THE
TOLLY
GROUP

No. 8269 February 1998

TOLLY TESTED

Solutions
White Paper

IBM’s Multiprotocol Switched Services Server

Building Reliable & Scalable ATM Networks

A study commissioned by IBM

Table of Contents

The Tolly Group provides strategic consulting, independent testing, and industry analysis. It offers a full range of services designed to furnish both vendor and end-user communities with authoritative, unbiased information. *Fortune* 1,000 companies look to The Tolly Group for vendor-independent assessments of critical corporate technologies. Leading manufacturers of computer and communications products engage The Tolly Group to test both pre-production and production equipment.

The Tolly Group is recognized worldwide for its expertise in assessing leading-edge technologies. By combining engineering-caliber test methodologies with informed interpretation, The Tolly Group consistently delivers meaningful analyses of technology solutions. The Tolly Group has published more than 400 product evaluations, network design features and columns in the industry's most prestigious publications.

Kevin Tolly is President and CEO of The Tolly Group. He is a leading industry analyst and is responsible for guiding the technology decisions of major vendor and end-user organizations. In his consulting work, Tolly has designed enterprise-wide networks for government agencies, banks, retailers, and manufacturers.

For more information on The Tolly Group's services, visit our World Wide Web site at ***http://www.tolly.com,*** email to info@tolly.com, call 800-933-1699 or 732-528-3300, or fax 732-528-1888.

Internetworking technology is an area of rapid growth and constant change. The Tolly Group conducts engineering-caliber testing in an effort to provide the internetworking industry with valuable information on current products and technology. While great care is taken to assure utmost accuracy, mistakes can occur. In no event shall The Tolly Group be liable for damages of any kind including direct, indirect, special, incidental, and consequential damages that may result from the use of information contained in this document. doc 8269 ver. 18Feb 98

Executive Summary

The world's most advanced campus networks are built using ATM. This fact is universally acknowledged. That ATM is based on a massive tome of multivendor standards and enjoys a high degree of cross-vendor interoperability is also well understood. From this, one might conclude that all ATM networks, regardless of equipment vendor, are the same and offer equal levels of support for building reliable and scalable networks. Such a conclusion would be incorrect.

Like the other industry standard specifications for token ring and Ethernet, the ATM Forum specifications define a required, minimum set of functionality. As with legacy LANs, implementations of the standard can and do vary in the ATM world. Value-added services are already becoming key differentiators in the ATM marketplace.

IBM commissioned The Tolly Group to evaluate how the IBM Multiprotocol Switched Services (MSS) device enhances the robustness and scalability of ATM-legacy LAN networks. The Tolly Group built an actual network consisting of ATM, Ethernet and token ring technology that used ATM's LAN Emulation (LANE) as the basis for logical connectivity. From this starting point, The Tolly Group developed several scenarios to demonstrate how the advanced functions provided by MSS enhance robustness and scalability in areas such as redundant services, broadcast management and route switching. Testing was conducted in December 1997.

Overview

When ATM is used as the core of a campus backbone it must, by definition, integrate with existing frame-based Ethernet, token ring and FDDI LANs. While ATM links can simply be used as high-speed transports between frame-based LANs, most network managers desire a higher degree of integration. They require that devices connected to ATM switches be able to communicate not only with other ATM-connected devices, but also with stations connected to legacy LANs. ATM Forum LAN Emulation defines a set of functions and services that provides transparent integration of ATM networks and legacy LANs. Using LANE, ATM-devices appear logically as an extension of the existing campus network. Groups of ATM LANE devices form an Emulated LANE (ELAN) that is linked both to other ELANs and real LANs via standard layer 2 switching and/or layer 3 routing services.

Designers of ATM LANE-based campus networks discover that because LANE is deliberately modeled after real, frame-based networks, network personnel may encounter some of the same scalability and robustness impediments found in traditional frame-based LANs. The nature of ATM, though, allows for innovative solutions not possible in the realm of legacy LANs.

In the course of testing, The Tolly Group established that IBM's MSS is able to enhance networks in several critical areas. Though it is instructive to review the details of each area contained in this report, the findings can be summarized as follows:

Redundant Services

A hot-standby MSS can be used to provide automatic redundancy of all MSS services including LAN services, IP gateway and layer 2/3 switching and routing. The secondary MSS, which can be located on the same switch or on a different switch, automatically detects the failure of the primary unit and, within seconds, assumes its identity and enables an identical set of services. Depending on the functions being used, many end-stations will not even be aware of the changeover.

Broadcast Management

Intelligence in MSS monitors higher-layer traffic and develops an awareness of the protocols operating at each station. For example, by observing that RIP frames are emanating from a particular station, MSS can infer that the station connected to that port is a router. MSS intercepts broadcasts, analyzes them and, using the information it has gathered about the stations connected to the ELAN, MSS directs the broadcast only to those stations for which it is meaningful. Not only does this save bandwidth, but it also avoids unnecessary end-station processing.

