MadgeOne

Multilayer IP/IPX switching in the LAN

A Technology White Paper

AUTHOR: Martin Taylor, Madge Networks, January 1997

Executive Summary

This paper is intended for network designers and planners who are faced with the problem of upgrading their LANs to support growing traffic loads and changing traffic patterns, and who are charged with keeping the costs of operating the LAN under tight control. The paper describes multilayer switching as a complete, self-contained, cost-effective and fully scalable solution for the LAN infrastructure. Multilayer switching combines the best characteristics of LAN switching and routing to provide higher performance at lower cost and with greater flexibility than traditional router-based LAN backbones, without sacrificing any of the security or management controls that are typically provided by routers.

About the Author

Martin Taylor is Vice President, Network Architecture for Madge Networks. He is responsible for the strategic planning of the company's products, including LAN switching hubs, ATM switches and LAN-ATM access switches, and is the principal designer of the MadgeOne architecture. Martin joined Madge Networks in 1991, and formerly held the position of Director of Product Marketing. Prior to joining Madge, Martin held business development positions in Local Area Networks and in fiber optic cabling systems with GPT Ltd, the UK's leading supplier of telecommunications equipment. Martin has a total of 11 years experience in the communications industry, and a further 6 years experience in MIS management. He can be reached on email at mtaylor@madge.com.

Contents

The MadgeOne Architecture

MadgeOne is an architecture for multiservice switched networks, which provides a scalable, high performance, cost-effective solution for the LAN infrastructure, and which evolves to provide future support for videoconferencing, voice telephony and real-time multimedia applications at the desktop.

The reader who is interested in learning more about how the MadgeOne architecture handles real-time communications is invited to read the Madge white paper “MadgeOne – Architecture for Multiservice Switched Networks”, July 1996.

In this paper we are going to focus on the principal technologies employed in the MadgeOne architecture to address the needs of existing and emerging data applications: multilayer/ multiprotocol switching.

Why Multilayer Switching?

When the first LANs were installed in enterprise networks in the early 1980s, 10 Megabits per second – even when shared among several hundred users – seemed an extraordinary amount of bandwidth. And compared with the 2400 and 9600 baud communications links commonly used to connect dumb terminals to mini and mainframe computers, 10 Mbps was indeed generous!

The inventors of those original LANs promised that the advent of open communications networking between distributed computing resources would fundamentally change the landscape of information processing systems, even if they weren't, at the time, very clear about the exact nature of this change. It took quite a few years for the full implications of open networking to emerge: but now the LAN is ubiquitous, and the new vocabulary of network computing is expanding at a breathtaking pace. Who could have foreseen three years ago that the graphical web browser would lead to such an explosive growth of Internet usage? Or that this same technology would be taken up with such enthusiasm in corporate and enterprise networks – now increasingly referred to as "intranets"?

As LANs move beyond the era of simple mainframe connectivity and file and print services to encompass client/server computing, workflow applications, intranet Web server access and real-time voice and video communications, so the technology of LANs has needed to undergo a continual overhaul. More and more powerful computing systems drive the network ever harder, and larger and more information-rich data objects create a continually growing network traffic load. Traffic patterns are changing too, as workgroup- or department-level server-based solutions give way to corporate servers, pushing unprecedented levels of traffic across the LAN backbone.

To keep up with this growth in network activity, LAN technologies have evolved in two dimensions – speed and segmentation. *Speed* simply delivers more bandwidth: Ethernet has added 100 Mbps and (soon) 1 Gbps variants, Token Ring went from 4 to 16 Mbps, while FDDI and ATM have extended the speed choice still further. *Segmentation* enables more of this bandwidth to be delivered to individual users or hosts. Segmentation technologies have evolved from local bridging through collapsed backbone routing to LAN switching; and since switching enables dedicated bandwidth to be delivered economically to each user, this is clearly where the future of high performance LANs lies.

But today, LAN switching does not provide a complete and general solution for large-scale LAN installations. This is because conventional LAN switching is not completely scalable. In most practical networks today, LAN switches must be used in conjunction with routers. We will come to the reasons for this later in this paper. But in practice, having to combine switches with routers to build LAN infrastructures has a major impact in a number of areas:

- From the viewpoint of the network's users, it divides LAN communications into two classes of performance. Packets traveling on direct switched connections enjoy the high speed and consistent journey times of the expressway, but packets that are

forced to pass through routers take the slow road, and suffer further delays when traffic is heavy.

