

**Mixed-Media LAN Emulation
over an
ATM Backbone Solution
Test Lab Experience**

Introduction	Page 4
Solution Design and Configuration	Page 5
Physical Network Design	Page 6
Logical Network Design	Page 8
Network Management	Page 17
Network Management on Windows NT Platform	Page 17
Network Management on the AIX Platform	Page 27
Fault-Tolerance Testing	Page 34
8270 Token-Ring Switch Failures	Page 34
8274 LAN Switch Failures	Page 35
8271 LAN Switch Failures	Page 36
MSS Server Failures	Page 37
8265 Failures	Page 38
Performance Testing	Page 40
Test Configuration	Page 40
Test Results	Page 41
Conclusion	Page 43
Definitions	Page 44
Address Table	Page 46
Related Web Sites	Page 48

April 30, 1999

Introduction

This document describes a tested solution for an enterprise campus network that provides both Token-Ring and Ethernet switching to the desktop and communication across the campus over an ATM backbone. It describes the design, configuration, management and performance of the solution as we implemented it in our test lab. We've included considerations we discovered and tips we developed to help others avoid problems. The document is organized into the following sections:

Solution Design and Configuration - describes the physical and logical network design and high-level configuration steps. For details about the steps we took to configure this network, see **Detailed Configuration Steps** at:

wwwidd.raleigh.ibm.com/netsolut/atm/mixed/steps/detailedsteps.html

Network Management - describes the network management products tested and some of the monitoring performed with this solution. It also provides details about how we installed these products.

Fault-Tolerance Testing - describes our redundancy and fault-tolerance testing and results.

Performance Testing - describes our performance testing and results.

Conclusion - summarizes our experience and provides general guidelines for implementing the design in your network.

Address Table - lists the IP and ATM addresses we used for configuring the devices in this solution.

Definitions - defines some key terms.

Related Web Sites - provides a list of the Web links that we refer to in this document.

The configuration described consists of two IBM 8270 Nways® LAN Switch Model 800s with MSS Clients, two 8271 Nways Ethernet LAN Switch Model 712s with ATM UFCs, two IBM 8274 Nways LAN RouteSwitches (one Model W93 and one Model W53), two IBM 8265 Nways ATM Switches and two IBM Nways Multiprotocol Switched Services (MSS) Servers. The MSS Server routes IP and IPX traffic between the subnets on the ELANs. SNA and NetBIOS are bridged between LANs over the ATM backbone. Appletalk is also bridged by the MSS between the Ethernet switches. We used IBM Nways Campus Manager for AIX® and Tivoli NetView®, as well as IBM Nways Workgroup Manager for Microsoft® Windows NT® for network management. The solution outlined is a common implementation, a simple LAN Emulation environment.

Note: We used the 8271 Nways Ethernet LAN Switch as the Ethernet edge device in this solution but the 8371 Multilayer Ethernet Switch is an excellent, new alternative that provides next generation ATM network services. For example, the 8371, along with MSS, delivers high performance and ease of installation with multiprotocol over ATM (MPOA) virtual one-hop routing over the ATM backbone.

Solution Design and Configuration

This section describes how we implemented the mixed (Token-Ring and Ethernet) LAN Emulation (LANE) over an ATM Backbone High-Level Design in the test lab. See **High Level Design** at:

wwwidd.raleigh.ibm.com/netsolut/atm/mixed/mm-hld.html

The high-level design includes a description of the proposed physical and logical high-level design for the desktop and backbone, an overview of the products, and an explanation of the resiliency and scalability of the solution. Before we began the actual implementation, we found it helpful to diagram the physical and logical design and the routing and bridging configuration. We also created a table with the ATM addresses, IP addresses and MAC addresses for the devices (see the "Address Table" section).

We used ImageNet's NetFormx network design tool to create physical and logical diagrams and document design attributes. We could click on any product in the diagrams to view attributes.

The design we implemented for this test environment is a simple configuration that would be relevant in most medium Token-Ring and/or Ethernet networks. Overall, it is best to keep your design as simple as you can. However, the concepts used in this configuration could be extended to more complex networks. Most of the tips in this document are relevant to any variation of this solution.