Inter-ELAN Connectivity

Routers often are relegated as the workhorses of campus networks, forging links with disparate workgroups. Often, though, these links become overtaxed. In a robust MSS environment, two new services — Route Switching and Super ELANs — provide the versatility to bypass congested router links and reduce latency for traffic passing between ELANs.

Route Switching

Using connection-oriented switched virtual circuit (SVC) services inherent to ATM, MSS allows individual stations (outfitted with appropriate network drivers) to communicate directly in a “zero-hop” fashion to partner stations on different ELANs. After the initial exchange (between stations) flows across a layer 2/3 bridge/router, an SVC is dynamically built allowing all further communications to bypass the bridge/router. Not only is station-to-station performance improved by avoiding the potential bridge/router bottleneck, but removing the hop reduces session latency. Finally, this offload reduces the processing burden on the bridge/router.

Super ELANs

Known as “Super ELANs,” MSS provides network managers with a simple, yet powerful way to logically link ELANs without requiring traffic to transit a layer 2/3 bridge/router. A Super ELAN is a customer-defined superset of ELANs that can be managed in real time. MSS uses SVCs that it sets up between ELANs to provide high-performance, zero-hop connectivity among all devices of the Super ELAN. Furthermore, this is done without any modification to client-station drivers.

Redundant Services

With the amount of mission critical data riding over enterprise ATM networks today, it's imperative that strategic solutions provide high availability. Today, that's often accomplished by offering both physical redundancy, and redundant ATM services. This redundancy, coupled with an automatic recovery scheme, allows the network's users to continue operating even if critical devices fail, or if the link between key devices fails.

Physical redundancy must include replication of components, as well as duplicate connections between devices. This physical redundancy is a prerequisite for a highly-available enterprise network, but doesn't address all of the requirements for high availability.

The other element needed to ensure high availability is redundancy of logical network elements. Logical redundancy ensures that a physical problem with a network device won't require a reconfiguration of either end stations or edge devices. When a physical problem occurs, the multi-layer logical identity of the failed device is automatically and dynamically re-established using a different physical device with only a brief disruption of service. That means any such redundancy needs to be completely automatic, relatively fast, and should extend beyond basic connectivity to fully support layer 2 and layer 3 services.

For example: IP devices are configured to use a single primary gateway to the rest of the IP world. If that gateway device goes away, connectivity is lost. So the capability of assuming the complete identity of the failed device (at the ATM level, the layer 2 MAC level, layer 3 IP subnet level, layer 3 IPX subnet level, etc…) is key.

For layer 2 redundancy within ATM, the need for logical redundancy is focused on LANE Services such as the LAN Emulation Server (LES), the LAN Emulation Configuration Server (LECS), and the Broadcast Unknown Server (BUS). For stations running LAN Emulation to initiate sessions to other LANE stations, those services must be available. Furthermore, if the LANE implementation offers more advanced functionality like ELAN-to-ELAN bridging, LANE services must be available for any communication (not just session initiation) to take place.

Before IBM MSS

While redundant solutions for multiple services are required in an enterprise network, a discussion of LAN Emulation provides ample insight into the issues.

End-stations and edge-devices participating in LAN Emulation establish contact with a particular set of LANE services (there could be several available in the network). The LANE clients are configured with the ATM address of the LANE services they will use. These LAN stations will only attempt to communicate with the LANE services address set up at the time of configuration.

Should the LANE services become unavailable, there are only two options: neither of which is automatic, fast or cost effective.

In order to set up redundant services on an ATM network, historically it was often necessary to deploy a backup ATM switch on standby in case the primary switch crashed. Since LANE services weren't physically separate from the ATM switch, the secondary switch used the same network hardware configuration just to backup the LANE services. Unfortunately, this was an expensive solution. The same approach was true for routers -- a fully-outfitted backup router sat as a watchdog, ready to step into the breech when a primary unit failed.

Another, although less popular, alternative was to pre-configure a secondary backup system in parallel to the primary system. If a failure occurred, the network manager had to physically reconfigure the ATM network's LANE clients. This required the network manager to reconfigure each of the LANE clients with the new LANE services information and possibly with layer 3 information as well. This resulted in a temporary outage while the clients were being reconfigured and often would take significant time depending on how many LANE clients had to be changed.

With IBM MSS

IBM has taken the idea of services redundancy to the next level with MSS. Not only does MSS provide a solution for layer 2 and layer 3 redundancy, it can also provide a redundant backup for all of the services that the MSS supports. Network managers are no longer faced with the less than ideal options described above. Instead, IBM has architected an automatic recovery feature into MSS.

Multiple MSS devices can be connected to the ATM network. While one MSS actively provides the full complement of LANE services and routing functions, a second MSS sits by monitoring the primary MSS as a hot-standby in the event of a failure. Should the primary MSS fail, the secondary MSS will detect a loss of contact with the primary MSS and automatically take over the identity of the primary MSS and enable LANE services and routing. The nice feature is that network managers don’t have to do anything – the stand-by MSS re-establishes connectivity with LAN emulation clients without any changes to client configurations or intervention by network managers. And new LANE clients register with the same server resource. Note: Since the watchdog MSS needs to have the identical configuration as the “watched” MSS, there is a one-to-one relationship between active and standby MSS systems. The backup MSS can also be connected to the same switch as the primary MSS or to a physically-separate switch elsewhere in the network.