- Switches and routers are separate pieces of equipment that must be separately purchased, configured and managed, and inevitably the costs of implementing a solution based on these multiple elements are greater than for a solution based on a single, integrated approach.
- The use of routers in the LAN results in additional administration to deal with moves and changes. For example, if a PC is moved from a LAN segment on one router port to a different LAN segment connected to another router port, it is usually necessary to allocate a new IP address to the PC and re-configure this address in the PC's software.

Multilayer switching is based on an intelligent combination of switching and routing technologies to provide a complete and integrated solution for all kinds of LAN infrastructures, which very effectively addresses each of these issues. That's why more and more network planners will soon be setting their sights on a migration to multilayer switching in the LAN.

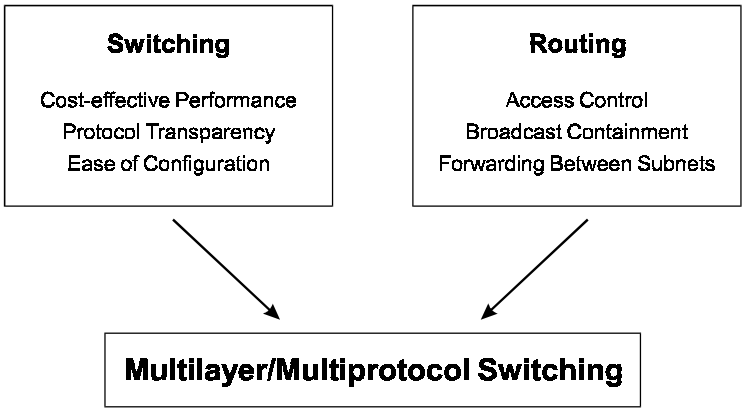

Figure 1: Convergence of Switching and Routing

Switching and Routing in the LAN

LAN switching has emerged as the most cost-effective way to segment shared LANs so that each user gets a greater share of the available bandwidth. Switching can be employed to link shared LAN segments together, or to deliver dedicated bandwidth to individual hosts, whether they be desktop PCs or server machines. Switching can also be used to link LAN technologies of different speeds – for example, connecting users with 10 Mbps Ethernet connections to a 100 Mbps Fast Ethernet uplink to the LAN backbone. The use of ATM with Ethernet or Token Ring LAN Emulation is a direct extension of the concept of LAN switching.

LAN switching operates at Layer 2 of the OSI 7-layer model of networking protocols, the Data Link layer. This means that LAN switches forward packets based on their Ethernet or Token Ring destination Media Access Control (MAC) address, or, in the case of source routed Token Ring, on the Routing Information Field (RIF). LAN switches are therefore transparent to Network Layer protocols such as the Internet Protocol (IP) or Novell's Internetwork Packet eXchange (IPX).

LAN switches have a number of physical ports for connection to LAN segments, typically in the range 8 to 128 or more, and they learn which ports are associated with which MAC destination addresses by extracting the source MAC address of each packet that is sent to the switch and learning the association between the MAC address and the port on which the packet was received. Because they are largely self-configuring, LAN switches are easy to install, configure and manage.

Routers, by contrast, operate at Layer 3 of the 7-layer model, the Network Layer, and they forward packets according to Network Layer address prefixes in conjunction with routing tables held in router memory. These routing tables are kept up-to-date with the aid of routing protocols which are used to exchange reachability information between routers. Routers are much more complex than LAN switches, and are certainly more costly. A good deal of software is involved with the processing of each packet through a router, so they are generally much slower than LAN switches, and harder to configure and manage.

LAN switches are simple, cost-effective and offer excellent performance. So when we need to increase the capacity of the LAN, why not migrate the entire LAN infrastructure to operate exclusively on LAN switches? The answer is that, in some cases, this is indeed possible – but more often than not there is some necessity for routing in the LAN, for one or more of the following reasons:

- Historically, the widespread use of routers in existing LAN infrastructures has led to the adoption of IP addressing schemes in which multiple IP subnets exist within the LAN. Packets can only be forwarded between subnets by a routing function, not by LAN switches. Replacing routers by switches would require that the IP addresses of all the stations on the LAN be changed so that all the addresses are on a single subnet identity, and many organizations would reject this as being impractical.

- Unlike LAN switches, most routers offer a range of access controls which can form a useful part of the enterprise LAN security policy. This includes the ability to block certain IP or IPX addresses or ranges of addresses from accessing certain resources. It is not always possible to provide this type of security in the end systems themselves, and therefore this kind of capability will continue to be needed by security-conscious organizations as part of the network itself.