Although we kept the configuration as simple as possible, we did take advantage of some specific support for improved network performance. Multiprotocol over ATM (MPOA) support is enabled for the Token-Ring portion of the network on the MSS Client and MSS Server by default. It allows more traffic without impacting MSS Server routing utilization. It enables local shortcuts within the switch and enables LANE shortcuts over the ATM interface.

Physical Network Design

The test lab configuration consisted of the following IBM networking products:

Two IBM 8265 Nways ATM Switches with operational code V3.3.5

Two IBM Nways MSS Server modules with software code V2R1 PTF3 (except for performance testing where we used V2R2)

Two IBM 8270 Nways LAN Switch (Token-Ring) Model 800s at V5.1B and MSS Client modules at V2R1 PTF4

Two IBM 8271 Nways Ethernet LAN Switch Model 712s at V3.22A with ATM uplinks at V1.07

One IBM 8274-W53 and one IBM 8274-W93 Nways LAN RouteSwitch at V3.1.8

An actual customer implementation could have met the same requirements using a configuration with a subset of these products. However, we chose to use a mix in the test lab to demonstrate the strengths and weaknesses of each device. For example, we used both 8271s and 8274s for Ethernet switching.

Our primary intention was to provide a modular and scalable design with a full redundant backbone in terms of its devices and their connections. The floor access devices (for example, the 8270, 8271 and 8274 switches) also have redundant links to the backbone. The main advantages of a modular design are that it can be tested and "proven" in a test lab and then repeated in a production environment without a lot of additional testing. This reliability is extremely important to financial institutions.

We envisioned that the test floor configuration could represent a three-building site with multiple telecommunications closets containing switches connecting into an ATM backbone as shown in Figure 1.

