How Did LANs Evolve to Multilayer Switching?

Anritsu

How Did LANs Evolve to Multilayer Switching?

Summary

In the past, internetworking devices could easily be categorized as either bridges, routers, switches, or hubs. Recently multilayer switches have emerged with capabilities that cross these traditional boundaries with abilities to meet new customer needs:

- Improved IP performance (speed, throughput, latency) for Internet and intranet applications,
- Improved reliability (more redundancy and simpler to employ, monitor, troubleshoot), and
- More capacity for scalability as network growth occurs (users, applications, traffic).

There is confusion in the marketplace, however, due to the variety of new multilayer switch products and inconsistent terminology used to describe them. Understanding how LAN architectures have evolved and how multilayer switches work helps clear up that confusion.

Background

Intranets and the Internet are becoming more important to the way in which enterprises develop new applications and conduct their businesses. This is motivated by the need to improve interconnection of office and remote PCs, provide better access to mainframe data, roll out client/server applications faster, and share information throughout the organization. As a result, the number of people connected to networks based on IP traffic and the scope of their usage are growing rapidly in nearly all organizations.

Traffic volume is increasing not only from heavier traditional uses (electronic mail, printing, and file sharing) but also from new or enhanced applications for collaboration (schedulers, project management, groupware) and information transfer (data warehousing, Internet access, and browser access to enterprise information). In fact, intranets are not just increasing the amount of traffic, they are also leading to entirely different patterns of traffic flow.

In Figure 1(a), for example, workgroup LANs were originally set up to enable local sharing of resources and then interconnected with relatively slow links. The old rule of thumb said that about 80% of the traffic would stay within the workgroup, and not more than 20% went out into the rest of the enterprise network. Traditional hubs and routers were adequate to handle this type of traffic.

The Internet Influence The influence of Internet technology combined with recent business trends such as downsizing, organizational flattening, cross-functional teams, and more general sharing of enterprise data have transformed this old-style network into Figure 1(b). Data warehousing and intranets have brought renewed interest in centralizing resources onto enterprise servers. In general, the amount of traffic among workgroups is increasing as a result of sharing enterprise resources (data, people, compute power) more effectively. These are rapidly driving traffic expectations to the opposite mix: only 20% stays local and 80% goes out of the workgroup.

Figure 1. Increasing Complexity of Enterprise Networks

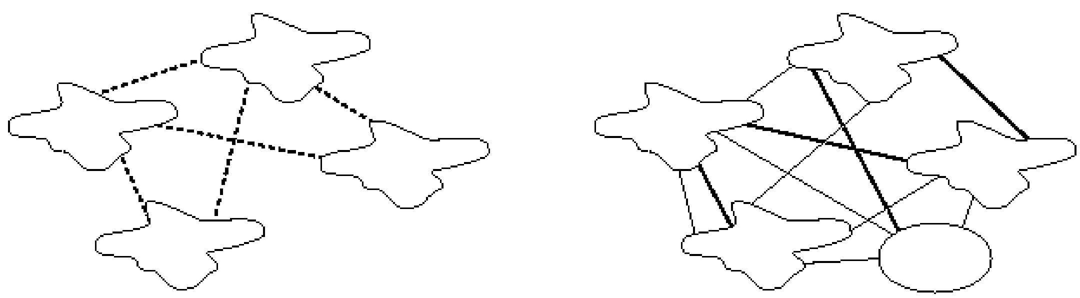

(a) Independent workgroups, loosely interconnected; most traffic stays local within each workgroup

(b) Workgroups use common server resources, so more traffic crosses workgroups or addresses central servers

With the increasing complexity and importance of enterprise networks, two additional issues gain major visibility from network managers: policies and performance. Policy issues include who has access to what, what bandwidth is made available to various users, how and where are things stored, what addressing schemes are used, and what protocols will be used. When networks are small or relationships relatively simple as in Figure 1(a), these policies can often be set and supervised in an ad hoc manner. Similarly, performance criteria such as throughput or reliability are straightforward. With the larger, more complex networks of Figure 1(b), organizational dependence increases. This requires more sophisticated policy management and higher levels of performance and reliability, which can be increasingly difficult to provide as the network gets larger. Furthermore, most of the new traffic is based on the IP protocol, and it makes sense now to have network devices that are optimized to handle such traffic.