- A fully switched LAN which has no routing in it would have to operate as a single large broadcast domain. This may result in unacceptable levels of broadcast traffic being seen by each end station connected to the network, causing congestion on the network adapters and slowing network response times to an unacceptable degree. The extent to which broadcast traffic loading is a problem in any particular environment is dependent on both the number of stations in the network and the types of protocols being used, but in the general case it is often considered necessary to split the network up into broadcast domains, which in turn require routing functionality to link them together.

- Some LANs contain a mix of technologies, such as Ethernet and Token Ring, between which it may be impractical to achieve switched connectivity at Layer 2 due to the problems of packet format translation and differences in maximum permitted packet size.

Alternative solutions do exist for the first three of these requirements for routing functionality in the LAN. Some organizations, recognizing the value of moving to a fully switched LAN, are changing to private IP addressing schemes which provide the flexibility to assign all stations in one site to a single subnet. Where Internet access is required, an address translation gateway is employed which also operates as a security firewall. Security within the LAN can be implemented exclusively at the level of the end system – depending on the application software or the network operating system – thereby removing the need for the network to implement access control. And the broadcast issue can be addressed by more sophisticated means than simply breaking the network up into broadcast domains. For example, some LAN switches can apply Layer 3 intelligence to the filtering of unnecessary broadcast traffic, as in Madge's Active Broadcast Control technology.

Nevertheless, in many real networks the practicalities of the situation will dictate that routing functionality is required, for one or more of the reasons we have identified. And this means that if conventional LAN switches are deployed, it is necessary to deploy also routers.

Further discussion on the coexistence of switching and routing in the LAN can be found in the Madge white paper "The Architecture of Switched LANs," March 1996.

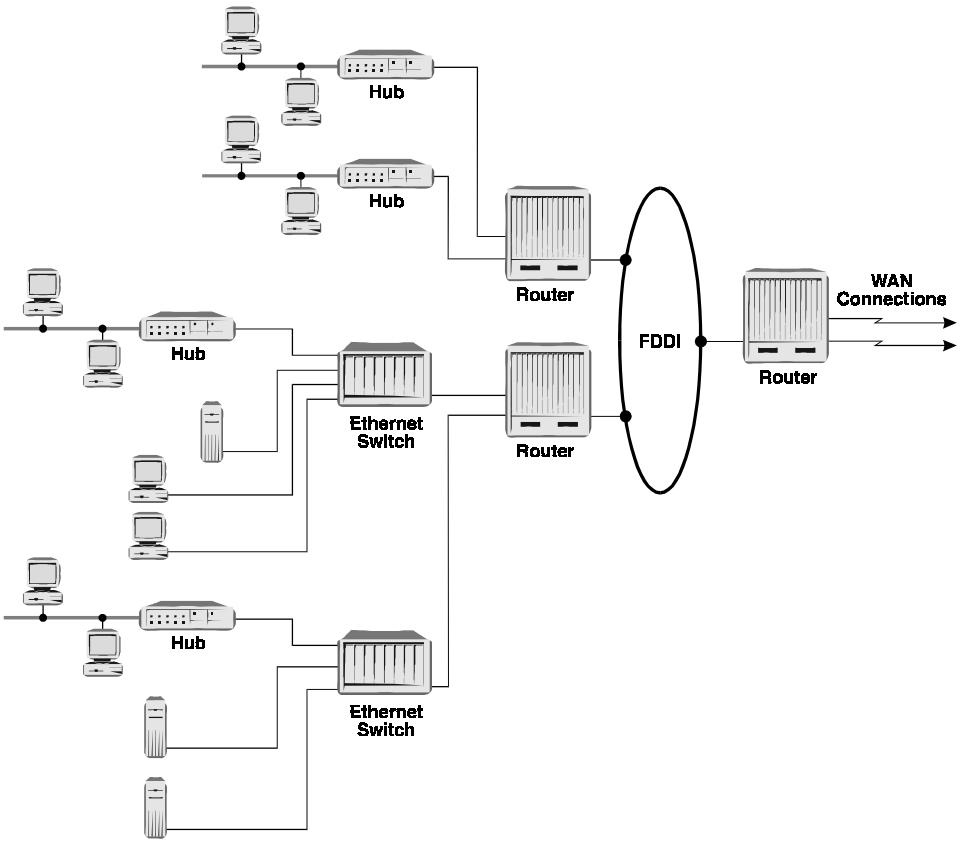

Figure 2: Routers and Switches in a LAN

The Multilayer Switch

We have already described multilayer switching as an intelligent combination of switching and routing technologies which provides a complete and integrated solution for all kinds of LAN infrastructures. Now let's get into some more detail.