Layer 3 routing functions are handled in much the same way as LANE services. The backup MSS takes over the IP and/or IPX address of the failed MSS and provides the layer 3 routing functionality for the clients of the failed MSS. The MSS' gateway address and all routing information remain the same and no changes to end user stations are required.

Network devices re-establish communications with the MSS and, depending on the protocol and service being used, may experience little or no disruption.

Functional Verification

OBJECTIVE

To verify the automatic recovery feature of MSS in an environment using both LAN Emulation (layer 2) and IP routing (layer 3) services.

Note: The Tolly Group also verified that the recovery operation is consistent whether the two MSS devices are connected to the same or different ATM switches.

A network was setup with two ELANs that were connected by the MSS using IP Routing (layer 3).

In order to verify that the network did recover from the loss of the primary MSS, two workstations were setup to continuously send an IP Ping to each other. This created a steady stream of traffic across the network which made it possible to determine when the network was down and how long the network required to come back up to operational status.

The Tolly Group’s goal was to verify that the same LANE and IP routing functionality was available after an induced “catastrophic failure” of the primary MSS. And that no operator intervention or network/client device reconfiguration was required to accomplish this.

TEST BED

An IBM 8265 ATM switch sat at the heart of the test bed, connected via OC-3 155 Mbit/s links to a primary MSS and a secondary MSS acting as a hot stand-by. An identical network configuration was loaded into each MSS, with one being configured as the primary MSS and the other as the secondary MSS. From the IBM 8265 ATM switch, 155 Mbit/s links connected a pair of IBM 8271 Ethernet switches, each of which supported 10BASE-T connections to downstream test workstations.

Logically, two Emulated LANs, ELAN CB-A and ELAN CB-B were created. Both ELANs contained 8271 Ethernet Switches configured as LANE clients and a single Ethernet workstation.

Redundant MSS Services

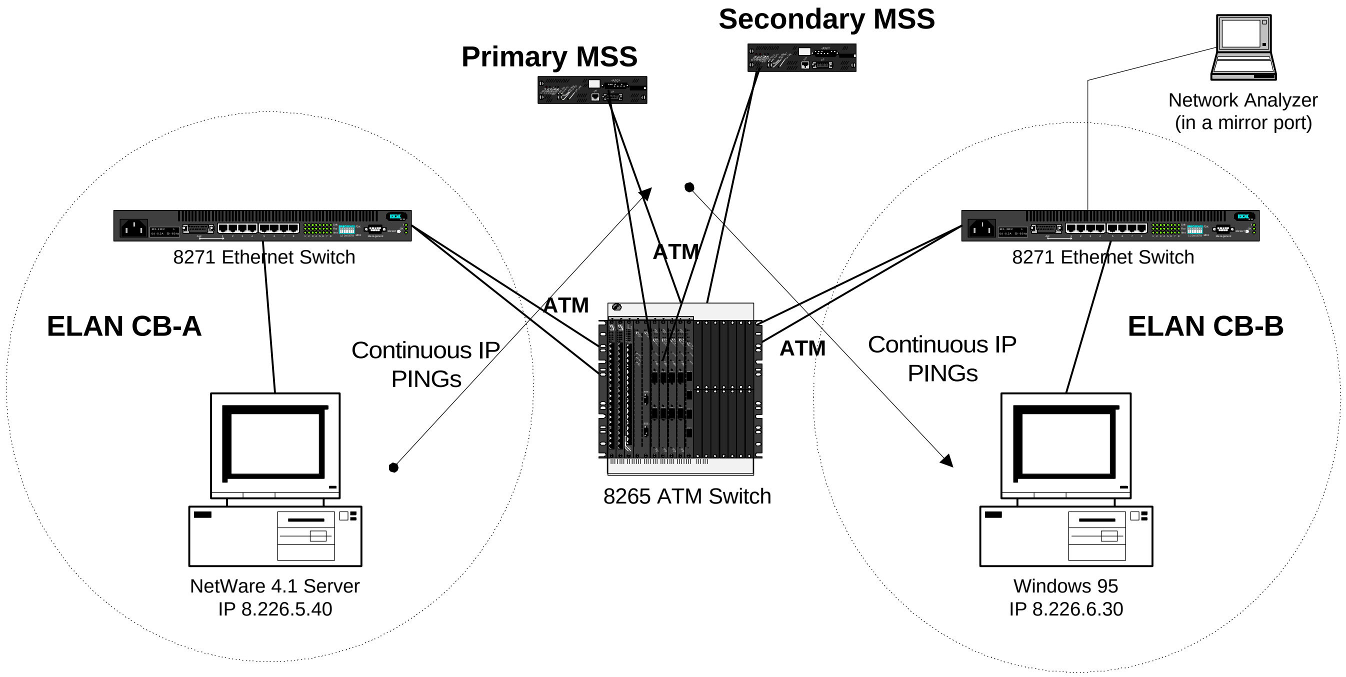

Test Procedure and Observations

The two Ethernet clients registered and joined their respective ELANs through the LEC in the Ethernet switch. The Ethernet switches implemented the LAN Emulation Client code on behalf of the native, legacy Ethernet stations. Furthermore, the MSS acted as a layer 3 IP router between the two logical LANs.