Modern LAN Architectures

Simple Shared Media A convenient way to understand today's architectural options is to look at how they have evolved to meet growing user needs. The first local area networks grew out of multipoint device links, and consisted of multiple computers tied together by shared transmission media, as in Figure 2(a). To make cabling more manageable and to serve clusters of devices, this quickly evolved into the form of Figure 2(b), where the box labeled "hub" functions as a medium access unit that meets the physical-layer signaling requirements for multiple devices. The detail of how a hub does its work depends on the underlying technology – for example, Ethernet or Token Ring. Shared media-and-hub approaches are particularly useful for a workgroup of people whose desktop PCs need to share local resources such as printers, files, or application servers.

Figure 2. Shared Media LAN

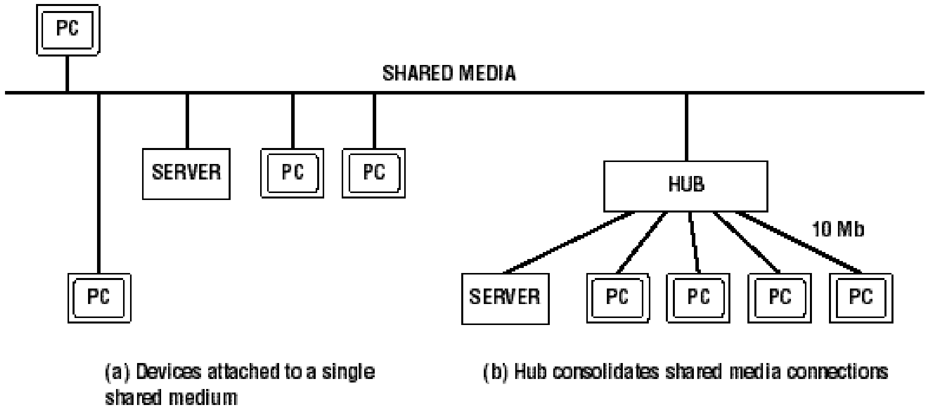

Simple interconnection of workgroups that need to communicate or share resources with each other usually starts by extending laterally at Layer 2 of the network, as in Figure 3. ("Layer 2" refers to the layer in the communications protocol that contains the physical address of a client or server station. It is called also called the data link layer or MAC layer and contains the address inspected by a bridge, switch, or PC NIC.)

This approach creates a large, "flat" network – its ability to scale up is limited by the transmission medium, the sharing discipline, addressing, or performance requirements. After reaching one of these limits, networks in the past were usually divided into pieces that get connected at Layer 3 by a router in order to allow further growth. ("Layer 3" refers to the layer in the communications protocol that contains the logical address of a client or server station. It is also called the network layer and contains the address (such as IP or IPX) inspected by a router that forwards it through the network.)

In the early days of LANs, Layer 2 switches did not yet exist or were very expensive. The hubs and shared media ignore the traffic addresses: all traffic appears everywhere on the network and creates a strain on the network's capacity.

Figure 3. Flat Interconnected Shared Workgroup Network

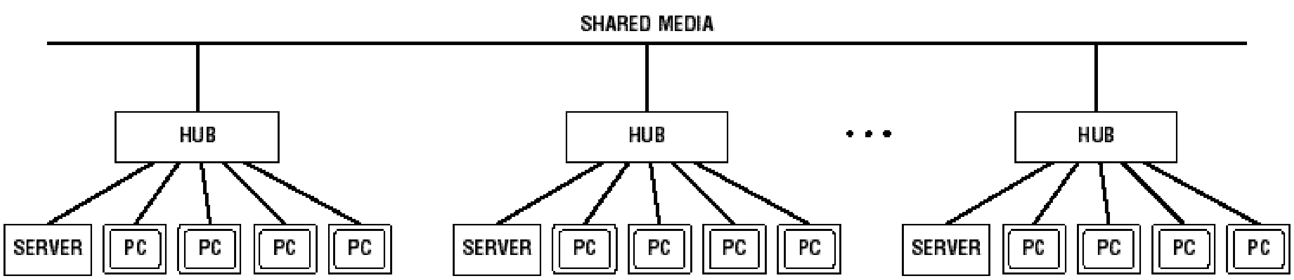

How Routers Help Figure 4 shows how networks are extended by using a router to create a hierarchy. A router calculates routes and forwards traffic only to the node containing the correct Layer 3 (or network) address — for example, based on its IP or IPX address. This allows each workgroup to have much more bandwidth, prevents broadcast traffic from flooding the entire network, and provides some security. For IP traffic, users are divided into different IP-address subnets that are connected at Layer 3 by routers. Traffic using Layer 2 (MAC) addressing is constrained to stay within the subnet.