A multilayer switch is a device which has multiple LAN ports, over which stations can communicate either by means of Layer 2 packet forwarding (as in conventional LAN switching), or by means of Layer 3 packet forwarding (as in conventional routing). The type of packet forwarding that is used in each case is whichever is needed for any given pair of stations to intercommunicate. In practice, this depends on whether they are members of the same subnet – in which case Layer 2 forwarding is used – or if they are members of different subnets, in which case Layer 3 forwarding comes into play.

Multilayer Switching Devices

A multilayer switching device can be logically viewed as a Layer 2 switching fabric which has a Layer 3 forwarding function attached to it by a high capacity connection. A number of LAN port interfaces are attached directly to the Layer 2 switching fabric. Just like a conventional router, the Layer 3 forwarding function has one or more IP addresses and MAC addresses associated with it that end stations use to send IP packets to it, for forwarding on to different subnets. Similarly, if the Layer 3 forwarding function supports other protocols such as IPX, it would look like an IPX router from the point of view of the end stations.

Let's consider how a pair of stations using IP would communicate across the multilayer switch, via, say, Ethernet segments on each side. The sending station starts out by knowing the IP address of the destination station, but not the MAC address which is needed to send the packet out on the Ethernet. It uses the Address Resolution Protocol (ARP) to determine this. The sending station compares its own IP address with the IP address of the target station and uses the subnet mask which is configured in its software to determine whether the destination station is a member of the same subnet as itself or not.

If the two stations are on the same subnet, the sending station will broadcast an ARP request identifying the destination IP address and requesting the station that owns this address to respond with its MAC address. On receipt of the response, the sending station will cache this address and use it to send the Ethernet packets to the destination. When these packets arrive at the multilayer switch, the Layer 2 switching fabric will look up the destination MAC address to establish which port to forward the packet on, and send it on its way.

If the two stations are on different subnets, the sending station will expect to transmit packets via the "default gateway" (which, confusingly, means a router), the IP address of which is configured into its software. This IP address will actually be the address that refers to the Layer 3 forwarding function in the multilayer switch. So when the station broadcasts an ARP request for the default gateway IP address, the multilayer switch

responds with the MAC address which corresponds to its Layer 3 forwarder. Then when the sending station starts transmitting Ethernet packets with this destination MAC address, the Layer 2 switching fabric will send these packets directly to the Layer 3 forwarder. At this point, the Layer 3 forwarder may need to broadcast an ARP request to obtain the MAC address of the final destination station, which it will store in a cache. As each packet is forwarded, the original MAC destination address (which refers to the Layer 3 forwarder itself) is stripped off and replaced by the new MAC address, which refers to the final destination station. The packet is then sent back into the Layer 2 switching fabric where the MAC address tables are used to direct the packet to the correct output port.

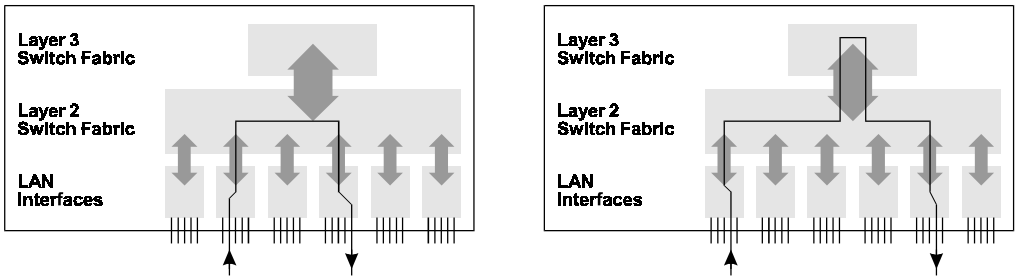

Figure 3: Layer 2 and Layer 3 Switching in a Multilayer Switch

Multilayer IP/IPX Switches

We have explained how a multilayer switch can operate at both Layer 2 and Layer 3 for communications using IP – the Internet Protocol. But IP is not the only protocol that may need to be processed at both Layer 2 and Layer 3 in the LAN. In fact Novell's Internetwork Packet eXchange (IPX) protocol is currently even more widely used in enterprise LANs than IP, and therefore ideally we need a solution for IPX also.

