerated recently and is expected to continue in order to ease administrative costs.

  Centralized server farms raise problems for the virtual workgroup model when vendor solutions do not provide the ability for a server to belong to more than one VLAN simultaneously. If overlapping VLANs are not possible, traffic between a centralized server and clients not belonging to that server's VLAN must traverse a router. However, if the switch incorporates built-in routing and is able to route inter-VLAN packets at wire speed, there is no performance advantage for overlapping VLANs over routing between VLANs to allow universal access to a centralized server. Remember, only inter-VLAN packets would need to be routed—not all packets. Several vendors support integrated routing as an alternative to overlapping VLANs.

  While workgroup VLANs may be extended to centralized server farms (for example, including a particular file server in a particular workgroup's VLAN), this is not always possible. In some networks, the MIS people who control the servers may want to place routers between the server farms and the rest of the network in order to create a separate administrative domain or to enhance network security via router access control lists. Depending on the vendor implementation, most switching products will not support VLANs that extend across routers (the exception to this would be "VLANs" that equate to IP multicast groups). It should be kept in mind that cordoning off servers with external routers conflicts with one of the reasons for utilizing switches and VLANs in the first place—to avoid the delay introduced by routers.

LAN switches supporting VLANs can be used to effectively control broadcast traffic, reducing the need for routing.

Reduction of Routing for Broadcast Containment

Even the most router-centric networking vendors have come to embrace the philosophy of "switch when you can, route when you must." Although switches certainly provide substantial performance enhancements over layer 3 packet forwarding (routing), as users learned years ago with bridges, switches normally do not filter LAN broadcast traffic; in general, they replicate it on all ports. This not only can cause large switched LAN environments to become flooded with broadcasts, it is also wasteful of precious wide area network bandwidth. As a result, users have traditionally been forced to partition their networks with routers that act as broadcast “firewalls.” Hence, simple switches alone do not allow users to phase out routers completely.

One of the primary benefits of VLANs is that LAN switches supporting VLANs can be used to effectively control broadcast traffic, reducing the need for routing. Broadcast traffic from servers and end-stations in a particular VLAN is replicated only on those switch ports connected to end-stations belonging to that VLAN. Broadcast traffic is blocked from ports with no end-stations belonging to that VLAN, in effect creating the same type of broadcast firewall that a router provides. Only packets that are destined for addresses outside the VLAN need to proceed to a router for forwarding.

There are multiple reasons for utilizing VLANs to reduce the need for routing in the network:

- *Higher Performance and Reduced Latency.* As the network expands, more and more routers are required to divide the network into broadcast domains. As the number of routers increase, latency begins to degrade network performance. A high degree of latency in the network is a problem now for many legacy applications, but it is particularly troublesome for newer applications that feature delay-sensitive multimedia and interactivity. Switches that employ VLANs can accomplish the same division of the network into broadcast domains, but can do so at latencies much lower than those of routers. In addition, performance, measured in packets per second, is usually much higher for switches than for traditional routers. However, it should be noted that there are some switches supporting network layer–defined VLANs that may not perform substantially faster than routers. Additionally, latency is also highly correlated to the number of hops a packet must traverse, no matter what internetworking device (switch or router) is located at each hop.
- *Ease of Administration.* Routers require much more complex configuration than switches; they are “administratively rich.” Reducing the number of routers in the network saves time spent on network management.
- *Cost.* Router ports are more expensive than switch ports. Also, by utilizing cheaper switch ports, switching and VLANs allow networks to be segmented at a lower cost than would be the case if routers alone were used for segmentation.

In comparing VLANs with routing, VLANs have their disadvantages as well. The most significant weakness is that VLANs have been, to date, single-vendor solutions and therefore may lead to switch vendor lock-in. The primary benefits of VLANs over routing are the creation of broadcast domains without the disadvantages of routing and a reduction in the cost of moves and changes in the network. Therefore, if neither of these is a problem, then the user organization may want to forgo VLANs and continue deploying a multivendor network backbone, segmented by a mix of a few routers and a relatively large number of simple switches.

Assuming a major implementation of VLANs, what is the role of routers in a network? Routers have two remaining responsibilities: to provide connectivity between VLANs, and to provide broadcast filtering capabilities for WAN links, where VLANs are generally not appropriate.

Routing Between VLANs. VLANs can be used to establish broadcast domains within the network as routers do, but they cannot forward traffic from one VLAN to another. Routing is still required for inter-VLAN traffic. Optimal VLAN deployment is predicated on keeping as much traffic from traversing the router as possible. Minimizing this traffic reduces the chance of the router developing into a bottleneck. As a result, the corollary to “switch when you can, route when you must” in a VLAN environment becomes “routing is used only to connect VLANs.”

Having said this, however, keep in mind that in some cases routing may not prove to be much of a bottleneck. As mentioned earlier, integrating routing functionality into the backbone switch eliminates this bottleneck if this routing is accomplished at high speed for inter-VLAN packets.