Multiple buildings or sites belonging to an enterprise can be connected by creating a network “backbone” that links to every router at every site. As long as these links stay within a local area scope, high-speed LAN technologies such as FDDI and ATM (and eventually Gigabit Ethernet) can be used for the backbone, and performance is not reduced by slower wide-area link technologies.

Figure 4. Hierarchical Routed Network

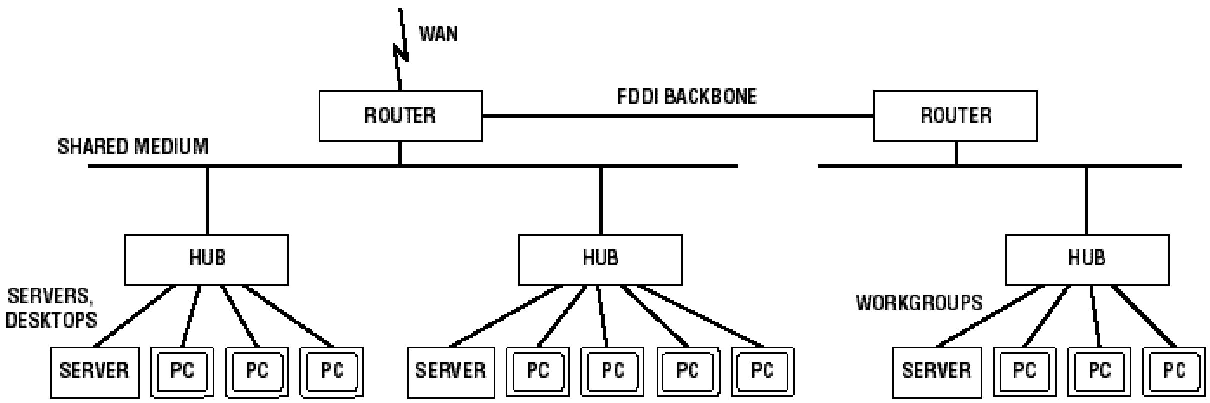

Routers provide additional policy advantages for the network manager by providing packet filtering for security and giving wide area access. But in performing these complex functions they also introduce processing delays that limit throughput and create higher latency. This is acceptable as long as the amount of traffic through the router is relatively small.

Switching Technology Enters the Scene The shared media approach of Figure 3 and Figure 4 LANs worked well for older PCs that could only drive traffic onto a LAN at a fraction of its carrying capacity. But client PCs are now so fast that a single PC can exceed the entire capacity of Ethernet or Token Ring; a single server or high-end workstation can exceed a Fast Ethernet link’s capacity. The next step in LAN architecture evolution, Figure 5, applies switching technology to reduce contention for the shared medium. A switch operates at Layer 2, directing traffic by its MAC addresses rather than routing by the network address at Layer 3. Whereas all the ports of a hub shared the same network segment, every port on a switch is a discrete segment and makes the full network bandwidth available there.

Today’s switches have a high enough port density that switch ports can be made available to individual PCs, giving each PC a full 10 or 100 Mbps bandwidth. To avoid simply pushing the previous performance bottleneck on shared network segments along to the switch/router interface, switches typically use a faster LAN “uplink” to the router such as FDDI or Fast Ethernet. In the future, more of these uplinks will be Gigabit Ethernet. Some switches offer multiple uplinks for additional capacity or to provide redundancy; load-sharing and aggregation over multiple links comes next.