In order to test for redundant services, an IP PING was started from a station on IP subnet 8.226.5.xx (representing ELAN CB-A) to another station on IP subnet 8.226.6.xx (representing ELAN CB-B). The PING, initiated on an Ethernet segment, subsequently was routed across the ATM segment and back onto the Ethernet on the far side. The goal was to establish an active session between the ELANs and monitor activity as the primary MSS failed. The PING ran continuously, making it possible to determine when the primary MSS failed and when the secondary MSS came online.

The successful execution of the PING command proved that both LANE services, as well as IP routing, were available and functioning. The Tolly Group took a network trace using a network analyzer to show the traffic flow and to determine when the network went down and the amount of elapsed time before the network came back. The primary MSS was then powered off to simulate a device failure.

At this point, the PING command failed to elicit a response from the partner station “across” the MSS. Detecting the failure of the primary MSS, the secondary MSS took over the identity of the primary MSS. The LAN Emulation clients automatically re-established contact with the replacement MSS and, six seconds later, traffic once again was flowing between the two end stations involved in the PING test.

Tests were repeated with both Ethernet and Token Ring edge networks. In both environments, the MSS automatic recovery proceeded identically. The test was also run with both MSS devices connected to a physically separate, but logically connected, ATM switch. The results were the same as the single-switch test.

Broadcast Management

Whether it is a NetWare (IPX) server advertising the resources it has available, an NT (NetBIOS) client initializing and announcing its presence or a PING (IP) trying to locate its partner, there are many times when a network station needs to communicate with a group of stations simultaneously. It is the broadcast frame that accomplishes this function. A frame sent as a broadcast is delivered to every other station on the network. This group is referred to as the "broadcast domain." Such an approach assures that every "interested party" will receive the information. When the network contains bridges or switches, the frames are propagated to all connected LANs. Layer 2 switching/bridging, by definition, has no knowledge of any higher protocol and cannot intelligently manage broadcasts.

Unfortunately, those frames are blindly delivered to every station on the LAN whether the frame is relevant or not. For a station running only NetBIOS, an IPX broadcast is meaningless. AppleTalk means nothing to a station running only TCP/IP. In many cases, then, the traffic propagated across the network is immediately discarded by the end-station once the determination is made that the information in the frame is not relevant.

While the impact is negligible in small, stand-alone networks, broadcasts quickly become an issue when switching (layer 2) is used to build large campus networks. Unnecessary broadcast traffic not only wastes network bandwidth, but burdens each end-station with pointless processing as it must accept and then discard irrelevant broadcasts. And, with a large, switched network, there are more end-stations that may possibly be sources of broadcast traffic.

Many network managers are so concerned about the ill effects of such broadcasts, they say it is a limiting factor when designing a switched campus network.

LAN Emulation, the fundamental approach to integrating ATM and legacy LANs, links ATM-connected devices together in the logical equivalent of a legacy LAN — an Emulated LAN (ELAN). Although built using state-of-the-art ATM technology, ELANs still handle broadcasts in the same manner as legacy LANs. Thus, broadcast concerns impact the scalability of an ELAN.

IBM’s MSS includes a Broadcast Management component that intelligently manages broadcast traffic by monitoring the higher layer protocol activity in the network. Broadcast management optimizes bandwidth and eliminates waste by directing broadcast traffic to appropriate network stations while suppressing its propagation to those stations that would simply discard it upon arrival.

Before IBM MSS

Traditionally, network managers had no means of curbing broadcast traffic within a bridged or LAN switched network, so frames were delivered to all devices. Today, ATM’s ELAN functions similarly and layer 2 and layer 3 broadcast frames are delivered to all stations participating in an ELAN, even though the broadcast traffic is likely intended for only a subset of the stations.

Many vendors’ ATM implementations allowed ELANs to be connected at layer 2 via a virtual ATM bridge. This additional functionality exacted a price in performance. With ELANs thus connected, the broadcast domain encompassed both ELANs and the broadcasting issue became even more important, since more stations were generating and receiving broadcasts.

At some point, the burden of handling broadcast traffic would become a limiting factor in the size of the ATM ELAN. While multiple ELANs could be constructed to contain the broadcast traffic, such segmentation had performance implications, as ELAN-to-ELAN traffic would need to be relayed by an intermediate switch or router.

