MadgeOne

Multilayer/Multiprotocol Switching over ATM

A TECHNOLOGY WHITE PAPER

AUTHOR: Martin Taylor, Madge Networks, April 1997

Executive Summary

This paper is intended for network designers and planners who are considering the deployment of ATM in the LAN to provide a scalable, fault-tolerant and cost-effective backbone with excellent investment protection. The paper describes how multilayer/multiprotocol switching over ATM meets the needs of the most demanding enterprise LAN environments, providing connectivity with existing desktop LAN technologies, supporting directly-connected end stations and flexibly accommodating existing Network Layer protocols and addressing schemes.

The paper focuses on the use of the emerging ATM Forum standard for MultiProtocol Over ATM to provide a distributed multilayer/multiprotocol switching solution, as a key element of the MadgeOne architecture. Comparisons are made with other techniques for supporting Network Layer protocols, including IP Switching from Ipsilon and Tag Switching from Cisco.

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.

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.

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

Prior to reading this white paper, the reader is strongly advised to refer to the Madge white paper “MadgeOne – Multilayer IP/IPX Switching in the LAN”, January 1997, to learn about the concepts of multilayer switching and the application of Virtual LANs. A knowledge of these terms and concepts is essential before tackling the subject of multilayer switching over ATM.

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.

Multilayer Switching and ATM

In the Madge white paper “MadgeOne – Multilayer IP/IPX Switching in the LAN”, January 1997, we describe a multilayer/multiprotocol switch as a device which combines the functions of LAN switching and multiprotocol routing in a single system. Such devices can provide very effective building blocks for large frame-based LAN infrastructures, connecting multiple LAN segments together and providing high-speed forwarding both within and between subnets.

Multilayer switching over ATM applies the same concept in a distributed fashion to an ATM network. Ethernet and Token Ring LAN segments can be connected to a backbone consisting of one or more ATM switches via multilayer/multiprotocol edge switches. End stations, which may be desktop PCs, workstations or servers, can be connected directly to the ATM switches. The ATM backbone will then enable communications to take place between and among any combination of end stations on Ethernet, Token Ring or ATM via direct switched connections, whether the respective end stations belong to the same IP subnet or not.

LAN Emulation Over ATM

ATM is now becoming widely adopted as a scalable, fault-tolerant, high performance solution for both LAN backbones and for demanding workgroup applications. For most enterprise LAN applications, the technique that is currently used to adapt the services provided by an ATM network to the needs of existing LAN applications is LAN Emulation over ATM (LANE).

We do not have space here to describe ATM or LAN Emulation in any detail. The reader who is interested in a more in-depth understanding of these subjects is invited to read the Madge white paper "LAN Emulation Over ATM", August 1996.

However, a few brief words of introduction on ATM might be helpful here. ATM is a switching technology for both LANs and WANs based on small, fixed-length packets known as cells. Unlike connectionless packet-based technologies such as Ethernet, ATM is connection-oriented, so ATM stations make use of a signaling protocol to request connections across the ATM network to the resources with which they wish to communicate.

LAN Emulation consists of software that resides in edge switches (which connect Ethernet and Token Ring LAN segments to ATM backbones) and in stations connected directly to ATM. This software is known as the "LAN Emulation Client" or LEC. The LEC software works together with additional software in the ATM network which provides some centralized services – the LAN Emulation Server, and the Broadcast & Unknown Server, collectively known as the LES/BUS. The purpose of all this software is to adapt the connection-oriented services provided by ATM to handle the transport of connectionless LAN traffic, including both Ethernet and Token Ring frames. The LEC works with the LES to resolve Ethernet and Token Ring MAC addresses to ATM addresses, so that the LEC can set up connections or "Virtual Circuits" to destinations within the ATM network. The BUS provides a forwarding service for broadcast and multicast packets, which is needed because ATM does not inherently provide the any-to-any multicasting that is needed to support LAN protocols such as ARP.