Figure 5. Switching Reduces Shared Media Contention

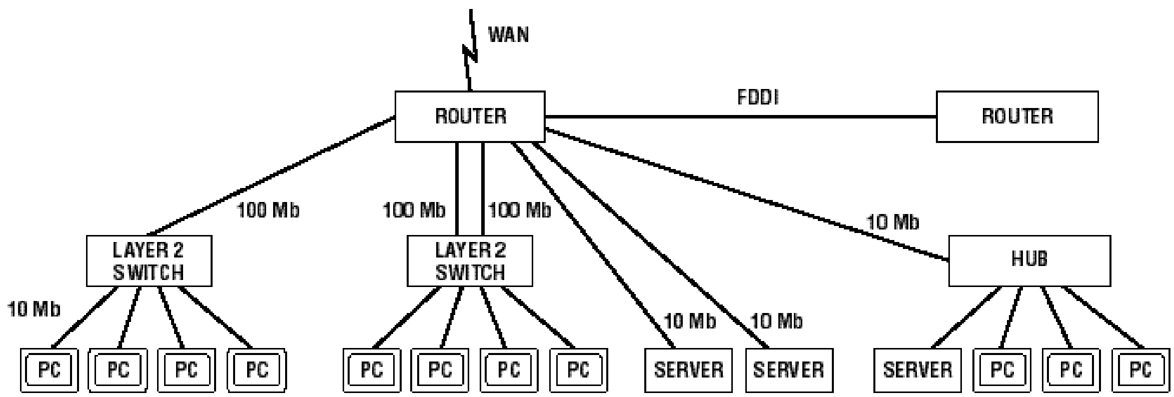

VLANs Simple Layer 2 switches provide basic bridging capabilities – they learn where each MAC address is and they forward traffic among various segments. Recent switch products have become more sophisticated, allowing each connected device to be placed into a logical group according to its physical point of connection (switch port), MAC address, or network protocol type. Such groups are called virtual LANs or VLANs. VLANs give network managers much greater flexibility to provide resource access and security, as well as good performance, to users that belong to the same functional group or department but are physically distributed around an enterprise. VLAN features often provide additional policy mechanisms that help support larger and more complex networks.

Dedicated Server Links Become Necessary Figure 5 also shows servers being moved out of individual workgroups and given direct links to the router so that they can support users that are located in multiple groups. This is a growing trend as organizations develop intranets to share information. Direct router connections on dedicated links give the servers greater communication bandwidth as well. This is crucial because server links often become bottlenecks as networks move toward an intranet architecture.

Backbones Collapse As more ports are added to connect increasing numbers of servers, segments, and switched subnets, the router begins to function as a "collapsed backbone". This term means that more and more local forwarding and interconnection services are provided within a single box, absorbing the discrete backbone links that were between boxes. The scope and complexity of routing functions make it difficult for traditional routers to be cost-effective and maintain good performance in this role, however. They typically cannot process traffic fast enough to keep up with the growing volume of activity, and their port costs are often much higher than that of switches.

Multilayer Switch Solutions

Three factors combined over the last two years to fuel the evolution of multilayer switching:

- Users need to get beyond the performance bottleneck of collapsed backbone routers and avoid the high cost of expanding them.
- IP traffic from intranet and Internet applications has increased dramatically.
- ASIC (Application-Specific Integrated Circuit) densities increased enough to allow economic implementation of complex routing functions directly in high-speed hardware instead of using the slow software techniques of traditional routers.

What Is It? Simply stated, multilayer switching is the combination of traditional Layer 2 switching with Layer 3 protocol routing in a single box, usually through a fast hardware implementation.

Multilayer switches aren't called "routers" because they are so much faster and less expensive. Also, some multilayer switches lack the modularity, flexibility, and port density usually associated with routers. Many are more limited than routers in the variety of traffic and routing protocols they support. Finally, multilayer switches generally don't support all the WAN interfaces handled by traditional routers.

Nomenclature Confusion Because multilayer switching is so new, there is still no industry standard on nomenclature. So there is little agreement among vendors, analysts, and editors about the specific meaning of terms such as multilayer switch, Layer 3 switch, IP switch, routing switch, switching router, and wirespeed router. Compounding the confusion, these different terms don't reflect differences in architecture or function as much as differing editorial and marketing policies of the involved parties.

Taking Advantage of Multilayer Switches The introduction of multilayer switching products leads to the next generic architecture shown in Figure 6. Some routing functions (which and how many depends on particular vendors) are built into the multilayer switch so that the traditional router can be moved out to the WAN boundary of a LAN. This saves the costs of further investments in expensive routers while maintaining existing edge functions such as wide-area network access. At the same time, network performance increases dramatically because most of the local traffic routing is now handled by the multilayer switch using high-speed hardware.