VLANs Over the WAN. Theoretically, VLANs can be extended across the WAN. However, this is generally not advised, since VLANs defined over the WAN will permit LAN broadcast traffic to consume expensive WAN bandwidth. Because routers filter broadcast traffic, they neatly solve this problem. However, if WAN bandwidth is free for a particular organization (for example, an electric utility with dark fiber installed in its right of way), then extending VLANs over a WAN can be considered. Finally, depending on how the they are constructed, IP multicast groups (functioning as "VLANs") can be effectively extended across the WAN, as well as the routers providing the WAN connections, without wasting WAN bandwidth.

Security

The ability of VLANs to create firewalls can also satisfy more stringent security requirements and thus replace much of the functionality of routers in this area. This is primarily true when VLANs are implemented in conjunction with private port switching. The only broadcast traffic on a single-user segment would be from that user's VLAN (that is, traffic intended for that user). Conversely, it would be impossible to "listen" to broadcast or unicast traffic not intended for that user (even by putting the workstation's network adapter in promiscuous mode), because such traffic does not physically traverse that segment.

VLANs and ATM

While the concept of VLANs originated with LAN switches, their use may need to be extended to environments where ATM networks and ATM-attached devices are also present. Combining VLANs with ATM networks creates a new set of issues for network managers, such as relating VLANs to ATM emulated LANs (ELANs), and determining where to place the routing function.

VLANs Transparent to ATM

In a LAN backbone with VLANs spanning more than one LAN switch, switches determine where frames have originated by the techniques discussed earlier in "Communicating VLAN Membership Information" (VLAN tables, frame tagging, and TDM). In an environment where ATM exists only in the backbone (that is, there are no ATM-connected end-stations), ATM permanent virtual circuits (PVCs) may be set up in a logical mesh to carry intra-VLAN traffic between these multiple LAN switches.

In this environment, any proprietary technique the vendor has employed is transparent to the ATM backbone. ATM switches do not have to be VLAN "aware." This means that ATM backbone switches could be from a different vendor than the LAN switches; ATM backbone switches could be selected without regard for VLAN functionality, allowing network managers to focus more on performance-related issues. As convenient as this situation sounds, it does not reflect reality for many network environments.

Complexity Arising with ATM-Attached Devices

Usually, organizations that implement ATM backbones would also like to connect workstations or, more likely, servers directly to those backbones. As soon as any logical end-station is connected via ATM, a new level of complexity arises. LAN Emulation must be introduced into the network to enable ATM-connected end-stations and non-ATM-connected end-stations to communicate.

LAN Emulation

With the introduction of ATM-connected end-stations, the network becomes a truly "mixed" environment, with two types of networks operating under fundamentally different technologies: connectionless LANs (Ethernet, Token Ring, FDDI, etc.) and connection-oriented ATM. This environment puts the responsibility on the ATM side of the network to "emulate" the characteristics of broadcast LANs and provide MAC-to-ATM address resolution.

The LAN Emulation (LANE) specification, standardized in 1995 by the ATM Forum, specifies how this emulation is accomplished in a multivendor environment. LANE specifies a LAN Emulation server (LES), which can be incorporated into one or more

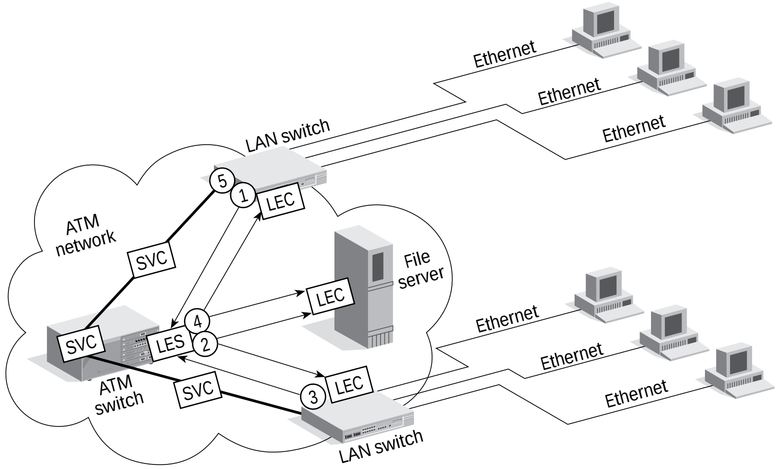

Figure 2. *LAN Emulation*

switches or a separate workstation to provide the MAC-to-ATM address resolution in conjunction with LAN Emulation clients (LECs), which are incorporated into ATM edge switches and ATM NICs.

Figure 2 briefly illustrates how LANE operates:

1. The LAN switch receives a frame from an Ethernet-connected end-station. This frame is destined for another Ethernet end-station across the ATM backbone. The LEC (which in this situation resides in the LAN switch) sends a MAC-to-ATM address resolution request to the LES (which in this case resides in an ATM switch).
2. The LES sends a multicast to all other LECs in the network.
3. Only the LEC that has the destination (MAC) address in its tables responds to the LES.
4. The LES then broadcasts this response to all other LECs.
5. The original LEC recognizes this response, learns the ATM address of the destination switch, and sets up a switched virtual circuit (SVC) to transport the frame via ATM cells as per AAL5, which governs segmentation and reassembly.