IPX has some different characteristics from IP, particularly in the area of addressing and subnetting. With IP, we have a strictly limited number of possible hosts in a single subnet because the entire IP address has only 4 bytes which have to be split between the network or subnet number, and the host number. By contrast, the IPX address consists of 10 bytes, of which 4 bytes is the network number and the remaining 6 bytes are copied from the LAN adapter's MAC address – therefore we don't have the same limitations as to how many stations can belong to a single subnet. Also, an IP subnet is defined as a group of stations that share the same network or subnet number in their address – which is normally configured into the station's software. Whereas an IPX subnet is effectively defined as all the stations in a broadcast domain, since the stations learn their network number from a server or servers, by means of a broadcast request.

These differences mean that, in principle, we can treat IP and IPX differently in the switched LAN from the point of view of Layer 3 forwarding. We could, for example,

define the entire switched LAN as a single broadcast domain which contains multiple IP subnets but which would contain only a single IPX subnet. In this case there would be no need for any Layer 3 forwarding function for IPX within the multilayer switch, since all IPX communications could be conducted with Layer 2 switching alone.

In practice, however, there may be good reasons for aligning IP and IPX subnetting. Foremost among these is the requirement to keep broadcasts under control. Since the broadcast protocols associated with IPX, the Service Advertising Protocol (SAP) and the Routing Information Protocol (RIP) are prone to generating excessive amounts of broadcast traffic, one way of dealing with this is to break the network up into a number of separate broadcast domains, each of which will contain one IP subnet and one IPX subnet. This preserves the same scheme of subnetting that is imposed when we use a conventional router to connect LAN segments together for IP and IPX communications, only with multilayer switching we can define these broadcast domains logically instead of physically. This technique is widely known under the name “Virtual LANs” – of which more later.

Clearly, if we are going to have both IP and IPX in the LAN, and we choose to divide the switched LAN into multiple separate broadcast domains, then we are going to need Layer 3 forwarding for both IP and IPX in the multilayer switch. This then becomes a “multilayer IP/IPX” switch.

IP and IPX are by no means the only protocols that are seen in today’s LANs. Other protocols include both “routable” and “non-routable” protocols. IP and IPX are both routable protocols which embody the concept of hierarchical Network Layer addressing with the need for Layer 3 forwarding between subnets. Of the other common routable protocols, DECnet is probably the most widely employed, although its use is declining. When considering how to deal with lesser-used routable protocols, we should bear in mind that it isn’t strictly necessary for a multilayer switch to have the capability for Layer 3 forwarding of every routable protocol in the LAN. If we are able selectively to forward these protocols at Layer 2 – effectively with a bridging function – then we may have a satisfactory solution, albeit one which forces the adoption of a single subnet within the switched LAN for such protocols.

Many LANs today are also handling non-routable protocols such as NetBIOS (or NetBEUI), DEC Local Area Transport (LAT) or IBM DLC. Since none of these protocols recognizes the concept of Layer 3 forwarding, they must be bridged at Layer 2 within a multilayer switch.

Multilayer Switches and Virtual LANs

We referred earlier to the technique of defining broadcast domains within the switched LAN as “Virtual LANs” with the promise of some further explanation. We can use the term broadcast domain or VLAN interchangeably since they really mean the same thing, so in exploring this topic further let’s use the term VLAN for the sake of brevity.

At the simplest level, VLANs are defined (via network management) as groups of ports on LAN switches which can mutually exchange unicast and broadcast packets. A

broadcast packet sent to the switch via one of the ports that belongs to the VLAN will be copied by the switch to all of the other ports that belong to the VLAN. Some LAN switches allow VLANs to be defined that include ports on multiple switches, although the ability to do this may depend on some additional protocols being used between switches to communicate VLAN membership details.

It is perfectly possible to operate a LAN based on multilayer switching without defining VLANs, but there are a couple of reasons why VLANs are likely to be needed in most network installations:

- Many existing networks rely on the Layer 3 forwarding function (in conventional routers) to provide a degree of security and access control. While the Layer 3 forwarder in a multilayer switch may be perfectly capable of providing the same degree of security and access control, in the absence of VLANs it would be easy for a user to bypass the Layer 3 forwarding function by reconfiguring his/her workstation's IP address to have the same subnet number as the networked resource he or she is trying to attack. Then the user would be able to access this resource via the Layer 2 switching fabric alone, where the checks cannot be made.

- Without the definition of VLANs, a LAN based on multilayer switching would be a single, large broadcast domain. Depending on the number of stations in the LAN and the types of protocols in use, this could result in problems due to excessive broadcast traffic. Some LAN switches can deal with this using intelligent broadcast control techniques, but if these are not available then VLANs must be used to counter the problem by breaking the network up into a number of separate broadcast domains.