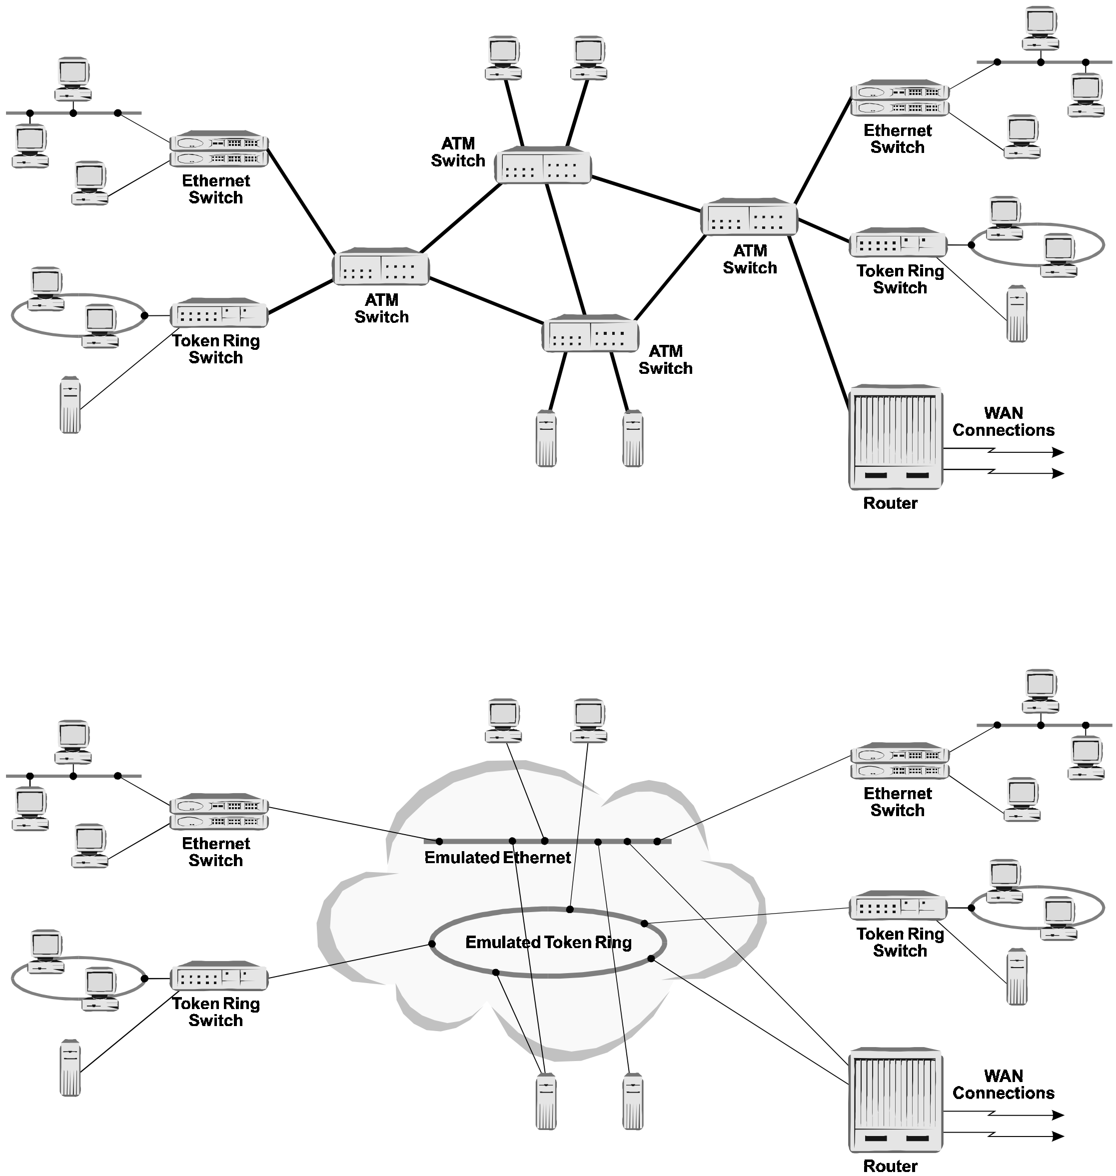

Figure 1: Physical and Logical Views of Emulated LANs

When ATM is used with LAN Emulation, we can create one or more virtual or emulated LANs (ELANs) within the ATM network that support the transport of Ethernet or Token Ring packets between and among edge switches and stations connected directly to ATM. In other words, the ATM network can emulate one or more LAN segments. When ATM is used with edge switches to connect physical Ethernet or Token Ring segments to an ATM backbone, we can link ELANs within the ATM backbone to VLANs within the switched LAN domain, so that ELANs and VLANs simply become extensions of one another.

LAN Emulation deals only with the Layer 2 functionality of a LAN. Where a LAN contains multiple subnets, LAN Emulation itself provides a solution only for communication within the subnets – and requires some kind of external Layer 3 forwarding function, or router, to provide communication between subnets.

We could choose to treat the ATM network with LANE simply as a group of one or more LAN segments, and use an external router as the device which performs Layer 3 forwarding between the emulated LAN segments. However, a new standard which is emerging from the ATM Forum goes one better than this: it enables the entire ATM network to be treated as a kind of distributed multilayer switch. The work on this standard is known as MPOA – MultiProtocol Over ATM.

MultiProtocol Over ATM

MPOA is designed to build upon and enhance LAN Emulation. Simply put, the objective of MPOA is to enable any combination of edge switches and ATM-connected end stations to intercommunicate directly using switched virtual circuits, without passing through any separate Layer 3 forwarding function in the ATM network, regardless of whether the stations that are intercommunicating belong to the same subnet or not. At the time of writing, the MPOA specification was out for Straw Ballot on its way to final approval by the ATM Forum.

LAN Emulation enables any combination of edge switches and ATM-connected end stations to intercommunicate directly using switched virtual circuits, provided that the intercommunicating end stations belong to the same subnet. If they belong to different subnets, then the stations would normally expect to send these packets via one or more routers.