With IBM MSS

The broadcast manager transparently imposes order on the previously uncontrolled broadcasts present in the network. The ATM LANE Broadcast Unknown Server (BUS) intercepts all broadcasts and, instead of propagating the broadcast to all members of the ELAN, directs them to the MSS Broadcast Manager. Broadcast management is aware not only of the MAC-layer addresses of each of its LAN Emulation Clients but has higher-layer protocol awareness as well. (For example, it recognizes a NetWare server as the source of Service Advertising Protocol, or SAP, broadcasts.) After determining the valid recipient(s) of the frame, the frame is forwarded to only those stations and, in effect, filtered from the remainder of the ELAN.

For example, in a TCP/IP network, the broadcast management function works as follows: MSS "self-learns" traffic flows based on client broadcasts; MSS consequently decides how to respond to different traffic flows from LANE clients. For instance, an IP Address Resolution Protocol (ARP) frame transmitted onto an ELAN would normally result in a layer 3 broadcast being sent throughout the broadcast domain (i.e., every member station of a particular ATM ELAN, as well as all bridged ELANs) to determine the MAC address of the destination station. With MSS, however, the ARP frame is intercepted by the MSS' BUS and forwarded to the MSS Broadcast Manager. Since MSS maps the MAC address of the destination station to its IP address, MSS redirects the ARP frame as an IP unicast frame to the destination station.

The table below shows examples of how the MSS Broadcast Manager handles some types of IP, IPX and NetBIOS broadcast traffic. The various broadcast frames are handled differently by the broadcast manager.

Examples of MSS Broadcast Manager's response to broadcasts

Network Protocol	Broadcast Frame Type	MSS Handling
IP	ARP	ARPs are directed as unicast frames to intended destinations.
IPX	RIP SAP	RIP frames are forwarded only to routers. SAP frames are forwarded only to servers.
NetBIOS	Add Name Query	Implements NetBIOS name caching. Filters repeated multicast frames.

Functional Verification

OBJECTIVE

To verify that an IP ARP frame — normally delivered as a broadcast frame to all members of the broadcast domain — is forwarded only to the appropriate destination station when the MSS Broadcast Manager is enabled.

Note: An ARP frame created by a PING between two workstations was used to demonstrate broadcast management. The Tolly Group selected a PING since it is a simple function that results in a broadcast frame being sent out to look for the destination station's address. Therefore, it is possible to observe the operation of the PING without having to filter through a lot of other network activity.

TEST BED

On the physical side, an IBM 8265 ATM switch was linked via OC-3 155 Mbit/s links to two IBM 8271 Ethernet Switches. The Ethernet switches were “edge devices” and implemented the ATM LAN Emulation Client (LEC) function. A single MSS hung off the 8265 switch via an ATM connection. Two workstations were each connected via 10Base-T Ethernet to the IBM 8271 Ethernet Switches. A network analyzer hung off a port on one of the IBM 8271 switches.

The entire test bed represented a single ELAN, representing an IP subnet and, logically one broadcast domain.