Large enterprises will not likely convert a huge legacy network to a multilayer switch architecture all at once. Existing equipment investments, reliability considerations, and differing traffic requirements on Ethernet and ATM subnetworks often create a need for multilayer switches to provide FDDI and ATM links to other portions of the existing enterprise network. Such links are shown in the right side of Figure 6. Multilayer switches from some vendors provide Ethernet-type links only, and are thus not suitable for these enterprise requirements.

Figure 6. Multilayer Switches Push Routers to the Network Edge

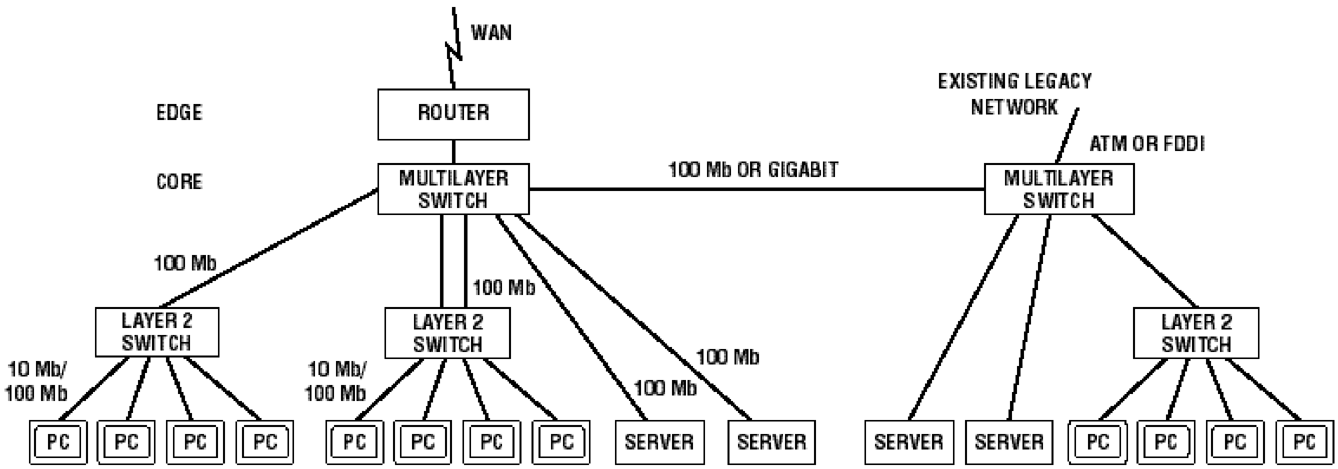

Collapsing the Backbone Again As enterprise networks evolve, or entirely new networks are deployed, the architecture shown in Figure 7 can take over. In this case, the backbone collapses into a single multilayer switch box that handles all traffic forwarding internally for maximum performance and reliability. If needed, several multilayer switches can be used, with long-distance Gigabit Ethernet fiber links between them. ATM and FDDI links can still be accommodated via appropriate modules installed in the multilayer switches. Existing routers still operate at the edges of the network providing WAN connections.