In looking at the path of traffic between an Ethernet-attached client and an ATM-attached server, the section that is governed by LANE extends from the LEC in the ATM interface of the LAN switch to the LEC residing in the server's ATM NIC. From the standpoint of either MAC driver, frames pass directly between them just as if they were connected by a non-ATM backbone, with each LEC acting as a proxy MAC address. VLANs defined by port group would treat the ATM interface on the LAN switch as just another Ethernet port, and all ATM-attached devices would then be members of that VLAN. In this way, VLANs could be deployed without regard to whether the ATM switches in the backbone are from the same vendor (so long as they support LANE).

However, from an administrative point of view, many organizations may not want to employ separate management software for the ATM backbone and may prefer to source both edge devices (LAN switches) and backbone devices (ATM switches) from the same vendor.

LANE can also allow for multiple ELANs by establishing more than one LEC in the ATM interfaces of participating devices (as well as a separate LES for each ELAN). Each LEC in the ATM interface of the LAN switch is treated as a separate logical Ethernet port, and each LEC in a single ATM-attached device is seen as a separate Ethernet-attached end-station. Therefore, multiple LECs in a single ATM-attached device can be members of different VLANs, allowing these VLANs to overlap at ATM-attached devices. Since LANE supports only ATM-attached devices,

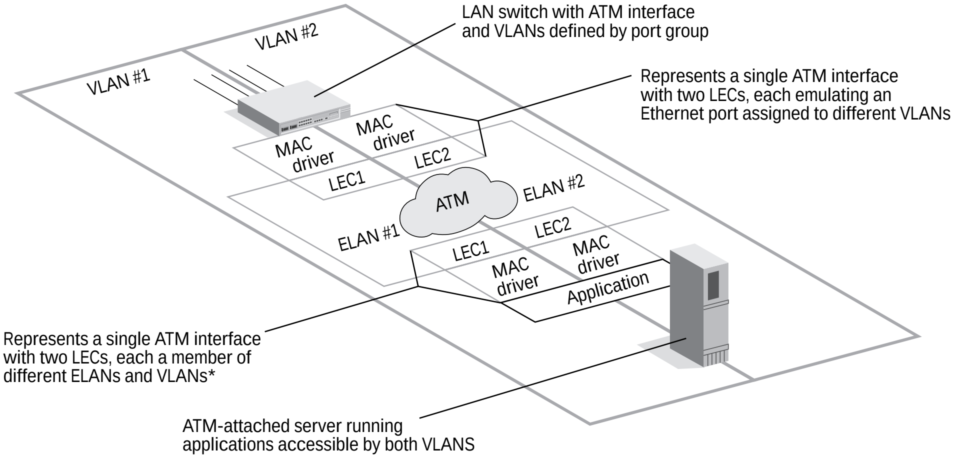
LAN switch with ATM interface
and VLANs defined by port group
VLAN #2
VLAN #1
Represents a single ATM interface
with two LECs, each emulating an
Ethernet port assigned to different VLANs
MAC
driver
MAC
driver
LEC1
LEC2
ATM
ELAN #2
ELAN #1
LEC1
LEC2
MAC
driver
MAC
driver
Application
Represents a single ATM interface
with two LECs, each a member of
different ELANs and VLANs*
ATM-attached server running
applications accessible by both VLANS

** Note: Each LEC on a single ATM interface* must *be on separate ELANs. They are shown here on separate VLANs only because their corresponding LECs on the ATM switch have been explicitly assigned to different VLANs.*

Figure 3. *VLANs as Supersets of ELANs*

while VLANs are defined for both ATM and non-ATM network devices, VLANs can be seen as supersets of ELANs (Figure 3).

With this structure, an ATM backbone can enable all end-stations from multiple VLANs to access a centralized server or servers without passing through a router by establishing a separate ELAN for each VLAN. Since most traffic in a network is between client and server, establishing VLANs that overlap at ATM-attached servers greatly reduces the number of packets that must be routed between VLANs. Of course, there is still likely to be a small amount of inter-VLAN traffic remaining. Therefore, a router is still required for traffic to pass from one VLAN to another (and, therefore, from one ELAN to another). Figure 4 depicts this type of structure.