Broadcast Manager

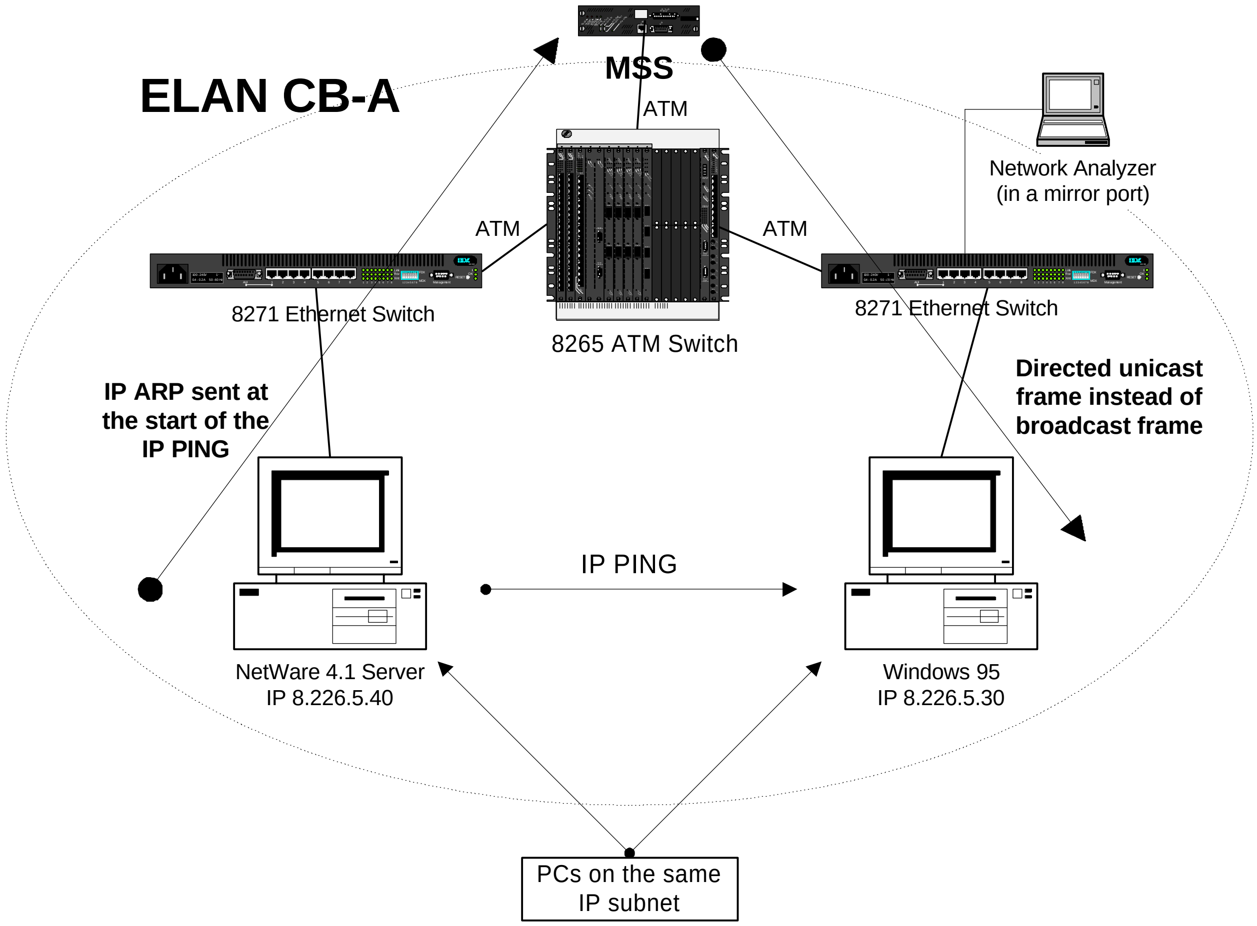

Test Procedure and Observations

The two 8271 Ethernet Switches were setup as LANE clients in the same ATM ELAN. (Note: The switches are ATM edge devices that implement LANE client functionality on behalf of the Ethernet devices connected to them.) The test of Broadcast Manager called for a PING to be sent between the two workstations. The two clients were on Ethernet and were each attached to an 8271 Ethernet Switch that provided connectivity to the ATM network. The Tolly Group attached an Ethernet analyzer to the same 8271 Ethernet Switch that supported the destination station of the IP PING. The analyzer attached to a mirrored port that displayed the same traffic that traveled across the destination station's port. The analyzer could then record and decode the traffic that was being sent to and from the destination station.

Initially, the test called for sending a PING between workstations with the MSS Broadcast Manager facility disabled, to learn how the network would respond. The PING was sent and the network analyzer recorded the resulting traffic. A broadcast (or ARP) frame looking for the address of the destination station was seen by the destination station. The destination responded to the broadcast and the PING completed successfully.

Next, The Tolly Group activated the MSS Broadcast Manager and executed the same test. In order to make certain that the PING function would result in the generation of a broadcast frame, the ARP caches on the workstations were cleared prior to the test. The PING was sent again with the network analyzer capturing all traffic headed for the destination station. The decode showed that this time no broadcast frame was sent to the broadcast domain, but instead MSS Broadcast Manager sent a unicast frame directly to the destination station. The receiving station responded and the PING completed successfully.

Inter-Elan Connectivity

Routers provide the indispensable links that transform “islands of connectivity” into a campus network. Those same routers, however, can themselves impede communications. Analysts tell us that traffic profiles are changing dramatically with 80% of the traffic now leaving the local LAN segment. Network managers are concerned that traffic crossing router hops will experience degraded performance as it competes for processing power and limited bandwidth associated with transiting routers. Routers can become a significant performance bottleneck and limit the overall effectiveness of campus-wide communications.

In order to offset potential routing bottlenecks, IBM has woven two new services into MSS: Route Switching and Super ELANs. One of the other major benefits of Route Switching and Super ELANs is reduced — actually eliminated — latency. Bypassing the router and utilizing Route Switching means that hop-by-hop forwarding is eliminated, as is the latency associated with that function. This results in a faster response time for accessing data on the network. Note though that Route Switching is a layer 3 service, while Super ELANs is a layer 2 bridging function.

Route Switching

Route Switching is an implementation of the ATM Forum's "Multiprotocol over ATM" (MPOA) specification and uses virtual channel connections (VCC) to enable network end stations belonging to different LANE ELANs to establish direct circuits between them, thus bypassing routers for all but for initial communication.

This "router bypass" allows the ATM network to scale since some routing is handled by LANE clients, rather than being directed at the router core. This, of course, reduces the router-as-bottleneck situation that would otherwise exist. From a migration and cost-savings point of view, the reduced burden on the network's core routers allows for either: a) downsizing of the central routers, or b) reducing the load on and extending the life of the existing routers.

Route Switching is based, in part, on the Internet Engineering Task Force's Next Hop Resolution Protocol (NHRP), which enables the packet forwarding function of intermediate routers along the data path to be bypassed. One of the real benefits of MSS, is that Route Switching enables MSS to extend true multiprotocol support to non-routable protocols — even for SNA and for NetBIOS —that are typically unsupported by pure NHRP.

Super ELANs

With change the only constant, even the most carefully considered network design can soon require modification. Geographically separate groups of users can, overnight, become part of the same virtual department. In ATM terms, ELANs with relatively few connectivity needs one day may need to be united the next.

