LANswitch Layer-3 Switching Overview

Madge Networks, Confidential Draft, August, 1996

Catalog No.142721 Rev.A

LANswitch Layer-3 Switching Overview

In this paper, we will describe current networks and explain how and why LAN internetworking is implemented. We will define the role of the router in a network as well as the alternative of using Layer-3 switching in the LAN instead of a router. And lastly, we will suggest the configuration of a flat network design which combines a flat address scheme with broadcast control. Some organizations will be unwilling to implement a flat network design, either because the change to their existing network structure would be too radical and distressing, or because network applications and/or architecture is Layer-3 dependent. Therefore, a high-speed intelligent interconnection is essential to meet network structure requirements.

Router routes according to Layer-3 addresses

Layer 5	Application	
Layer 4	Transport	
Layer 3	Network	IP, IPX DECnet, and others
Layer 2	MAC	Ethernet, Token Ring, FDDI, NetBIOS
Layer 1	Physical	

The main difference between Layer-3 and Layer-2 forwarding decisions is the address upon which these decisions are based. A MAC layer device looks at the Ethernet or Token Ring address, while a Network layer device looks at the IP, IPX, or any other network layer protocol address in order to determine where to forward the packet. The Layer-2 device makes a binary decision - the address is either local or non-local. In contrast, the Layer-3 device makes much more intelligent forwarding decisions based on knowledge of network topology and interconnections. In addition, the Layer-3 device must have the intelligence to analyze various types of information and the computing power needed to insert and process new field values in the data packet. This results in two kinds of forwarding devices: a simple and very fast MAC device (hardware based) and a complex, somewhat slower Network device (software based). Since a traditional router already has adequate computing power, it can be used for policy based filtering, protocol conversion, and different media connections. The router’s role is important when packets must cross LAN-WAN borders.

Routers keep an updated picture of the networks they support. To do this, a router periodically communicates its local connections to all the other Layer-3 devices so routing tables throughout the network can be dynamically maintained. This requires yet another level of intelligence and computing capability.

Migrating Router-Based Topologies to High-Speed Networking

The popular way to improve network performance has been to segment the enterprise network into many, smaller subnets which were interconnected by routers in a collapsed or distributed routed topology. Until recently, there was no better alternative for reducing LAN congestion and eliminating broadcast storms.

In the beginning, routers offered a straightforward internetworking solution. Soon thereafter, the demand for network based services grew by leaps and bounds, and routers were called upon to cope with this rapid expansion in the number of network users. As a result, network topologies based on routers became even more prevalent, and customers invested large percentages of their networking budgets in these highly complicated and intelligent devices for segregating LANs. However, routers require the customers to re-configure the network according to the router's architectural restrictions rather than transparently install it for enhanced performance. That is why today, so many enterprise networks are built around logically defined subnets that are tightly coupled with physical router ports.

These networks were designed according to Layer-3 requirements, restrictions and limitations rather than configured in the way that would best meet application needs. The advantage of the routing approach was advanced internetworking and filtering services, while the disadvantages included limited flexibility, physical constraints and high latency, lower performance connections. Over the years, routers have become more costly and more complex devices. The router's natural application is internetworking - connecting different networks together (e.g., LANs and WANs, or Ethernet and Token Ring). They do not operate at high-speed or with low-latency; they are very expensive, and their performance does not support the new generation of data-intensive and delay-sensitive applications (i.e., voice & video). However, the majority of enterprise networks are already configured according to the router paradigm so any migration strategy must consider the current situation as well as the future target.

Switching to ATM pushes routers to the edge

LAN Switches were initially introduced to speed up LAN traffic between the growing number of segments. The majority of LAN traffic does not need the resourceful router and its overhead. Therefore, switches can deliver transparent, fast and efficient data transfer within and between physical network segments.

LAN switches are not bound by upper layer restrictions and therefore, are better able to support network applications that require high-speed, low-latency transmission and reduced physical dependencies. LAN switches exhibit virtually no performance variance whether servicing few or many segments and users. They can also dedicate guaranteed bandwidth to critical network resources without sacrificing switching speed advantages.