Routing Between Emulated LANs and/or VLANs

Since routing remains necessary in any mixed ATM/shared media environment to forward inter-VLAN traffic, network designers are faced with the question of where to locate the router functionality. The following are four architectural solutions to the problem of where to locate the routing functionality: edge

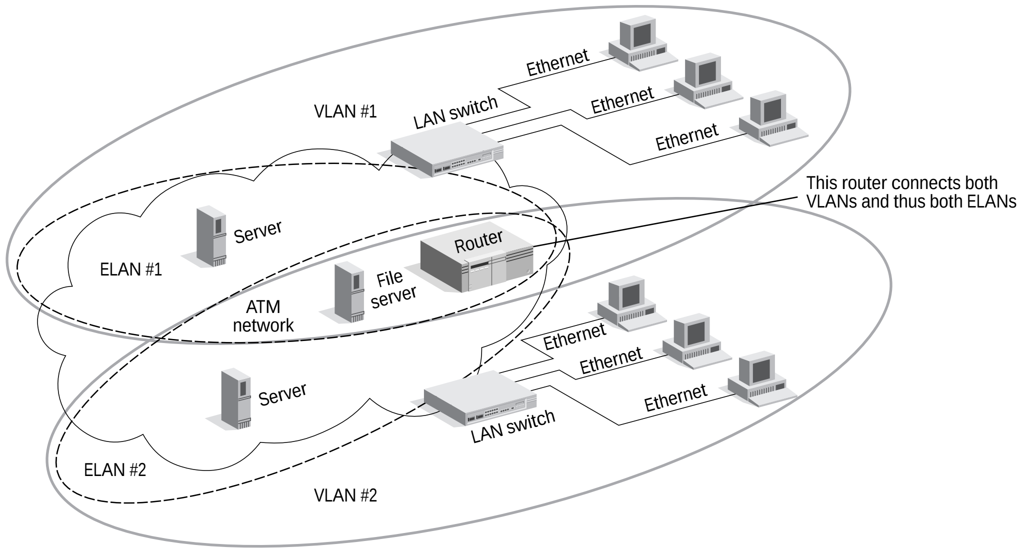

Figure 4. *Router Connecting Overlapping VLANs/ELANs*

routing, the "one-armed" router, the route server, and MPOA.

Edge Routing. Basically, edge routing dictates that the routing function across the ATM backbone be incorporated into each LAN switch at the "edge" of the ATM backbone. Traffic within VLANs can be switched across the ATM backbone with minimal delay, while inter-VLAN packets are processed by the routing function built into the switch. In this way, an inter-VLAN packet does not have to make a special trip to an external router, eliminating a time-consuming extra hop.

There are three other major advantages to this architecture. First, unlike solutions that have centralized routing, there is no single point of failure with edge routing architectures. Second, several solutions featuring edge routing are available today. Third, edge routing will function in multivendor environments if each vendor's equipment supports LAN Emulation.

The primary disadvantage of edge routing is the difficulty of managing multiple physical devices relative to having centralized management of a consolidated router/routing function. Additionally, edge routing solutions may be more expensive than centralized routing solutions made up of a centralized router and multiple, less-expensive edge switches.

The One-Armed Router. The concept of the so-called "one-armed router" has become particularly attractive because it removes the more processing-intensive, higher-latency routing function from the primary data path. A one-armed router sits off the side of an ATM backbone switch with a single ATM link, allowing packets that do not need to traverse the router to pass through the ATM backbone unimpeded. Another advantage of the one-armed router is that, relative to other configurations, it is less complex to configure and administer.

The key to the one-armed router structure, shown in Figure 5, is to keep as much traffic as possible out of the one-armed router. By structuring VLANs to support the 80/20 rule (so that 80 percent of the traffic remains within each VLAN), the router is not required to handle most traffic. For this to work well, optimal configuration of VLANs to minimize inter-VLAN traffic (traffic passing through the one-armed router) is critical. There are several vendors presently shipping one-armed router solutions.

One of the disadvantages of the one-armed router is that it represents a single point of failure in the network. For this reason, two or more redundant one-armed routers are generally preferred. However, perhaps the most significant drawback of the one-armed router is that its one arm can develop into a bottleneck if VLAN traffic does not support the 80/20 rule. This can occur particularly in networks with large amounts of peer-to-peer traffic.

The Route Server. The route server model (see Figure 6) is *physically* similar to the one-armed router model, but *logically* very different in that it breaks up the routing function into distributed parts. In a one-armed router configuration, a packet from VLAN A heading to