Let us imagine a pair of edge switches which are serving Ethernet-connected end stations that belong to different subnets. Both edge switches are connected to an ATM LAN backbone. To communicate across the backbone (to a different subnet), an Ethernet station will look for the "default gateway" or router that is to act as the Layer 3 forwarding function. MPOA defines a MultiProtocol Server (MPS) function in the ATM network, which behaves just like a router and which provides a Layer 3 forwarding function. The Ethernet station is configured with the address of the MPS as its default gateway. It therefore sends any packets destined for a station on a different subnet to the MPS, via the LANE function in the edge switch. The MPS will, in turn, forward these Ethernet packets using LANE to the edge switch which serves the destination station.

The MPS provides a forwarding function between subnets by "pretending" to be a router. But in the example we have described, the packets passing between the two edge switches are not traveling on a direct switched virtual circuit – they are having to go via the Layer 3 forwarding function in the MPS which must re-assemble cells into complete packets before sending them on. Not only does this slow the traffic down, but the MPS is likely to become a bottleneck if traffic flows between subnets are heavy.

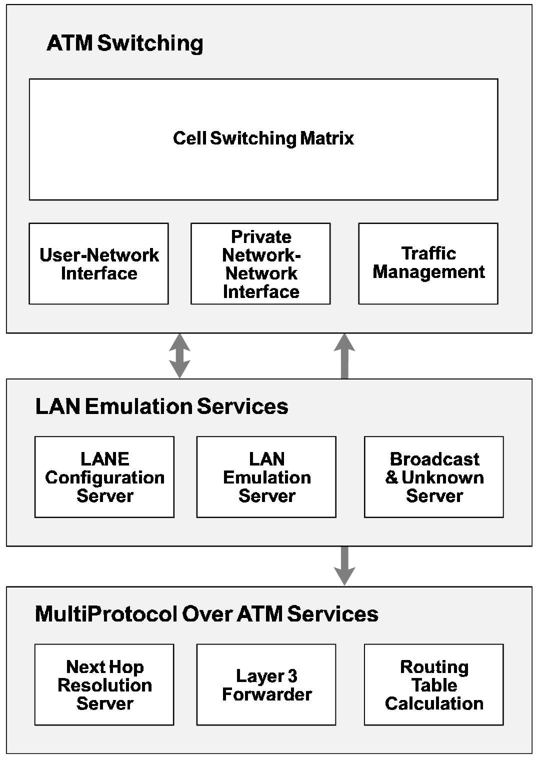

Figure 2: ATM Switch incorporating LANE and MPOA Services

The solution to this is provided by cooperation between the MPS and the MultiProtocol Client (MPC) software in edge switches and ATM-connected end stations – which is based on the LAN Emulation Client (LEC) software, but with some new functions. The MPC software is intelligent enough to recognize when it is sending traffic to a Layer 3 forwarding function (the MPS) rather than on a direct switched virtual circuit to the final ATM destination. When the MPC sees this happening, it issues a request to the MPS to report back to it the ATM address of the final ATM destination. It does this using a protocol called Next Hop Resolution Protocol (NHRP), which talks to a function in the MPS called the Next Hop Server (NHS). The NHS responds to this request with the appropriate ATM address, which the MPC then uses to set up a direct switched virtual circuit to the final destination within the ATM network, whether it be an edge switch or an ATM-connected end station.

With the direct switched virtual circuit in place, the MPC stops sending packets to the MPS for forwarding, and sends them instead on the direct connection to their final destination within the ATM network. Packets will now travel faster to their destinations and will avoid having to pass through any Layer 3 forwarding functions within the ATM network that could become bottlenecks.