Now let's see how we would use VLANs in practice with multilayer switching. To keep things simple, we'll start with the assumption that each IP subnet would have its own VLAN, and that each VLAN will contain one, and only one, IP subnet.

All we have to do is to decide which ports on which multilayer switches are going to belong to which IP subnets. In an existing routed LAN this is easy, since each LAN segment that is connected to a router will already have its own subnet identity. When we connect a LAN segment to a multilayer switch, we may want to break the segment into several smaller segments in order to improve performance. So we may have a number of ports on each multilayer switch which relate to the same subnet identity. We simply need to define a VLAN around each group of ports that have the same subnet identity, and then "connect" this VLAN to a logical "router port" within the Layer 3 forwarder in the multilayer switch. All this is configured, of course, via the network management console.

Defining VLANs in this way means that we have addressed the security issue, because users cannot communicate with any resources that are not in the same VLAN as themselves without passing through the Layer 3 forwarding function. We have also addressed the broadcast issue, since our broadcast domains are now no larger than they were when we used conventional routers to connect shared LAN segments together.

However, this approach to VLANs does not enable us to take advantage of one of the most talked-about benefits of VLANs, the ability to handle moves and changes without having to re-configure IP addresses in end stations. In the example we described, if a user moves his PC from one switch port to another, and the new switch port is not defined as belonging to the same VLAN as the old switch port, then the user is going to need a new IP address to match the subnet assignment of the switch port he is now connected to – otherwise he will not be able to communicate across the LAN.

One solution to this is to use a network management application which tracks moves and changes (by observing the MAC addresses learned on each switch port and their corresponding IP or IPX addresses) and which then re-configures VLAN boundaries automatically to ensure that users maintain the connectivity they need, without having to issue them with new IP addresses. To maintain security, this kind of application must be able to request confirmation of any changes to VLAN boundaries it is proposing to make.

Another rather simpler solution is to place all the workgroup segments and subnets in one big VLAN, and the resources that need to be secured in another VLAN. Then users can move around between any of the workgroup segments without needing changes of IP address and without there being any need to update VLAN boundaries. But all traffic to and from the secured resources is forced to pass through the Layer 3 forwarding function, where the appropriate security checks can be applied. The only potential area of concern with this solution is the size of the broadcast domain that contains all the workgroup segments, and the possibility that excessive broadcast traffic might cause problems. The extent to which this applies in any particular case can easily be established by experimentation, or by extrapolation from measurements of broadcast traffic on the existing network.

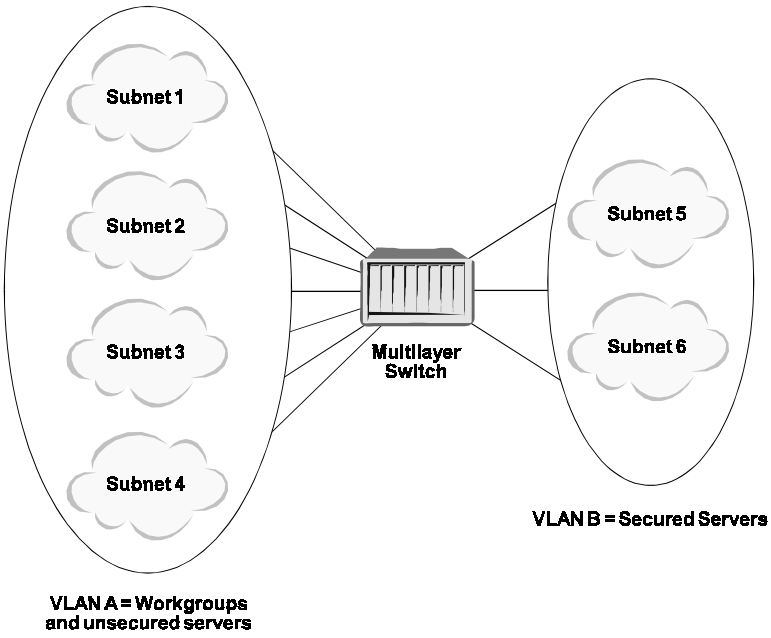

Figure 4: Combining Ease of Moves and Changes with Secure Access Control

Multilayer Switches and Routing Protocols

We have described a multilayer switch as a device which combines Layer 2 switching (for example, Ethernet switching or Token Ring switching) with Layer 3 forwarding (for example, IP or IPX routing) so as to provide a complete solution for the needs of high performance LANs. However, we have not so far made reference to the routing protocols which are needed to update the routing tables used to make Layer 3 forwarding decisions.