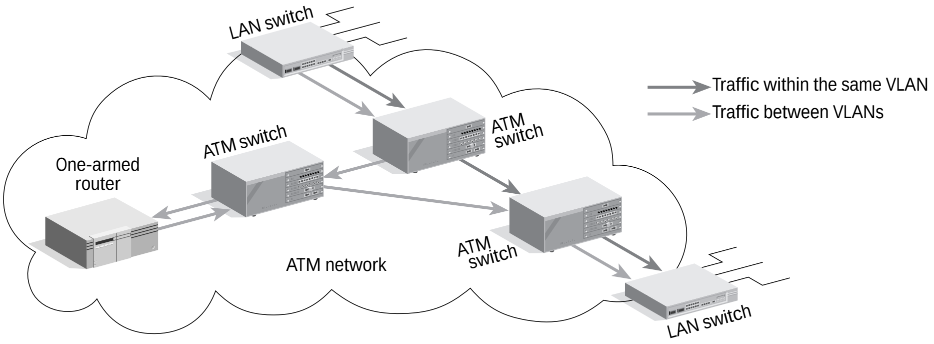

Figure 5. *One-Armed Router*

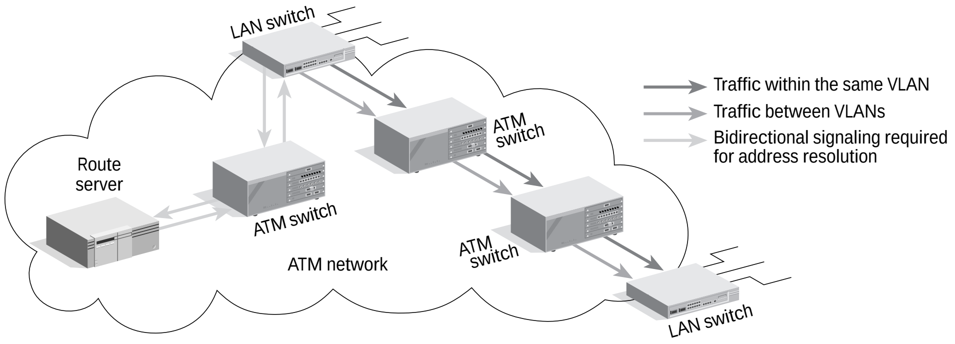

Figure 6. *Route Server*

VLAN B is sent to the one-armed router, where it waits for address resolution, path calculation, establishment of a connection across the ATM backbone, and, finally, transmission. In a route server scheme, the same packet waits in the cache of the LAN switch at the edge of the ATM backbone before transmission. In this process, the packet itself never traverses a router. The only traffic to and from the route server is the signaling required to set up a connection between LAN switches across the ATM backbone. The advantage is that less routed traffic must be diverted to the route server, often reducing the number of hops required through the backbone. Also, overall traffic across the route server's one arm is reduced.

There are, of course, disadvantages to the route server approach as well. First, initial vendor implementations are strictly proprietary and do not support standard routing protocols. Secondly, at this point available route servers only support IP. Of course, the route server shares one of the one-armed router's drawbacks in that it can be a single point of failure, but, as with the one-armed router, this problem can be mitigated through redundancy. Finally, because a route server architecture requires LAN switches to have a certain level of routing functionality, route server solutions tend to be more expensive and more complex to configure than the relatively simple LAN switches deployed in the one-armed router architecture.

MPOA. There is at least one development that may eventually standardize the route server approach. The Multiprotocol over ATM (MPOA) standards working group of the ATM Forum is currently working out the details of an implementation model for MPOA service. While a variety of models have been proposed, MPOA is expected to provide direct virtual circuit connectivity between ATM-network-attached devices that may belong to different routing subnets. In other words, MPOA can let logical end-stations that are part of different ELANs communicate directly across an ATM network without requiring an intervening router.

Since ELANs are subsets of VLANs, MPOA holds the promise of enabling an ATM backbone to connect VLANs without the need for an external router. MPOA can be considered an enhancement beyond LANE that integrates routing functionality into the LAN-ATM edge switch. All inter-VLAN traffic would be able to leverage this capability, and network latency would be reduced.

An MPOA standard is not expected to be finalized until at least 1997, and the initial implementation will most likely support only TCP/IP. It should be noted that some of the disadvantages of the route server approach, such as cost and management complexity, would remain in MPOA solutions.

VLANs and DHCP: Overlapping Solutions

With Microsoft's recent introduction of the Dynamic Host Configuration Protocol (DHCP), users now have another alternative for reducing the workload associated with administration of workstation IP address. Unfortunately, DHCP can actually conflict with VLAN implementation, especially with layer-3, IP-based VLANs.

DHCP Functionality

When considering the ability of VLANs to deal with ever-changing networks, it should be remembered that most of the difficulty in supporting adds, moves, and changes occurs in IP networks. In order to deal with the problem of reconfiguring IP addresses, Microsoft has developed DHCP, a TCP/IP-based solution incorporated into the Windows NT™ server and most Windows® clients.

Rather than establishing location-independent broadcast domains as VLANs do, DHCP dynamically allocates IP addresses to logical end-stations for fixed periods of time. When the DHCP server detects a workstation whose physical location no longer corresponds to its allocated IP address, it simply allocates that end-station a new address. By doing so, DHCP enables workstations to be moved from subnet to subnet without the network administrator having to manually configure the workstation's IP address or update host table information.

The element of DHCP that equates most closely to VLAN functionality is the network administrator's ability to specify a range of IP addresses available for a particular logical workgroup. These logical groups are termed "scopes" in the Microsoft lexicon. However, scopes should not be equated with VLANs, because members of a single scope are still bound by their physical subnet, although there can be multiple scopes residing in each subnet. Consequently, DHCP implementation may reduce the labor-intensive administration of TCP/IP networks, but DHCP alone does not control network broadcasts in the same way that VLANs do.