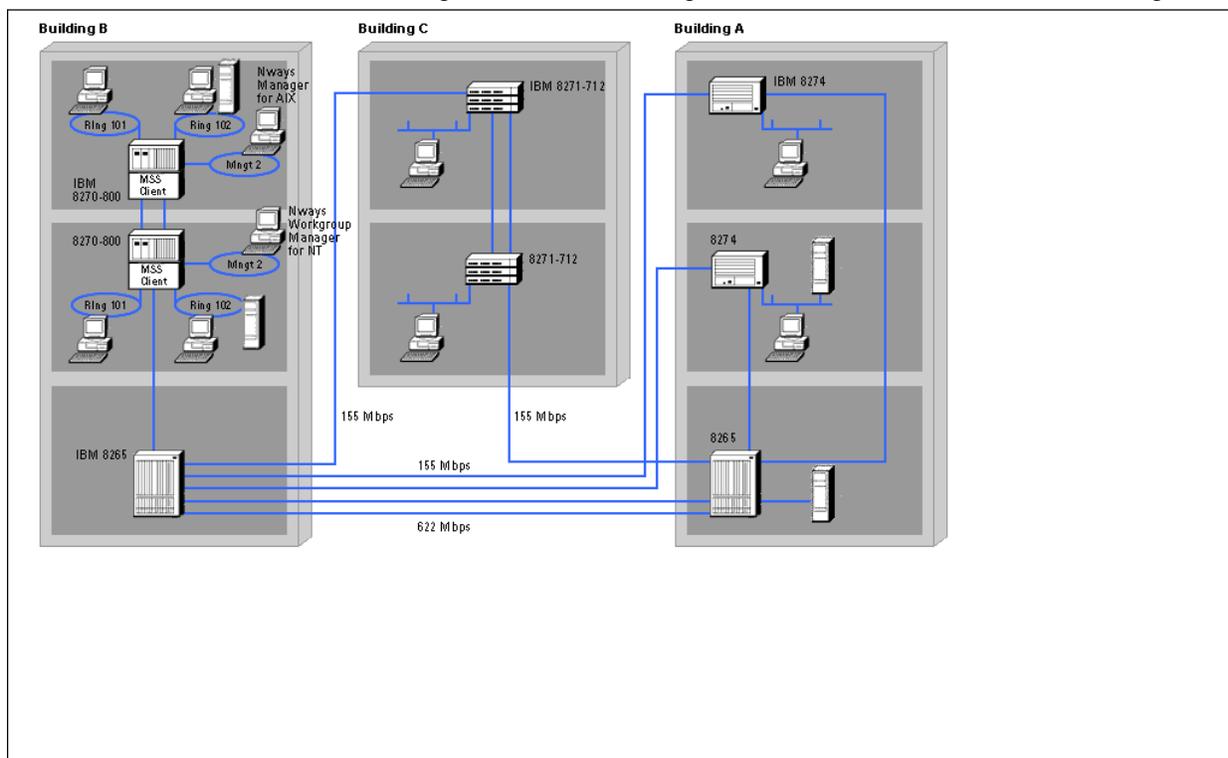


Figure 1. Network Overview

The telecommunications closets in Building B support a large concentration of users. IBM 8226 Multistation Access Units and 8238 Token-Ring Stackable Hubs connect the PCs to LAN segments connected to the 8270 switches in the second and third floor telecommunications closets. The 8270 on the second floor is ATM-attached to an 8265 on the first floor and the 8270 on the third floor is ATM-attached to an 8265 in Building A. Furthermore, 8270 switches are connected to each other through a TokenPipe to provide extra resiliency. For example, if the fiber link connecting the third floor 8270 to the backbone is broken, this 8270 can still communicate with the rest of the network through its TokenPipe link to the other 8270 on the second floor. PCs, servers, network management stations, and other heavily used devices can be directly connected to switched ports on the 8270 switches, providing 32-Mbps, dedicated bandwidth to the directly connected device.

Most of the servers are located in Building A. Two 8274s with redundant power supplies and management modules, one on the third floor and one on the second floor, provide redundancy for the users in their domains. Each 8274 is ATM-attached to a different 8265 backbone switch, so if any one switch fails, a backup path exists through the network. 8224 hubs connect the user LAN segments to the 8274 switched ports. However, you can also use switched 8274 ports with the 10/100 autosense function to directly connect PCs, servers or other heavily used devices and provide a maximum bandwidth of up to 200 Mbps.

Building C is a remote site though still reachable with either multimode or single-mode fiber, supporting a small concentration of users. If the distance causes constraints, consider some form of wide area network (WAN) connectivity. For example, employing routers and leased lines is a common solution. This design is a campus solution but it can apply to sites with WAN connections. PCs are directly connected to two 8271 Ethernet LAN switches that are connected to an 8265 backbone switch through an ATM fiber connection. Resilient links for the ATM uplinks for these two switches in Building C are provided by a VLAN trunking link (VTL) over copper cabling.

The backbone shown here consists of two 8265 Nways ATM Switches in the first-floor telecommunications closets of Buildings A and B. They are located in different places in case of a physical disaster. The 8265 in Buildings A and B have an MSS Server module. The two ATM backbone switches are connected with two fiber connections that would preferably be run in separate paths. As a general design guideline, place fiber cables in diverse ducts between the buildings unless it is safe to do otherwise.

Our test lab connections represented 622-Mbps connections (one single-mode fiber and the other multimode fiber) from the 8265 in Building A to the 8265 in Building B. High-throughput servers are ATM-attached to the switches. These 8265 ATM switches were also each equipped with redundant CPSW (control point switch) modules, power controllers and power supplies for a highly fault-tolerant backbone configuration. The outage of one of the 8265 switches will not affect the end users adversely because the second 8265 will provide backup services as well as its own normal services. The 8265s communicate with each other through the PNNI-Phase 1 protocol, which is a dynamic ATM routing protocol together with load-balancing feature. This provides a highly efficient form of connectivity between the PNNI nodes. Figure 2 shows an overview of the physical connections.