Figure 7. Multilayer Switch Provides Collapsed Backbone Solution

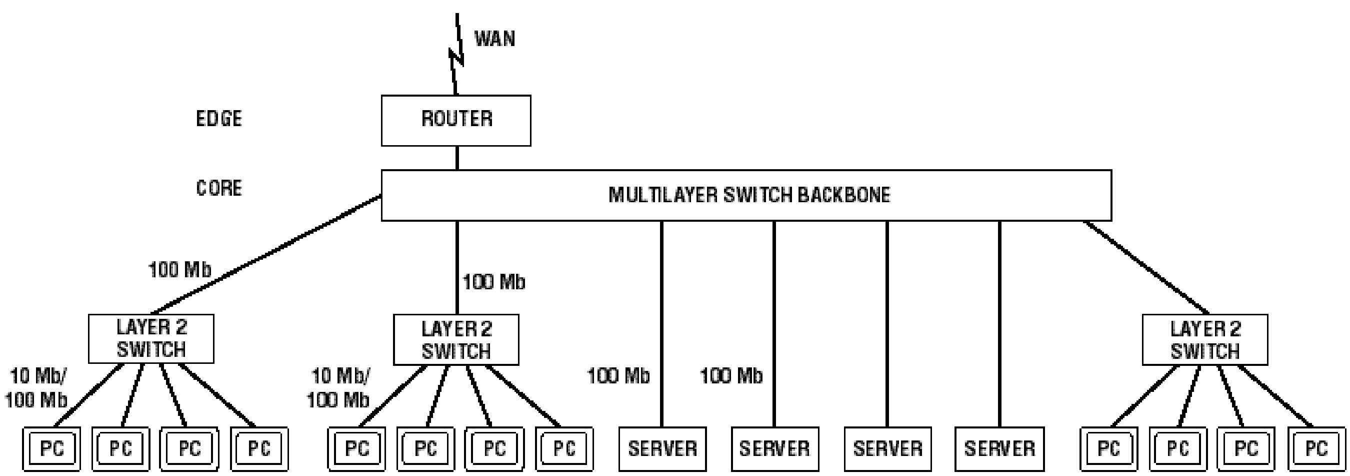

The multilayer switch serves multiple aggregating and connecting functions, using Layer 2 switching when possible to improve performance.

Adding the WAN Links As was noted earlier, there is no technical reason why multilayer switches cannot replaces routers entirely provided they have the required WAN-related features of the router and are able to handle all the types of traffic needed by the network.

The hybrid nature of this type of sophisticated multilayer switch solution is illustrated in Figure 8. On the left side the multilayer switch backbone replaces the previous router by adding a WAN interface, while on the right side the previously separate Layer 3 and Layer 2 switches are replaced by a single multilayer switch.

This type of integrated solution from a single vendor can greatly improve performance, reliability, and ease of management all at once. The viability of this architecture depends on a product's routing capabilities, WAN interface capabilities, performance characteristics, price per port, and influence on overall network ownership costs.

The Desktop Question Should multilayer switching, including its routing functions, be pushed all the way to the desktop? This can be economically justified in situations where an individual user has significant cross-subnet (or cross-VLAN) traffic and requirements for high bandwidth or high quality of service. Appropriately designed multilayer switches with high enough port density can forward this traffic at wire speed and provide excellent performance even for routed traffic – a definite bonus at the desktop. These kinds of connections are shown in the lower right portion of Figure 8.

Figure 8. Multilayer Switch Meets Multiple Network Needs

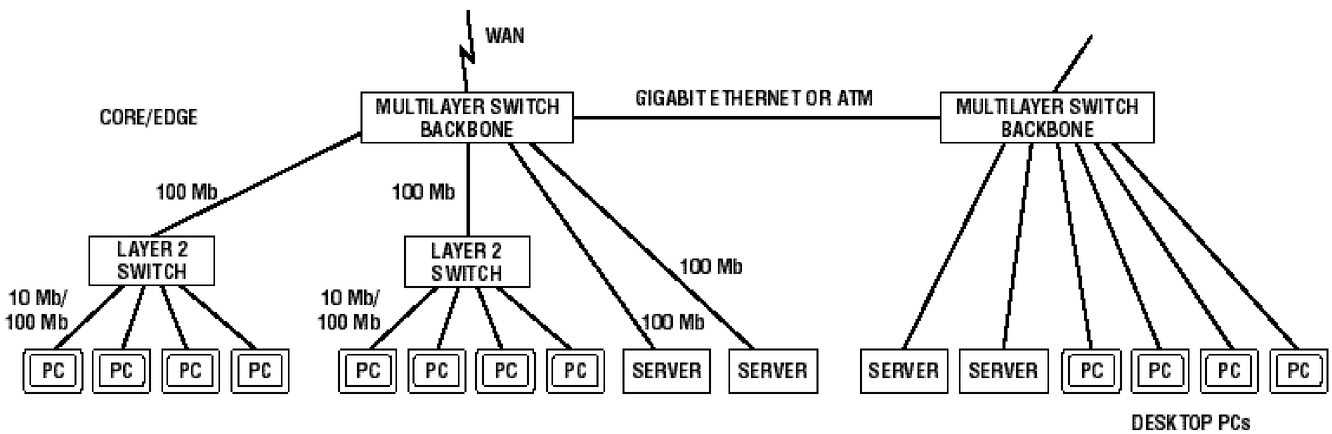

Multilayer switch with WAN interface displaces router entirely.

The MultiFlow™ 1000 Multilayer Switch from Anritsu

Based on extensive bridge and router development experience combined with sophisticated custom ASIC technology, the Anritsu MultiFlow 1000 multilayer switch has been designed for low latency and high throughput to meet customer needs for high-performance networks.