Best Use for Each

In what types of network environments should VLANs be implemented, and in what types of network environments does DHCP make the most sense? Since DHCP is solely an IP-based solution, it has little appeal in environments where IP users are a minority, since all non-TCP/IP clients would be excluded from scope membership. In particular, network environments where non-TCP/IP protocols are required for mission-critical applications may benefit more from VLAN implementation, since VLANs can be used to contain multi-protocol broadcast traffic.

However, for smaller, purely TCP/IP network environments (under 500 nodes), DHCP alone may suffice. By simply having fewer total network nodes and fewer physical subnets, the need to establish fully location-independent logical groups is greatly reduced. Additionally, for medium-sized organizations that, for whatever reason, do not support location-independent workgroups, VLANs lose much of their appeal when compared to DHCP.

There is one area in which VLANs and DHCP do not compete: reducing the necessity for routing in the network. Although DHCP servers dynamically maintain address tables, they lack routing functionality and cannot create broadcast domains. Therefore, DHCP has no impact on an organization's need for routing in the network. In environments where the containment of broadcast traffic without having to resort to routers is a major requirement, VLANs are a better solution.

Overlap Between DHCP and VLANs

It what ways can DHCP and VLANs work together, and in what situations do they represent competitive solutions?

DHCP and layer-3, IP-based VLANs clearly represent competitive solutions because of addressing problems that stem from implementing layer 3–based VLANs in conjunction with DHCP. If a client workstation physically moves to a new subnet, the DHCP server will allocate a new IP address for that workstation. Yet, this workstation's VLAN membership is based on the old IP address. Therefore, the network administrator would have to manually update the client's IP address in the switch's VLAN tables. This would eliminate the primary benefit of DHCP and one of the primary benefits of IP-based VLANs. In summary, these two solutions represent an either/or proposition for most network environments.

Implementing VLANs defined by MAC-layer address in conjunction with DHCP is a somewhat more plausible solution. However, DHCP together with MAC-based VLANs would create a two-tiered, redundant matrix of logical groups (MAC address–based VLANs and DHCP scopes). Having two tiers of logical groups would make otherwise easy-to-manage, “drag-and-drop” moves, adds, and changes unnecessarily difficult and might entail more labor-intensive network administration than if neither solution was implemented.

Port group–based VLANs and DHCP can coexist, and their joint implementation can even be complementary. As stated earlier, when users in VLANs based purely on port groups move from one port group to another, their VLAN membership changes. In a non-DHCP environment where IP subnets correspond one-to-one with VLANs, users who move from one port group to another would still need to have their workstation reconfigured to reflect their new IP subnet. Implementing DHCP would make this reconfiguration automatic. The port group–based VLANs, of course, provide the broadcast containment that DHCP implementation alone does not. In this way, DHCP and port-group-based VLANs can work together to accomplish both broadcast containment and automation of moves and changes.

> The choice of approach will have a substantial impact on the overall network architecture, and may even affect the management structure and business model of the organization.

Port group–based VLANs and DHCP, in conjunction with deployment of architectures that reduce the need for external routing of inter-VLAN traffic (such as multiple VLAN memberhip or integrating routing into the switch), represent a fairly complete short- to medium-term solution, which will alleviate the most pressing problems faced in many network environments.

VLAN Architectures Going Forward

Due to the trends toward server centralization, enterprise-wide e-mail, and collaborative applications, various network resources will need to be made available to users regardless of their VLAN membership. Ideally, this access should be provided without most user traffic having to traverse a router.

Organizations that implement VLANs recognize the need for certain logical end-stations (for example, centralized servers) to communicate with multiple VLANs on a regular basis, either through overlapping VLANs (in which network-attached end-stations simultaneously belong to more than one VLAN) or via integrated routing that can process inter-VLAN packets at wire speed. From a strategic standpoint, these organizations have two ways to deploy VLANs: an “infrastructural” VLAN implementation or a “service-based” VLAN implementation. The choice of approach will have a substantial impact on the overall network architecture, and may even affect the management structure and business model of the organization.

Infrastructural VLANs

An infrastructural approach to VLANs is based on the functional groups (that is, the departments, workgroups, sections, etc.) that make up the organization. Each functional group, such as accounting, sales, and engineering, is assigned to its own uniquely defined VLAN. Based on the 80/20 rule, the majority of network traffic is assumed to be within these functional groups, and thus within each VLAN. In this model, VLAN overlap occurs at network resources that must be shared by multiple workgroups. These resources are normally servers, but could also include printers, routers providing WAN access, workstations functioning as gateways, and so forth.