With LAN switches, there is no restriction on the number of users that can be configured in a logical networked group, nor is there any limitation on the physical segmentation within the group. This makes the LAN switch a better tool for creating network topologies that match user requirements. It also allows Virtual LANs to be implemented independent of the addressing scheme imposed by the network layer. A properly designed switch can mobilize a variety of mechanisms such as Virtual LANs, Switch RMON monitoring, network management, per-packet prioritization, and others, to efficiently handle broadcasts and accomplish the LAN segmentation that was traditionally the domain of routers.

Migration to ATM switching technology should be viewed as a strategic plan adopted by organizations to support future network applications. ATM is flexible and scalable; its transfer services support all types of traffic including data, voice and video; and it serves LAN as well as WAN protocols.

ATM differs from all other protocols on the market in that is was designed from the ground up, to cope with network application needs by using a "cell-forwarding via switching" methodology. In other words, both "switching" and "routing" functions are native to the ATM protocol. In ATM, routing tasks are performed only in the connection establishment phase. Once established, the packet stream is forwarded over the switched connection at a very high speed, and does not pass through routers. In a logical sense, routing is accomplished separately from actual data transfer. Physically, routers are kept away from the main stream of traffic, and conventional routing functions are relegated to the edge of the LAN where they support routing for the legacy installed base. ATM routing architecture is therefore very different from today's architectures, and ATM routers will have a much diminished role in tomorrow's networks.

Merging Switching and Layer-3 Functions[1]

In light of the rapid development of networked applications and the growing demand for higher speed and lower latency data communications, LAN switches have become popular because they deliver the required efficiency. Traditional routers have become less popular because they are no longer viewed as the appropriate solution for data- intensive applications. Router logic imposes physical topology constraints which do not give users the option to configure the network in the way that would achieve the best performance.

For example, consider a network with a server farm, where servers are consolidated in a central location and each one is given a dedicated router port. With this configuration, all client/server traffic travels via the router and incurring all the latency and overhead inherent in the routing process. In contrast, Layer-2 switches perform a very limited set of operations, deliver much faster performance, and are less expensive to design and implement than a router.

Solving network performance issues using Layer-2 switches is sufficient when traffic flows with no Network Layer restrictions. However, whenever data must travel from one Layer-3 group to another, the packets *must* go through a routing function rather than a switching function. This performance constraint can be overcome using two complementary approaches.

1. LAN switches can be used to expand the number of users per logical group to match the maximum limits allowed by the network layer addressing scheme. This creates a "flatter" network configuration which means that a larger percentage of the network traffic can be switched rather than routed.

2. The second approach is simply to reduce the processing overhead and increase the forwarding speed of those packets which still need to cross a router. Although we assume that networks will move progressively towards a flatter topology, routing will still be needed for sessions between Virtual LANs. Forwarding traffic between Virtual LANs is based on the upper layer protocol which automatically directs the traffic to the routing entity. If the routing entity has a virtual link to each subnet/Virtual LAN it will be able to forward data between LANs at a faster rate. The faster these links operate and the faster the data is internally processed and routed, the better the network will perform.

---

[1] Third layer support is required when applications are running over network layer protocol stack.

While this approach provides an adequate performance solution for conventional applications, it might not be sufficient for the multitude of very demanding network applications that are currently emerging.

The LANswitch Approach: The best of both worlds: fast forwarding and routing intelligence

The concept of Layer-3 switching calls for splitting Layer-3 routing functions into two distinct processes: Layer-3 Routing and Layer-3 Switching.

Layer-3 Routing

The first process is CPU based, computing intensive and time-consuming. It handles every aspect of the routing and address processing except the actual forwarding of the data. It's main output is a simple forwarding database which is used for the actual data transfer.

Layer-3 Switching

The second process performs the actual data transfer. It uses the forwarding table previously calculated by the first process. If this forwarding process is limited to one or two address schemes (IP/IPX), it can be implemented in hardware.

Since actual data flow is managed by the Layer-3 Switching process, the result is fast and inexpensive switching based on Layer-3 information. The Layer-3 Routing process is consulted only in the case of an unknown destination, or when the information the Layer-2 switch is insufficient.