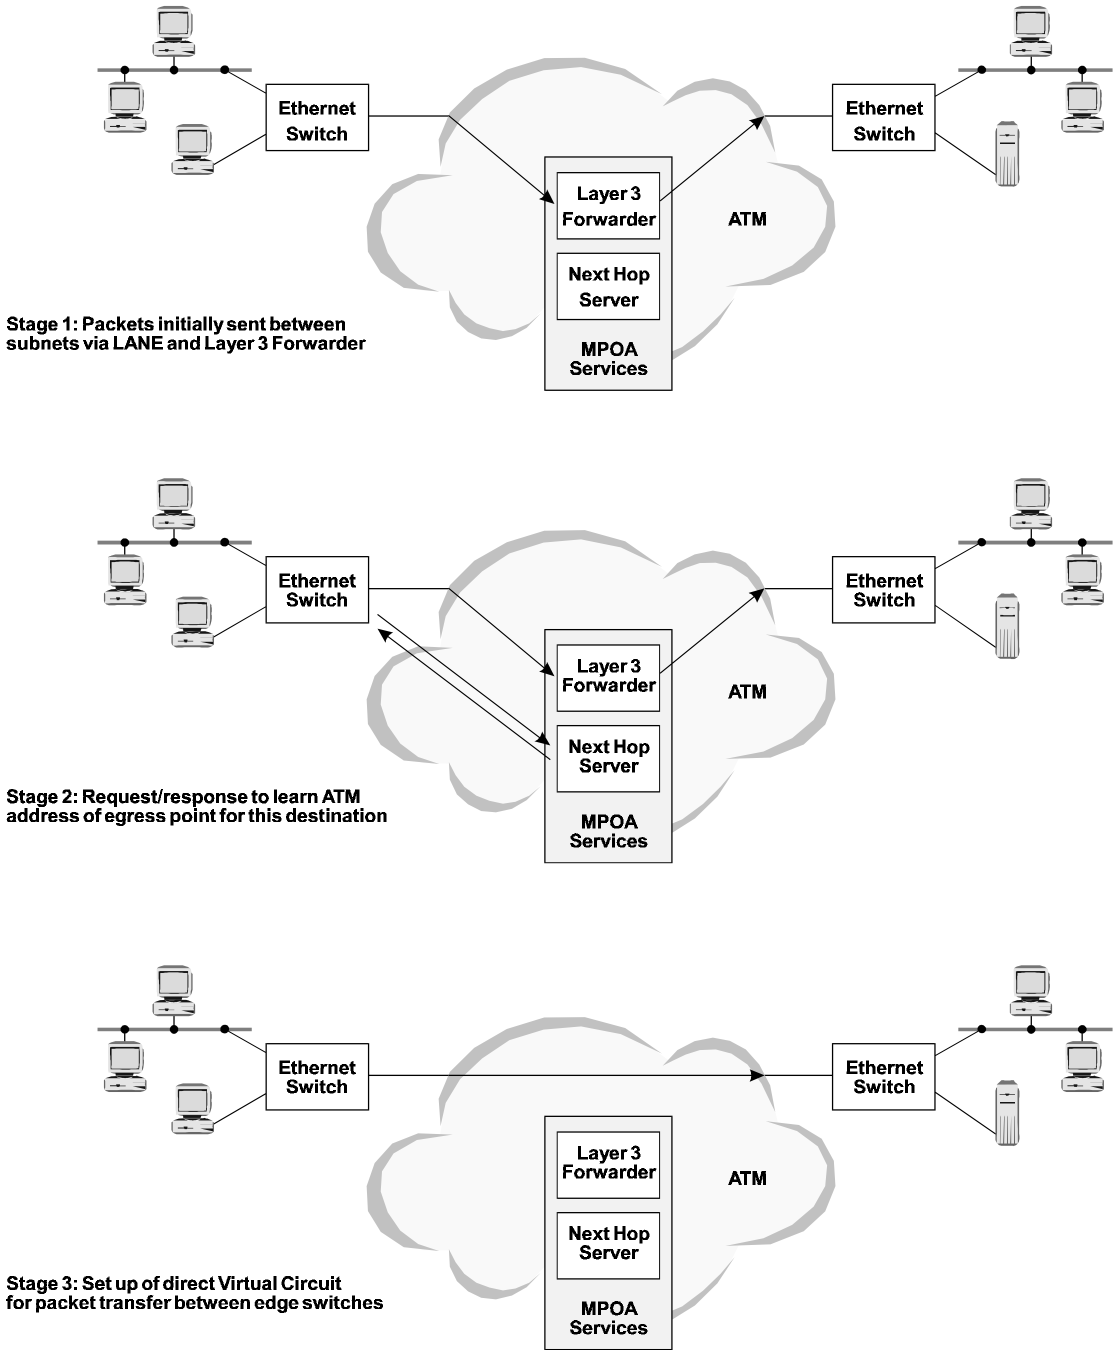

Figure 3: MPOA supporting communications between stations on different subnets

Note that, once a direct switched virtual circuit has been set up between these two edge switches, other stations connected to these switches are able to share it. So if one edge switch supports a server farm and the other edge switch supports a group of desktop PCs, then all communications between this group of desktop PCs and the server farm can travel via the same direct switched virtual circuit.

Distributed Multilayer Switching with MPOA

An ATM network with MPOA software in edge switches and directly-connected end stations behaves effectively like a distributed multilayer switch. Stations connected to this network can communicate with one another based on either Layer 2 or Layer 3 switching, depending on whether they are members of the same subnet or not. This means that we have the potential flexibility to move stations around from one physical switch port to any other, without having to change the station IP addresses. We also have the ability to define security policies, by placing secured resources on different subnets from users, and by placing the subnets in different emulated LANs. This ensures that communications between unsecured and secured parts of the network must pass through the routing functionality of the MPS, where IP-level connection privileges and access lists may be defined by the network administrator.

Just as in discrete multilayer switching devices, there is a need for distributed multilayer switching to participate in routing protocol exchanges with conventional routers. This is handled by the MPS function, which may support RIP or OSPF or both. Each MPS function within the ATM network looks like a peer router to any conventional routers connected to the network, so that distributed multilayer switching can happily coexist with routers in practical networks – even to the extent that conventional routers could be used to provide an alternative path between LAN segments that are connected by means of distributed multilayer switching.

Distributed multilayer switching is able to provide an excellent solution for switched connectivity between Ethernet and Token Ring. The one key requirement here is that Ethernet and Token Ring stations are defined as belonging to different subnets, since we can't have Ethernet and Token Ring traffic within the same emulated LAN. We may also have to make sure that Token Ring stations are limited to a maximum packet size of 1500 bytes, since some MPOA implementations may not support packet fragmentation for communication between subnet technologies with different maximum packet sizes.