Using the Super ELAN capability of IBM's MSS, this can be accomplished simply and cleanly via a management station function. A "superset" of ELANs can be dynamically created and modified such that all members of all ELANs achieve "zero hop" connectivity with each other. Essentially, Super ELANs provide ELAN users with a substitute for the layer 2 bridging functionality in an installed base of conventional routers. As with Route Switching, routers are bypassed and, thus, that potential bottleneck is avoided. Furthermore, this can be accomplished without requiring any modifications to client station drivers. MSS creates and maintains the connections necessary for inter-ELAN communications to occur without the need for routers.

Before IBM MSS

Typically, in an enterprise-class network, core routers pass traffic traveling across an ATM backbone to an ELAN or IP subnet. When the traffic is routed (at layer 3), every router in the hop-by-hop data path looks at every frame to determine where they each should be delivered. However, as networks grow in size, and as LAN switches pass vast amounts of data to backbone routers, the resulting load overcomes the core routers, creating performance bottlenecks. Furthermore, aggregate traffic levels are limited to the bandwidth of the physical interfaces into and out of the core routers.

Eventually, as networks grow and routers become overtaxed, they require an upgrade, or additional routers must be added to handle the increased load. Implementing such a solution, however, is an expensive proposition.

With IBM MSS

Route Switching and Super ELANs, by contrast, allow client stations to bypass routers (except for initial processing). This reduces the overall reliance on core routers. Client stations outfitted with Route Switching drivers have the capability to communicate with stations in other ELANs or subnets without requiring a router to process every packet. As mentioned previously, Route Switching intelligence allows an ATM SVC to be set up between LANE clients after MSS processes the initial route discovery between the devices.

Super ELANs work in a similar fashion but connect entire ELANs with one another.

In order to implement route switching, a piece of code that supplements the network adapter driver is deployed on an end-user's computer. (Note: Route Switching software can be deployed incrementally. It is not necessary to implement the drivers on all LAN clients.) When traffic begins to flow between two stations, the route switching client initiates layer 3 call setup by requesting the MSS to dedicate a VCC between the clients. In a routed network, this VCC creates a short cut for data flowing between the clients, eliminating the hop-by-hop routing that each frame would otherwise traverse. Thus, route switching is more efficient than routing since it allows end stations to communicate across a VCC between the source and destination stations at layer 2. In fact, route switching can be thought of as zero-hop routing.

Functional Verification

OBJECTIVE

To verify that MSS' Route Switching enables direct communications between two network stations of different ELANs in a relatively complex, layer 3 environment. (Note: The verification approach for Super ELANs is quite similar except that no drivers are required for individual stations. Verification of Super ELANs was conducted with successful results.)

In order to demonstrate the potential performance gains of route switching, it is beneficial to compare "before" performance with that of a heavily-loaded router. The Tolly Group used a frame generator to introduce extra traffic into the router from ELANs not under test.

The Tolly Group tested the performance of Route Switching by determining the latency of a continuous PING across this network both with and without Route Switching enabled.

TEST BED

The environment for the Route Switching verification consisted of an IBM 8265 ATM Switch connected by 155 Mbit/s OC-3 links to a pair of IBM 8271 Ethernet Switches. An OC-3 pipe also connected an IBM MSS to the 8265 ATM Switch. Each 8271 Ethernet Switch had one client station attached via 10Base-T Ethernet. An Ethernet analyzer also hung off a mirror port on one of the 8271 Ethernet Switches, enabling it to monitor all traffic going to or coming from the locally attached client.

Both client stations were running TCP/IP and belonged to different IP subnets and to different ATM ELANS. The first client was attached to ELAN CB-A and IP subnet 8.226.5.xx, while the second client was attached to ELAN CB-B and IP subnet 8.226.6.xx. The IBM MSS was configured to route IP traffic between the two subnets.

The same logical test bed was used for latency measurements, with one addition. A second IBM MSS was connected to the 8265 ATM Switch and was used explicitly as a packet generator, pumping out about 45,000 pps. Its job was to pass traffic to the MSS making the routing decisions. While this traffic rate is quite high, the goal was to simulate an overburdened router. (Note: Without route switching enabled, the MSS functions as a standard core router.)