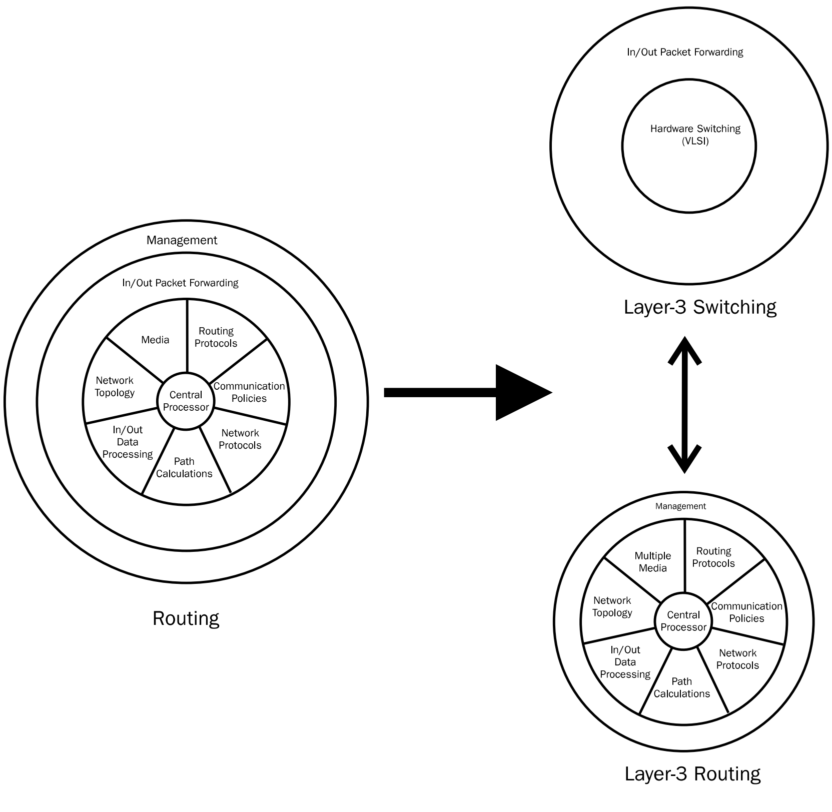

Figure 1 Splitting the Routing Function into Layer-3 Switching and Layer-3 Routing

Switch Architecture and Layer-3 Switching

One way to accomplish this complementary approach to Layer-3 switching is to adopt the new design scheme of splitting a router into two distinct functions: fast forwarding and routing intelligence. One half of the split router performs all of the traditional routing functions. The other half comprises a fast forwarding switching engine which performs the minimal set of routines required to accurately switch the data to its destination. When designed properly and integrated with the switching architecture, this set of routines can be implemented in hardware (VLSI) so that it transparently binds the source with the destination and minimizes the latency added by the routing function. The ability to support transparent switching depends on the specific design of the split-router as well as the architecture of the switch with which it interfaces.

The more fully the switch architecture supports a “flat” network topology, the larger the potential broadcast domain, and therefore, the more easily Layer-3 switching can be implemented.

Only a distributed switch architecture based on Layer-2 switching, combined with a very fast backplane operating independently from the switch, will be able to provide instantaneous communication between source and destination ports. Trying to integrate this design into an “unsuitable” switching core, will not deliver the promised speed.

Madge LANswitch Layer-3 Direction

Madge LANswitch architecture has every component required for high performance Layer-3 switching. Flatter network design coupled with efficient broadcast handling is the right approach for better network services. Layer 3 switching has a role in certain circumstances. It is useful for switched networks which require inter-Virtual LAN communication, and for those customers who choose to keep their large numbers of desktops unchanged and don’t want to modify their layer three structure today. As inter-subnet traffic levels grow so that they can no longer be efficiently handled by a router, Layer-3 switching should be considered. One of the notable advantages of integrated fast-forwarding (switching) based on Layer-3 protocols is the switching between Virtual LANs. Layer-3 switching provides a very efficient solution that includes those router functions that are still necessary for the network without investing the funds and paying performance penalties for unnecessary features.