MPOA and Multiprotocol Switching

The work on the MPOA standard being conducted by the ATM Forum is intended to address the needs of any Layer 3 networking protocol. The only such protocol which is both widely used and fully defined in the public domain is the Internet Protocol (IP), so the MPOA standard is likely to be focused on IP, together with some guidelines on how other routable protocols can be handled.

This clearly raises questions about the applicability of multilayer switching based on MPOA to real enterprise LANs, where there are likely to be routable protocols other than IP in use – particularly IPX, but perhaps also DECnet, AppleTalk or even OSI.

For protocols that are very widely used, particularly IPX, we can expect to see de facto standard implementations, based on informational appendices to the main body of the MPOA standard. Leading vendors in the ATM industry are likely to organize interoperability testing events for MPOA, where both IP and IPX interoperability will be established.

There are two solutions for dealing with routable protocols that are not supported at Layer 3 by MPOA. The first solution, which addresses also the needs of non-routable protocols such as NetBIOS, LAT and DLC, is to place all resources that need to communicate using protocols other than IP into a single, large emulated LAN. This provides Layer 2 connectivity for all protocols – in effect a bridged environment. MPOA provides Layer 3 connectivity between subnets for IP, while all other protocols must be defined around a single subnet across the entire LAN.

A single large emulated LAN may well raise concerns about excessive broadcast traffic; but these concerns can be addressed with the aid of intelligent broadcast filtering within the BUS function of LAN Emulation, such as that provided by Madge's Active Broadcast Control technology. The requirement for security of IP-based communications can be met by placing any sensitive IP resources into a separate ELAN, forcing all traffic to and from these resources to pass through the Layer 3 switching function of MPOA, where access control can be applied.

The alternative solution is to use conventional routers connected to the ATM network to provide Layer 3 forwarding between multiple ELANs for all routable protocols other than IP – plus, if necessary, bridging between ELANs for non-routable protocols. This solution may be preferable in environments where security is a particularly important consideration.

Alternative Architectures for ATM LANs

Standardized techniques that provide scalable solutions for IP networking over ATM have been a long time coming. The standard for LAN Emulation Over ATM was approved by the ATM Forum in May 1995, but LANE alone does not deal with the need for communication between IP subnets over ATM, and requires external routers to do this. The goals agreed by the ATM Forum for MPOA were ambitious, and the solution is complex, so it has taken a long time to arrive at something workable.

In the meantime, it is no surprise that proprietary solutions for handling IP over ATM have come onto the market. There is a clear need here which, arguably, was not being met in a timely manner by the standards process. The best-known of these proprietary solutions is that proposed by Ipsilon, Inc. under the name "IP Switching".

More recently, Cisco have put forward another approach known as "Tag Switching". Cisco has submitted proposals to the Internet Engineering Task Force (IETF) for a set of standards around this approach, and other vendors have been invited to contribute to this process. As a result, IBM has come forward with a technology know as Aggregate Route-based IP Switching (ARIS) which has many similarities to Tag Switching. Other vendors have also expressed an interest in this topic, under the general heading of "Multiprotocol Label Switching".

Since both IP Switching and Tag Switching have the potential to provide a solution for IP over ATM in the enterprise LAN, it is worth taking a moment to examine these approaches and to discuss how effectively they meet the needs of corporate and enterprise LANs.

IP Switching

The Ipsilon approach uses special edge routers combined with IP routing software running on each ATM switch, and effectively creates a kind of distributed virtual router from these elements. Conventional ATM signaling and routing protocols are completely discarded, and are replaced by a combination of IP routing protocols and a new and proprietary protocol: the Ipsilon Flow Management Protocol (IFMP).

The edge devices behave like routers in the sense that they map IP datagrams into ATM cell streams. However, there is no need for the edge routers to participate in IP routing protocols since they just forward all traffic to the ATM switch to which they are connected. Each ATM switch has a full router software load, and communicates with other ATM switches and routers external to the ATM network using routing protocols such as RIP and OSPF.

When a station connected to an Ethernet segment starts sending IP datagrams to an edge device for forwarding via the ATM network, the edge device sends these packets to the router software in its neighboring ATM switch via a “well known” virtual circuit number. The router software in the switch reassembles the IP datagrams from the incoming cell stream and uses its routing table to decide which port to forward them on. The next hop may be another ATM switch or it may be an edge device – whichever it is, the datagrams are sent on their way on another well-known VC. Thus, initially, the IP datagrams proceed hop-by-hop through the ATM network, being reassembled from cells into packets in the router software at each switch, and then being segmented into cells again for forwarding to the next hop. Eventually when they reach the edge device which serves the destination subnet, they will be transmitted as frames on the appropriate LAN segment to the destination station.