Route Switching Over Ethernet

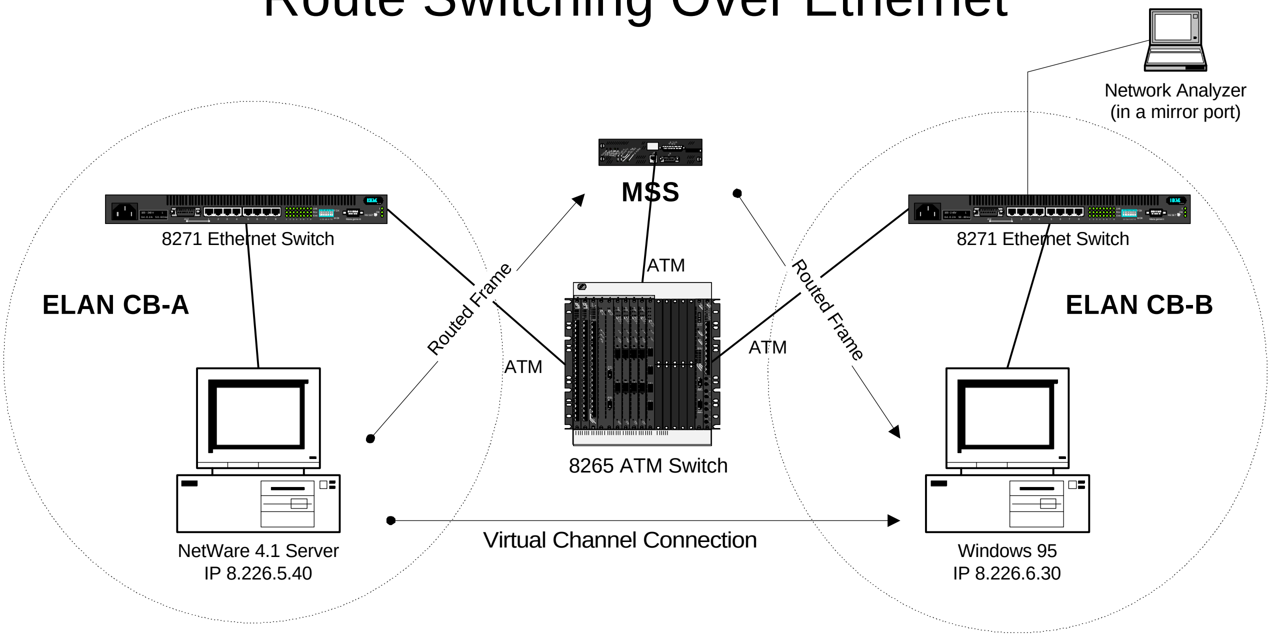

Test Procedure and Observations

Both workstations' Route Switching drivers were loaded in three different configurations in order to verify proper operation of the Route Switching facility.

- Neither client had Route Switching drivers loaded. This served as a control test to verify that traffic would be routed across the network correctly.
- The first client, attached to ELAN CB-A had Route Switching loaded, the second client on ELAN CB-B did not. This demonstrated that Route Switching could run uni-directionally.

- Both clients had Route Switching loaded. This proved that bi-directional traffic could be route switched.

These three configurations tested all possible scenarios for implementing Route Switching across the test network.

For each configuration, The Tolly Group engineers initiated a continuous PING between the two stations. The source and destination MAC addresses of each frame were noted using the network analyzer connected to the Ethernet switch. The analyzer monitored all traffic going to or coming from the ELAN CB-B client. From the source and destination MAC addresses of the captured frames, The Tolly Group determined whether the frame traveled across a short cut path, bypassing intermediate routers, or whether it was routed through the MSS.

Initially, The Tolly Group engineers conducted the baseline test without Route Switching loaded on either client; the resulting data and traces from the network showed that all frames were routed on layer 3 between the clients.

After generating a baseline, Route Switching was enabled on the ELAN CB-A client but not loaded on the CB-B client. The initial frame from the ELAN CB-A client went to the MSS, which set up a uni-directional VCC between clients from CB-A to CB-B. The analyzer determined that all traffic sent from the ELAN CB-A client to the CB-B client was transported at layer 2 across the VCC and no longer routed. Once it sets up this VCC, the client caches the route for future use. The analyzer also showed that all of the traffic from the CB-B client back to the CB-A client passed through the MSS at layer 3 since the CB-B client didn't have the Route Switching code loaded.

Finally, The Tolly Group activated Route Switching on both clients. Once again the initial frame from the client sending data instructed the MSS to set up a VCC between the client stations. This time the VCC was bi-directional, allowing layer 2 transfers in both directions. Using the analyzer, we were able to determine that traffic to and from the ELAN CB-B client traveled at layer 2 after the initial frames had set up the VCC.

The flexibility of being able to install the Route Switching drivers on particular client stations or on all stations allows network managers to add Route Switching to their networks in a phased approach. Route Switching can be enabled on servers initially and later implemented on workstations. This provides a flexible method of rolling out Route Switching to a network as time and dollars permit. Moreover, the client function can also be deployed on network edge devices, relieving users of having to open up their desktops or their servers for the software.

The Tolly Group engineers then turned their attention to testing for latency. Using the second MSS to generate traffic, the network was tested for latency on both routed traffic and then for route switched traffic between the two clients. With a continuous IP PING bouncing between the clients, The Tolly Group used a network analyzer to measure latency. All of the latency tests were conducted in a routed environment with the central router slightly congested (around 40%).

Without Route Switching drivers loaded, the latency of a single frame from the IP PING across the network was measured at around 25ms, but varied slightly due to network congestion. Once Route Switching was enabled, the latency of the same IP PING dropped to less than 10ms and showed no variation since the congestion on the router was no longer an issue.