Routers employ a variety of routing protocols to exchange information about network topology and the reachability of subnets. Examples of standard routing protocols are Routing Information Protocol (RIP) and Open Shortest Path First (OSPF). The information which is exchanged by these protocols is processed by the router to maintain the routing tables which map IP address prefixes to router ports.

The Layer 3 forwarding function in a multilayer switch needs exactly the same kinds of routing tables as a conventional router. Furthermore, multilayer switches need to be able to operate in networks that contain conventional routers. It follows that multilayer switches must be able to participate in the normal exchange of information that takes place between routers via the routing protocols.

In practice there are two different approaches to the support of routing protocols in multilayer switches. With the “self-contained” approach, the Layer 3 forwarding function in each multilayer switch engages in the routing protocols just as if it were a

conventional router. In the “route server” approach, a central function in the network engages in the routing protocols on behalf of one or more multilayer switches, and uses some new and additional protocols to communicate routing table updates to the multilayer switches.

The route server approach has the potential to offer a lower cost solution in large networks with many switches, because it can reduce the complexity of the multilayer switch by moving responsibility for the routing protocols elsewhere. However, in most practical situations large networks can be built by combining simple Layer 2 switches with a much smaller number of multilayer switches. When two stations that belong to different subnets are inter-communicating across a number of switch hops, only one of the switches in the path needs to be a multilayer switch to perform the Layer 3 forwarding function, so it makes sense to use conventional Layer 2 LAN switches in the workgroup, with multilayer switches in the backbone. In this case, the cost saved by simplifying the multilayer switches may be more than offset by the complexity that comes with a separate route server function. Furthermore, there are no standards in place for the protocols used to distribute routing table updates from the route server in frame-based networks, so these kinds of solutions are highly proprietary.

Whichever of these two techniques are employed – self-contained routing protocol support in each multilayer switch, or distributed routing based on route servers – the Layer 3 forwarding function of multilayer switching makes use of conventional routing protocols and is therefore compatible with existing routers. This means that multilayer switches can be deployed in networks that contain routers, and that the multilayer switches will look just like peer routers from the point of view of any existing routers. Schemes that provide fault tolerance based on multiple redundant routes can therefore be implemented with any mix of routers and multilayer switches.