It capitalizes on the growing dominance of IP traffic in enterprise networks, using special-purpose hardware to maximize IP routing performance. The MultiFlow 1000 is a wire-speed Layer 2 switch that includes full packet-by-packet Layer 3 routing capabilities to forward IP, IPX, and AppleTalk traffic and run standard routing protocols including RIP and OSPF. Full IP multicast capabilities are included for optimum multimedia support. Built-in VLAN capabilities and compliance with IEEE 802.1Q VLAN tagging provide the network segregation features and industry VLAN compatibility that network managers need.

The MultiFlow 1000 is designed to provide customers with a highly reliable, scalable, and easy-to-manage network infrastructure. Extensive support for SNMP-based network management (MIB II, Bridge, RMON, and Enterprise MIBs), Web-based management, and port mirroring make it easy to monitor and analyze the network. Built-in Layer 2 and Layer 3 switching capabilities mean that network design can be simpler, with fewer different kinds of network elements, so that the whole network is easier to understand and to extend.

The MultiFlow 1000 is a chassis-based switch so it provides configuration flexibility and allows for easy future network enhancement. It is a highly-capable device with a fast backplane capable of forwarding IP packets at 500 Kpps and switching at wire speed on each interface. Interface modules support 10BASE-T, 100BASE-TX, 100BASE-FX, 100VG-AnyLAN, 1000BASE-SX Gigabit Ethernet, FDDI, and ATM connections. The router module has its own slot and can be added separately, allowing users to start with lower-cost Layer 2 switching and upgrade to a multilayer switching solution later.

This combination of capabilities makes the MultiFlow 1000 an outstanding product to meet customer needs in a variety of applications:

- As a 100 Mbps or Gigabit backbone, it can connect workgroups across a campus while providing high-performance server links using Fast Ethernet, Gigabit Ethernet, ATM, or FDDI.
- Within a workgroup, 100Mbps channels can be allocated to hubs, Layer 2 switches, servers, or dedicated to individual desktops as required.
- Where an existing router acts as a collapsed backbone, it can be used to push the router to the network edge for WAN access, relieving local router performance bottlenecks and forestalling further router port investments.

Recommendations and Summary

As scope and performance requirements for enterprise networks grow and products proliferate, network design seems to become more complex. The following guidelines can be helpful in sorting through all the options to achieve higher throughput, reduced latency, more scalability, and easier management:

1. As more capacity is needed at the desktop, segment large 10 Mbps shared networks by connecting hubs to Layer 2 switch ports. Desktops or "power users" that need additional network capacity should receive dedicated switch connections. Move high-end workstations to shared 100 Mbps segments, and then on to dedicated 100 Mbps switch ports as needs grow. Use multilayer switches when addressing, broadcast, or other scalability issues outgrow the capabilities of Layer 2 switches.

2. Move any servers still on shared 10 Mbps segments to 10 or 100 Mbps switch ports based on their capability and access demands. Again, use 100 Mbps multilayer switch ports connected directly to servers to get beyond the limitations of Layer 2 switching and provide better IP traffic handling.

3. Avoid adding uplink ports from Layer 2 switches to existing routers (creating or expanding collapsed router backbones). Instead, put multilayer switches in the LAN backbone to aggregate traffic from Layer 2 switches or to replace Layer 2 switches. This pushes any existing routers to the network edge – for example, to maintain WAN access — but forestalls further router investment.

In short, multilayer switching products with high-speed hardware IP routing offer significant improvements in scalability and performance over simple Layer 2 switches (that don't handle IP traffic routing) or traditional routers (that have low throughput and high latency).

Notes

Notes

Anritsu Company
1155 East Collins Boulevard
Richardson, TX 75081
Tel: 800-ANRITSU (267-4878) or 972-644-1777
Fax: 972-671-1877
E-mail: moreinfo@naro.us.anritsu.com
URL: www.anritsu.com

Anritsu Corporation
5-10-27, Minamiazabu, Minato-ku
Tokyo, 106 Japan
Tel: +81-3-3473-7203 (Direct)
Fax: +81-3-3442-0235
URL: www.anritsu.co.jp/index_e.html

Printed in the U.S.A.
PN MultiFlow-E-E-1.1