The router software in each ATM switch is able to detect when a number of packets are received in quick succession from a particular source which are addressed to a particular destination. This is referred to as a “long-lived flow”. When the router software detects such a long-lived flow, it sends a message (using IFMP) to its upstream neighbor – which may be the edge device at which the flow is entering the ATM network, or it may be the router software in an upstream ATM switch – and requests the upstream neighbor to send IP datagrams on a different virtual circuit from the well-known virtual circuit. In fact it assigns a VC number for this flow on the link to the upstream device and tells the upstream neighbor what VC number to use. The flow then becomes a “labeled flow”.

Each ATM switch independently carries out this flow detection and asks its upstream neighbor to label the flow by giving it a specific VC number. The downstream edge device, where the flow leaves the ATM network and goes onto a LAN segment, also does this flow detection and requests its neighboring upstream ATM switch to label the flow.

Now, whenever the routing software in any given ATM switch sees a labeled flow coming into it for forwarding as an outgoing labeled flow, it instructs the ATM switch hardware to stop sending the incoming flow to the routing software, but instead to forward the cells that make up the flow directly on the virtual circuit number which represents the outgoing flow label. In effect, the routing software disconnects itself from the forwarding path and allows the ATM switch hardware to take over this task.

So, once the routing software in each ATM switch in the path has detected the existence of a flow, requested its upstream neighbor to label the flow, noticed that it is forwarding an incoming labeled flow to an outgoing labeled flow, and then instructed the ATM switch hardware to “join” the two labeled flows together across the switch – we have achieved an end-to-end switch path across the ATM network for this particular flow.

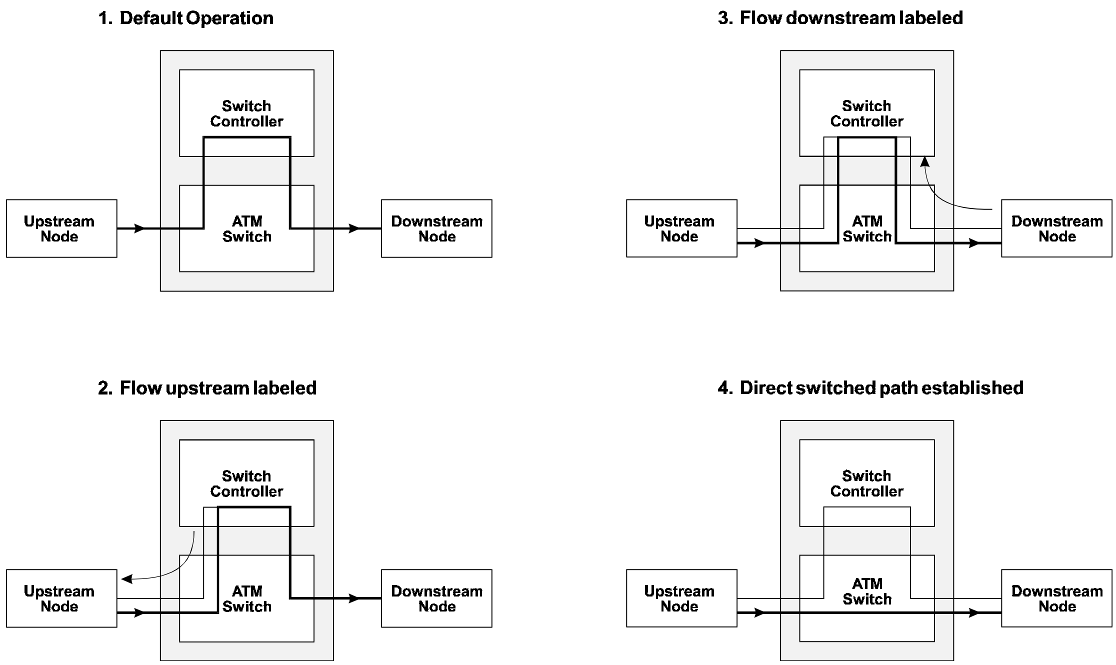

Figure 4: Operation of Ipsilon IP Switching and Flow Management Protocol

The example we have used to describe the operation of Ipsilon's IP Switching was based on traffic flowing across an ATM network between two edge devices, which connect end stations via LAN segments such as Ethernet or Fast Ethernet. However, IP Switching can also be supported at end stations connected directly to ATM switches. In this case, the end station must be able to send IP datagrams mapped into ATM cells according to Ipsilon's specifications, and support Ipsilon's IFMP protocol for labeling flows.

The advantages of Ipsilon's IP Switching approach over MPOA appear to be:

- IP Switching can provide a solution today for networks that require a very high capacity IP backbone, whereas solutions that implement the MPOA standard won't be available until late 1997.

- IP Switching is based on familiar IP routing protocols and therefore the learning curve for users can be expected to be less steep than for MPOA, which is based on ATM routing and signaling protocols.

- IP Switching involves less overall software complexity than MPOA, which may mean that practical implementations are more reliable than MPOA – at least until vendors have debugged their MPOA software.