The amount of VLAN overlap in the infrastructural model is minimal, involving only servers rather than user workstations—making VLAN administration relatively straightforward. In general, this approach fits

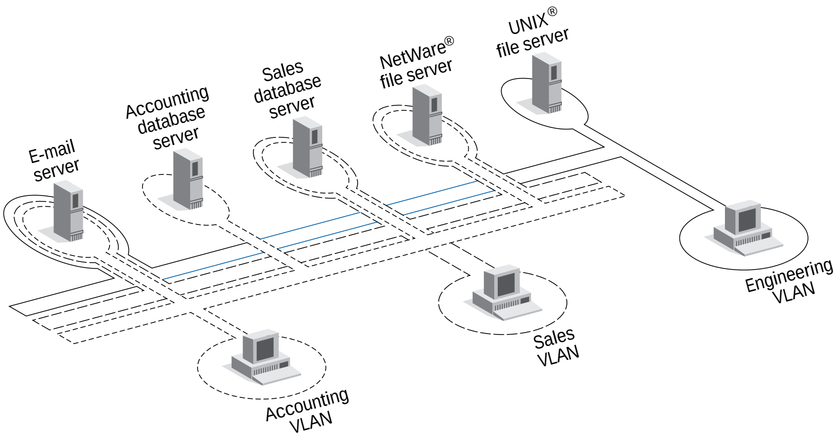

Figure 7. *Infrastructural VLANs*

well in those organizations that maintain clean, discrete organizational boundaries. The infrastructural model is also the approach most easily enabled by presently available solutions and fits more easily with networks deployed today. Moreover, this approach does not require network administrators to alter how they view the network, and entails a lower cost of deployment. For these reasons, most organizations should begin with an infrastructural approach to VLAN implementation.

As can be seen in the example in Figure 7, the e-mail server is a member of all of the departments' VLANs, while the accounting database server is only a member of the accounting VLAN.

Service-Based VLANs

A service-based approach to VLAN implementation looks, not at organizational or functional groups, but at individual user access to servers and applications—that is, network resources. In this model, each VLAN corresponds to a server or service on the network. Servers do not belong to multiple VLANs—groups of users do. In a typical organization, all users would belong to the e-mail server's VLAN, while only a specified group such as the accounting department plus top-level executives would be members of the accounting database server's VLAN.

By its nature, the service-based approach creates a much more complex set of VLAN membership relationships to be managed. Given the level of most VLAN visualization tools presently available, a large number of overlapping VLANs using the service-based approach could generate incomprehensible multilevel network diagrams at a management console. Therefore, to be practical, service-based VLAN solutions must include a high level of automatic configuration features. However, in response to the types of applications organizations want to deploy in the future, as well as the shift away from traditional, more rigid organizational structures, the trend in VLAN implementation will be toward the service-based approach. Figure 8 depicts the service-based VLAN model.

As bandwidth to the desktop increases and as vendor solutions become available to better manage greater VLAN overlap, the size of the groups that belong to a particular set of VLANs may become smaller and smaller. At the same time, the number of these groups becomes larger and larger, to the point where each individual could have a customized mix of services delivered to his or her workstation. Taking that concept a step further, control over what services are delivered at a given time could be left up to each individual user. At that point, the network structure begins to take on the multiple-channel characteristics of a cable TV (CATV) network. In fact, at this stage, this model finds the greatest degree of similarity in VLANs defined by IP multicast group—each workstation has the choice of which IP multicast or "channel" it wants to belong to.

In such a future environment, VLANs lose the characteristics of static or semistatic broadcast domains defined by the network manager, and become channels to which users subscribe. Users simply sign up for the applications they need delivered to them at a particular time. Application use could be accounted for, enabling precise and automated chargeback for network services. Network managers could also retain control in order to block access to specific channels by certain users for security purposes.

VLAN Migration Strategies

As this paper has demonstrated, there are many factors to be considered in VLAN implementation: technological, architectural, and organizational. Given the effects of VLANs on network architecture, organizational structure, and even the business model of some organizations, it is difficult to deploy VLAN technology solely as a tactical solution, only where and when it is needed. However, this does not imply an all-or-nothing strategy in which the network architecture is transformed overnight from one based on physical subnets and router-based segmentation to one of service-based VLANs.

What steps are necessary before applying VLANs to an enterprise network? Initially, VLANs should be seen as a solution to at least one of two problems:

- Containment of broadcast traffic to minimize dependence on routers
- Reduction in the cost of network moves and changes

An organization where broadcast traffic is not yet a problem, or where the cost of network moves and changes is tolerable, may want to forgo implementing VLANs for the time being. However, the majority of large enterprise networks are now experiencing one or both of these problems.

In organizations that are rapidly replacing routers with switches and may soon face broadcast traffic containment issues, another element of the network architecture should be considered: the degree to which the network has evolved toward a single user/port switched LAN architecture. If the majority of users are still on shared LAN segments, the ability of VLANs to contain broadcasts is greatly reduced. If multiple users belonged to different VLANs on the same shared LAN segment, that segment would receive broadcasts from each VLAN—defeating the goal of broadcast containment.

Having determined that VLANs need to be a part of network planning in the immediate future, server access, server location, and application utilization must all be thoroughly analyzed to determine the nature of traffic flow in the network. This analysis should answer the remaining questions about where VLAN broadcast domains should be deployed, what role ATM needs to play, and where the routing function should to be placed.