Layer-3 switching is yet another step towards virtual networking. Second layer switching eliminates the physical constraints by permitting flexible, location independent network design. Third layer switching is the new approach to Layer-3 functionality that goes one step further. It frees the network from logical constraints imposed by network layer architecture.

The optimal Layer-3 switch architecture should eliminate all dependency on a central processing point (i.e., a central router). Preferably, the routing and switching functions should be completely integrated into a unified product for better performance and reliability. By definition, an integrated unit is not dependent on a central router, nor is it affected by the speed and status of the links connecting the router to the Layer-3 switch. A switch which incorporates both routing and Layer-3 switching is better able to maintain an accurate, real-time database. In the Madge Layer-3 architecture, both the switching and the routing functionality enjoy the full internal bus capacity. An external router is limited to the speed of the physical interface. In fact, if designed properly, the Layer-3 switch should be able to guarantee *one hop* between switches meaning *only one fast-forward layer-3 process is required when communication is either local or spans multiple hubs*.

A distributed switch architecture which integrates Layer-3 functionality guarantees redundancy in case of failure and allows for load sharing between multiple Layer-3 switches. The design should provide fault-tolerance options, and should preferably distribute routing functions amongst the switches. Fault-tolerance ensure no single-point of failure, while distributed routing allows for better performance by better utilizing available resources.

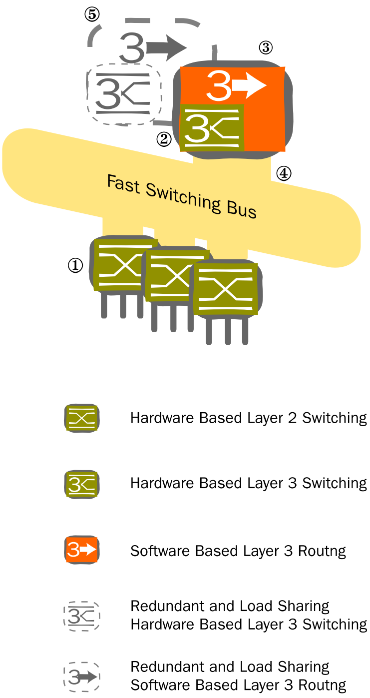

Figure 2 The Madge LANswitch Layer-2 and 3 Switching Integration

Figure 2 explains how the Madge LANswitch Layer-2 and 3 Switching Integration works:

1. Distributed Layer-2 switching is the Core of the Madge switching architecture. MAC layer switching will always be the first choice.

2. Layer-3 switching is the next level of the architecture. Packet switching based on Layer-3 addressing is achieved with a hardware based Layer-3 switch which complements the MAC layer switching when wire-speed communication between Virtual LANs or subnets is required.

3. Layer-3 routing is the highest architectural level. It supplements the Layer-3 switching as well as providing Layer-3 functionality.

4. A unique part of the architecture is the direct attachment of each and every component of the architecture directly to the core of the switch- the Fast Switching Bus. Access to the switching fabric is therefore real-time, unrestricted and highly reliable.

5. The switching architecture is fully distributed and redundant. Figure 2 shows the redundant and load-sharing aspect of the Layer-3 component. The architecture allows for multiple Layer-3 switches as well as Layer-3 routing with load sharing.