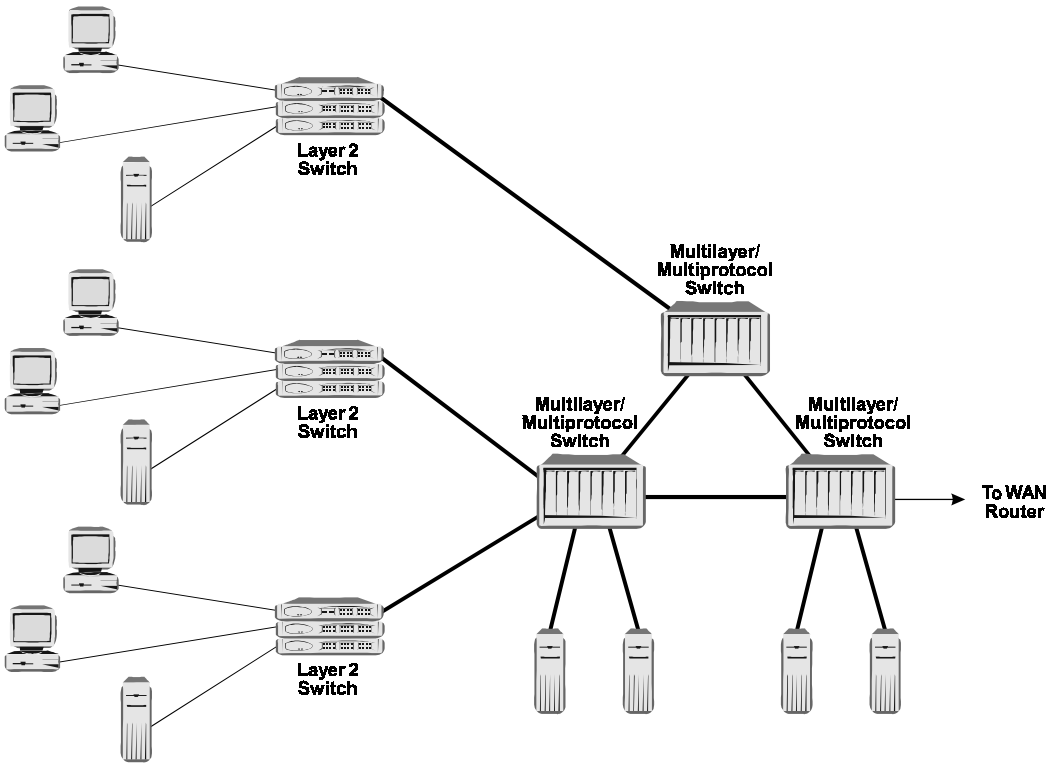

Figure 5: Combining LAN Switches with Multilayer IP/IPX Switches

Multilayer Switch Performance

LAN switches generally achieve very good price/performance by applying hardware-based frame forwarding techniques to the process of moving packets from one LAN segment to another. This is relatively easy to engineer, since the forwarding decision is based on a simple MAC address (or RIF field) look-up table, and the forwarding process involves no change to the content of the packet.

Routers, on the other hand, generally exhibit much lower forwarding rates at considerably higher cost. Three reasons for this are:

- The packet forwarding decision in a router is considerably more complex than that in a LAN switch.
- A number of changes have to made to the contents of each packet received from one LAN segment before it can be forwarded to another – for example, the entire MAC packet header has to be replaced, and the Time-To-Live field has to be updated in the IP packet header.
- If packet filtering is used to implement network security, each packet may have to be compared with a number of complex filtering criteria.

This additional complexity has meant that many routers implement the packet forwarding process mostly or entirely in software – which accounts for the inferior price/performance characteristics of routers compared with LAN switches.

The Layer 3 forwarding function in a multilayer switch is subject to the same complexities and packet processing responsibilities as exist in a conventional router. Therefore it is not safe to assume that a product described as a "multilayer switch" automatically offers higher performance than a conventional router for forwarding packets between subnets. Performance depends very much on the manner in which the Layer 3 forwarding function has been implemented.

If the Layer 3 forwarding function is implemented entirely in hardware, then a multilayer switch should be able to show the same high performance whether it is switching at Layer 2 or Layer 3. If, as in most routers, the Layer 3 forwarding function is implemented largely in software, then the multilayer switch is likely to perform no better than a typical router when forwarding at Layer 3, and in many cases its performance may be substantially worse.

Real multilayer switch products are likely to show wide variations in the performance of their Layer 3 forwarding. In this area particularly, all switches are not created equal.

Conclusion

In this paper we have described multilayer IP/IPX switching as a complete solution for scalable, high performance, cost-effective LAN infrastructures. The advantages of this approach over a classical router-based backbone include much improved price/performance characteristics, greater flexibility to cope with growing traffic loads and changing traffic patterns, elimination of the need for IP address re-assignment during moves and changes, and fewer different kinds of devices to manage in the network.

The approach we described to multilayer IP/IPX switching complies with the existing standards for LAN switching and for routing, and does not introduce any new or proprietary protocols into the LAN. Therefore multilayer IP/IPX switches can be readily supported in multivendor networks, and furthermore, they are compatible with equipment that is installed in existing networks. This means both an easy migration path to a more effective networking solution – with protection of existing networking investments – and the reassurance that this solution will not create undesirable lock-in to a single vendor.

There can be little doubt that multilayer IP/IPX switching is the right way forward for large and medium-sized corporate and enterprise LANs. Where routers are already installed in the LAN, further investment in router upgrades can be frozen, and the additional capacity needed in the LAN can be provided by multilayer IP/IPX switches working alongside the routers. As traffic loads continue to grow, changing traffic patterns bring new stresses onto the backbone, and new classes of applications demand support for real-time traffic on the LAN, multilayer IP/IPX switching will provide the simplest and most cost-effective answers to the problem.

Madge Networks

Americas
2310 North First Street
San Jose,
CA 95131-1011
United States
Tel +1 408 955 0700
Fax +1 408 955 0970
http://www.madge.com

Asia, Australia & New Zealand
12/F Li Po Chun Chambers
189 Des Voeux Road,
Central,
Hong Kong
Tel +852 2593 9888
Fax +852 2519 8022
http://www.madge.com

Europe, Middle East & Africa
Knaves Beech Business Park
Loudwater, High Wycombe
Bucks HP10 9Qz
England
Tel +44 1628 858000
Fax +44 1628 858011
http://www.madge.com

Japan
Mita NN Building
1-23, Shiba 4-chome
Minato-ku, Tokyo 108
Japan
Tel +81 3 5232 3281
Fax +81 3 5232 3208
http://www.madge.com