Because of the limitations of present VLAN technology, initial VLANs are likely to employ an infrastructural approach.

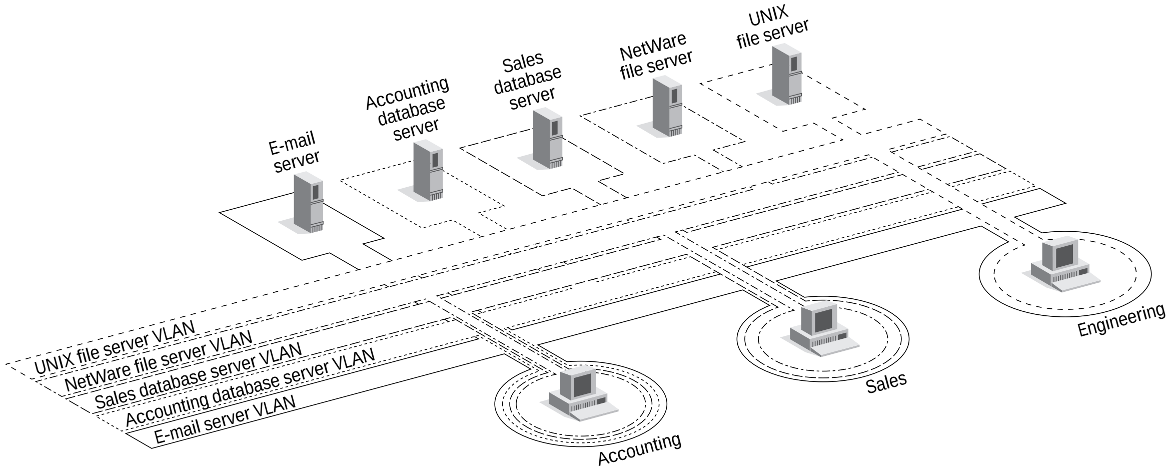

Figure 8. *Service-Based VLANs*

However, as vendor solutions develop, many organizations will want to consider migration toward a more service-based model, which will more easily let users subscribe to various network services.

This concept of user-controlled subscribership, as opposed to administrator-controlled membership, is augmented by NICs with built-in VLAN functionality operating in environments with a single user per switch port. In this scheme, the NIC driver dynamically tells the switch which multicast groups or VLANs it wants to belong to. Certainly, this type of distributed VLAN control leverages the increasing processing power of the desktop and enables a higher degree of other, related functionality such as automatic VLAN configuration and traffic monitoring. In addition, agents residing in each NIC will enable the workstation to collect and report information on specific application usage (rather than just simple layer 2 traffic statistics in the case of RMON1). This capability facilitates the automated chargeback for network services described earlier for service-based VLANs.

If individual users control VLAN membership, what about security? Clearly, users cannot be allowed to simply subscribe to any network service they wish. The network administrator must be able to establish policies that define which users have access to what resources and what class of service each user is entitled to. One solution to the security problem may come in the form of an authentication server. These servers may well develop into the primary method by which the VLANs of the future are defined. Authentication servers define VLAN membership by user ID (password or other authentication device) rather than by MAC address or IP address. Defining VLANs in this way greatly increases flexibility and also implies a certain level of integration of VLANs with the network operating system, which typically asks the user for a password anyway to allow or deny access to network resources. One of the primary advantages of authentication servers is that they allow the user to take his or her VLAN anywhere, without regard to which workstation or protocol is being used.

The analysis of network traffic, applications usage, server access, and so on that is necessary in the VLAN migration process, and which will be greatly furthered by the implementation of RMON2, may simply produce VLANs that correspond to functional teams or departments. On the other hand, if migration is undertaken with a holistic view of the capabilities of VLAN technology, and the network designers ask the question, “Who *should* talk to whom?” rather than “Who *is* talking to whom?,” it may become apparent that fundamental process and organizational changes are needed. Many organizations are making such changes: trends such as flatter hierarchies, revamped workflows, and innovative business models are helping to fully leverage the possibilities of emerging applications.

Conclusion

The concept of service-based VLAN technology holds the potential for harmonizing many of today’s organizational and managerial changes with the structural and technological developments in the network. Despite the promise of this vision, VLAN implementation must solve real-world problems in order to be financially justified. Organizations that have deployed or are planning to deploy large numbers of switch ports, dividing the network into smaller segments to increase bandwidth per user, can make a very strong case for VLAN implementation in order to contain broadcasts. However, any organization that expends substantial resources dealing with moves and changes in the network may also be able to justify VLAN implementation. This is simply because VLANs, if implemented as part of a strategic solution, may be able to substantially reduce the cost of dealing with moves and changes. For these organizations, the switching infrastructure upon which most VLAN solutions are based can be seen as an added, and quite valuable, benefit. □

Printed in U.S.A.

200374-001 5/96