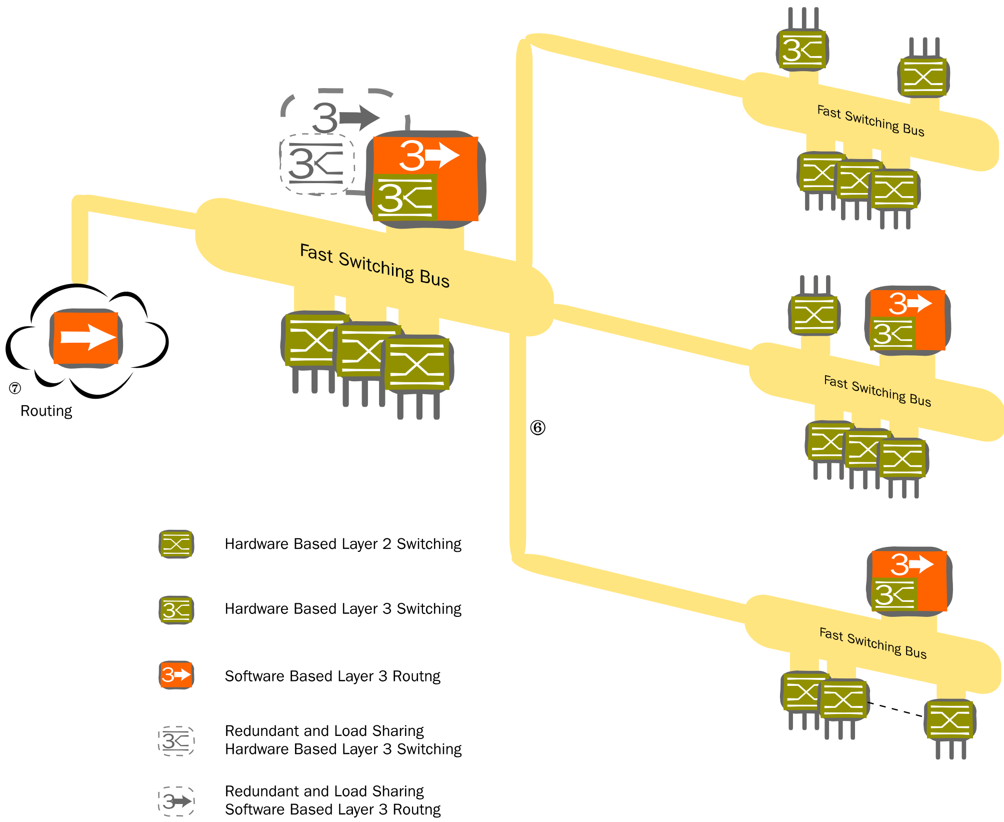

Figure 3 The Madge LANswitch Switching Architecture

Figure 3 shows the Madge LANswitch Layer-2 and 3 switching architecture in a wider context. (The numbers below refer to figure 3)

6. The Layer-3 intelligence guarantees one hop Layer-3 switching. The packet enters the Layer-3 logic only once, and is then forwarded directly to the end stations.

7. The design is fully compliant with all Layer-3 routing standards. Compatibility with standard routers is assured.

Summary

As networks continue their migration to LAN switching and ATM, traditional routers are losing their significance in data forwarding to switches, which provide a better solution for data and performance requirements.

Layer 3 switching is a step forward in the direction of switched virtual networks. It allows switches to better address both new emerging applications and existing networks. It expands the capabilities of hardware based switches to encompass functions that are currently handled by slower, software based routers.

While routing requires the division of a network into multiple logical subnets, a layer-3 switch based network can be much flatter and consequently faster. Logical groupings can be set up as needed based on purely organizational requirements, in the form of virtual LANs. A good design of switched network should minimize the time sensitive traffic that needs to be routed. When routing is needed, layer 3 functionality provides fast and virtually delay-free connectivity between Virtual LANs, keeping the separate broadcast domains

Switches place users who are connected to different physical segments on one virtual segment. Network design is no longer physically dependent. When high speed communication between different logical groups is required, Layer 3 switches go one step further. They give the freedom to place users who are defined on different logical subnets or virtual LANs on one virtual segment.

Madge Networks

Americas	**Asia, Australia & New Zealand**	**Europe, Middle East & Africa**	**Japan**
2310 North First Street	64-01 Central Plaza	Knaves Beech Business Park	Mita NN Building
San Jose	18 Harbour Road	Loudwater , High Wycombe	1-23, Shiba 4-chome
CA 95131/1011	Wanchai	Bucks HP10 9QZ	Minato-ku, Tokyo 108
United States	Hong Kong	England	Japan
Tel+1 408 955 0700	Tel+852 2593 9888	Tel+44 1628 858000	Tel+81 3 5232 3281
Fax+1 408 955 0970	Fax+852 2519 8022	Fax+44 1628 858011	Fax+81 3 5232 3208
http://www.madge.com	http://www.madge.com	http://www.madge.com	http://www.madge.com

Trademarks appearing in this document are the property of their respective owners