Ipsilon's IP Switching can be thought of as a "quick fix" for supporting IP over an ATM backbone. It may be appealing to those who are in a serious hurry to solve their IP congestion problems, but for those who are prepared to wait, MPOA has a number of important advantages over IP Switching:

- MPOA offers a solution for all protocols, not just IP, through its ability to provide multilayer switching. Non-routable protocols, and routable protocols that are not supported at Layer 3 by MPOA, are handled via LAN Emulation. IP Switching is a pure

Layer 3 solution for IP only, and requires all other protocols to be encapsulated in IP for transport over the ATM backbone.

- MPOA enables IP subnets to extend across the ATM backbone, because it provides multilayer switching. This enables network administrators to handle moves and changes without assigning a new IP address to each user that moves to a different LAN segment, and it also enables users to "roam" in the LAN with the same IP address. IP Switching, as a pure Layer 3 solution, does not permit IP subnets to extend across the ATM backbone, and therefore suffers from the same issues around moves and changes as a classical router-based backbone.

- MPOA makes use of ATM routing protocols (PNNI) which are based on OSPF, but which are enhanced to take into account any reservations of bandwidth that may have been made on each link by applications that use ATM's Quality of Service, typically for real-time voice or video communications. IP Switching uses conventional IP routing protocols which are not "QoS aware". This means that data flows could end up being routed over connections between ATM switches that are already heavily committed to handling real-time traffic, rather than via alternative routes on which there is far more available bandwidth.

- MPOA is far more efficient than IP Switching in its use of ATM network resources, in terms of signaling, processing within the network, and usage of virtual circuits. IP Switching detects flows and establishes the switched path across the ATM network on a hop-by-hop basis, whereas MPOA detects flows once, at the entry point to the ATM network, and establishes the switched path end-to-end with one single signaling request. This means that switched connections are established faster, so users see quicker network response times. Also, IP Switching places each and every flow on its own, separate virtual circuit, whereas MPOA allows many flows to share the same virtual circuit across the network – and this means faster recovery time in the event of a switch or link failure that causes re-routing to take place.

MPOA has taken longer to develop than IP Switching because it has far more ambitious goals. Those goals are concerned with meeting the real needs of large, multiprotocol enterprise LAN backbones, adding greater flexibility and addressing future requirements for the integration of real-time traffic alongside data. IP Switching will, perhaps, meet some of the short-term needs of networks which have a very high predominance of IP traffic. But it seems clear that IP Switching will be eclipsed by MPOA once standards-compliant MPOA solutions start shipping in 1997.

Tag Switching

The impetus behind Cisco's announcement of Tag Switching was the need to provide a highly scalable solution for the Internet that makes best use of fast and cost-effective Layer 2 switching technologies at the core of the network.

The principles of Tag Switching can be applied to any Layer 2 switching technology, including ATM, Frame Relay and even, perhaps, Ethernet switching. We will focus here only on how Tag Switching operates in an ATM environment.

An ATM-based Tag Switching network comprises a mesh of ATM switches which are running Cisco's Internetwork Operating System (IOS) software, surrounded by routers also running IOS. Tag Switching requires an enhanced version of IOS which supports the Tag Distribution Protocol (TDP).

The edge routers and the ATM switches make use of conventional routing protocols to exchange reachability information, so that both the edge routers and ATM switches can build routing tables which associate IP address prefixes with router or switch ports. But in addition, the software running in the ATM switches and the edge routers assigns virtual circuit numbers on each incoming port for each set of IP address prefixes that maps to each of its outgoing ports, and communicates this information via the Tag Distribution Protocol to the upstream neighbor on the relevant incoming port. Each edge router, receiving this information from the network, enters the virtual circuit numbers as tag identities against the appropriate IP address prefixes in its routing tables. Each ATM switch correlates the virtual circuit numbers and IP address prefixes in its routing tables, and creates a mapping of virtual circuit number on each incoming port to outgoing port and outgoing virtual circuit number, thereby effectively programming the ATM switch fabric with pre-established paths.

The net result of all this is that the network effectively configures itself to create a full mesh of switched connections between all the edge routers. Furthermore, each edge router learns a set of mappings between IP address prefixes and virtual circuit numbers.

Now when an edge router receives IP datagrams to be forwarded across the ATM network, it can map IP destination address prefixes directly to an ATM virtual circuit numbers which represent pre-established paths across the ATM network to the appropriate exit points. Each IP datagram can then be converted to ATM cells that will traverse the appropriate path across the network, with switching taking place end-to-end, until the cells arrive at the appropriate edge router to exit the ATM network. Here, they are re-assembled into IP datagrams for onward routing to their final destination.

For very large IP networks like the Internet, Tag Switching has some very substantial advantages over Ipsilon's IP Switching, namely:

- With Tag Switching, every IP datagram, whether it is a "one-off", a short-lived flow or a long-lived flow, is switched as a cell stream end-to-end without the need for intermediate re-assembly and onward routing. IP Switching makes use of hop-by-hop forwarding for "one-off" datagrams and short-lived flows, which is much slower and less efficient.

- Tag Switching can carry many different flows, representing sessions between different source-destination address pairings, over one virtual circuit. IP Switching requires each long-lived flow to be assigned its own, unique virtual circuit. If IP Switching is applied in the Internet, the very large numbers of separate flows that will be seen will result in an explosion of virtual circuit numbers, which most real ATM switches won't be able to handle for lack of hardware or memory resources.

- Tag Switching uses conventional routing protocols, combined with the Tag Distribution Protocol, to pre-establish paths across the network based on network topology. These paths only change if the topology changes. IP Switching, on the other hand, establishes a switched path across the network dynamically, for each long-lived flow detected. This means that the amount of signaling activity seen with IP Switching is many orders of magnitude higher than for Tag Switching, and this signaling activity is likely to consume an unacceptable amount of switch bandwidth and processing power in a very large network.

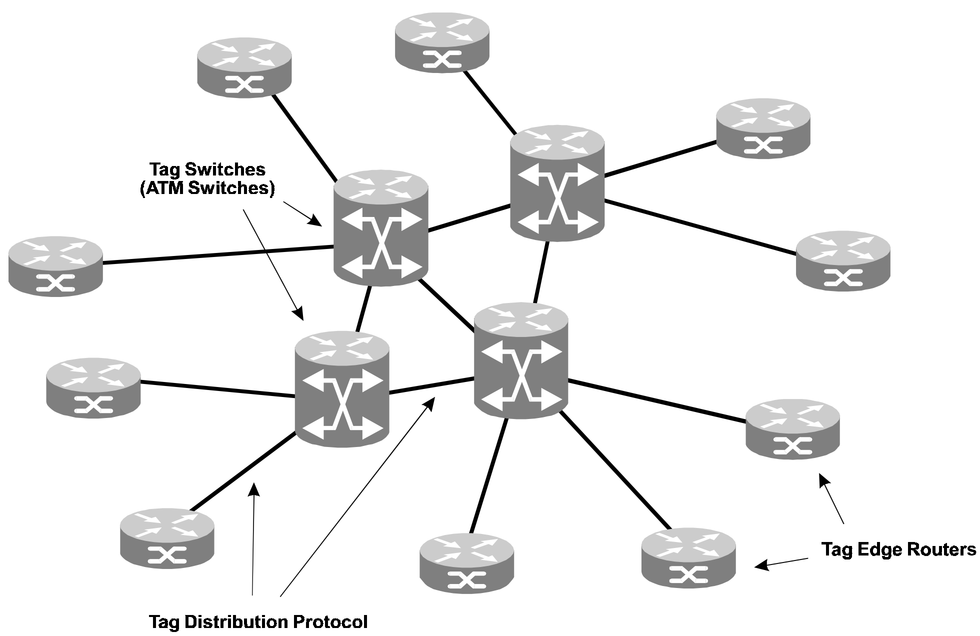

Figure 5: Elements of Tag Switching

The superiority of Tag Switching over IP Switching for very large networks such as the Internet is clear. IP Switching was, perhaps, not intended by its designers to address the needs of the Internet. In fact, MPOA was not intended by the ATM Forum to address such large networks, but MPOA is far more scalable than IP Switching, for reasons we have already discussed.

What is perhaps more relevant here is the comparison of Tag Switching with MPOA for corporate and enterprise LAN applications. Since Cisco has stated publicly that it will support both Tag Switching and MPOA, because they are complementary, it is not clear that Cisco really favors the application of Tag Switching in corporate LANs. MPOA appears to be superior to Tag Switching for this type of application for the following reasons:

- Tag Switching has many of the same problems as IP Switching in an enterprise LAN environment, which we have already discussed. It does not provide a solution for protocols other than IP, it does not handle IP subnets extending across the ATM network, and it does not use QoS-aware routing protocols. So Tag Switching doesn't match MPOA for flexibility, for multiprotocol support or for ability to integrate real-time traffic.

- Tag Switching uses full-featured routers as the edge devices to connect LAN segments to the ATM network. The high costs of these devices will compare unfavorably with the simpler MPOA edge switches, and will result in a much more expensive solution overall.

- It is not clear how Tag Switching can support end stations such as servers or high power clients connected directly to the ATM network. If end stations need to emulate edge routers to participate in Tag Switching, then it may not be practical to connect server farms or workgroups directly to ATM because of the complexity of the software required. MPOA is designed to support any mix of ATM-connected end stations and edge devices, and most ATM adapter card vendors are expected to provide MPOA client software drivers for their cards.

The case for Tag Switching in the Internet is very clear. The key advantage of Tag Switching – its immense scalability – is a big plus here, while the lack of multiprotocol support and the problems of direct ATM end station connectivity are not really disadvantages in the context of the Internet. But the requirements of the corporate or enterprise LAN environment demand a solution with somewhat different characteristics – which is where MPOA comes in.

Conclusion

In this paper we have described multilayer/multiprotocol switching over ATM 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 technique for achieving multilayer/multiprotocol switching over ATM will be fully specified in the ATM Forum MultiProtocol Over ATM standard. This means that multilayer/multiprotocol switching over ATM can be supported in multivendor networks. Furthermore, the MPOA standard is 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.

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/multiprotocol switching over ATM will bring much-needed relief by pushing back the frontiers of performance, efficiency and flexibility in the LAN.