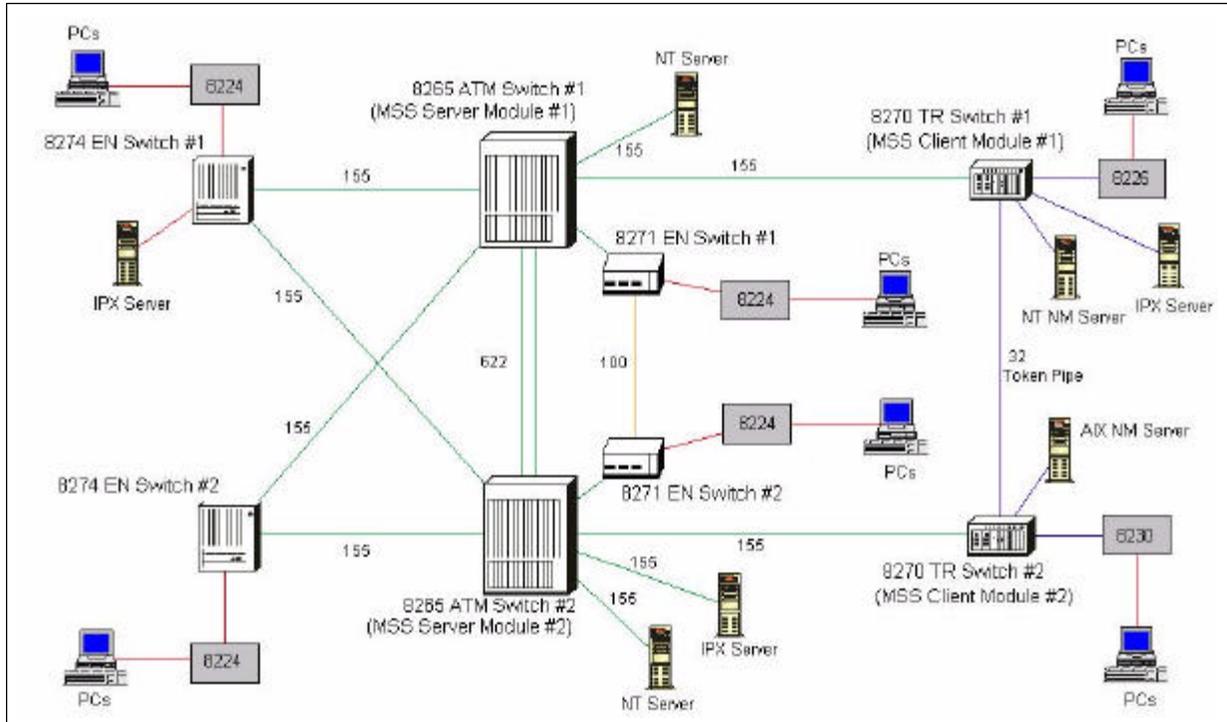


Figure 2. Physical Network Connections

Logical Network Design

A major advantage of an ATM network as discussed in previous designs is that the physical and logical design of the network are separated. This means that as long as an ATM-attached device has access to the ATM infrastructure, it can join any ELAN. Our configuration has four ELANs. However the concepts used in our configuration apply to configurations with any number of ELANs.

As you can see in the logical network diagram (Figure 3), the MSS Server routes IP traffic between the IP subnets and IPX traffic between the IPX networks in our network. Basically, in our configuration, routable protocols (IP and IPX) were routed, which eliminated sending Layer-2 broadcasts to all the ELANs.

However, the non-routable protocols (such as SNA and NetBIOS) were bridged by the MSS Server. The MSS Server also provides translational bridging between Ethernet and Token-Ring segments as well as the traditional bridging functions like source-route and transparent bridging. MSS also bridged AppleTalk traffic, because routing of AppleTalk was unnecessary with the existence of only one Ethernet ELAN. Furthermore, we enabled Spanning Tree algorithms in our network devices to facilitate the implementation of dual links to 8265s from the 8270 and 8274 LAN switches.

Tip: Translational bridging is a processor-intensive operation. It is important to limit the amount of translational bridging that must be performed. As much as possible, group in the same LAN type resources that need to be accessed with protocols that must be bridged.

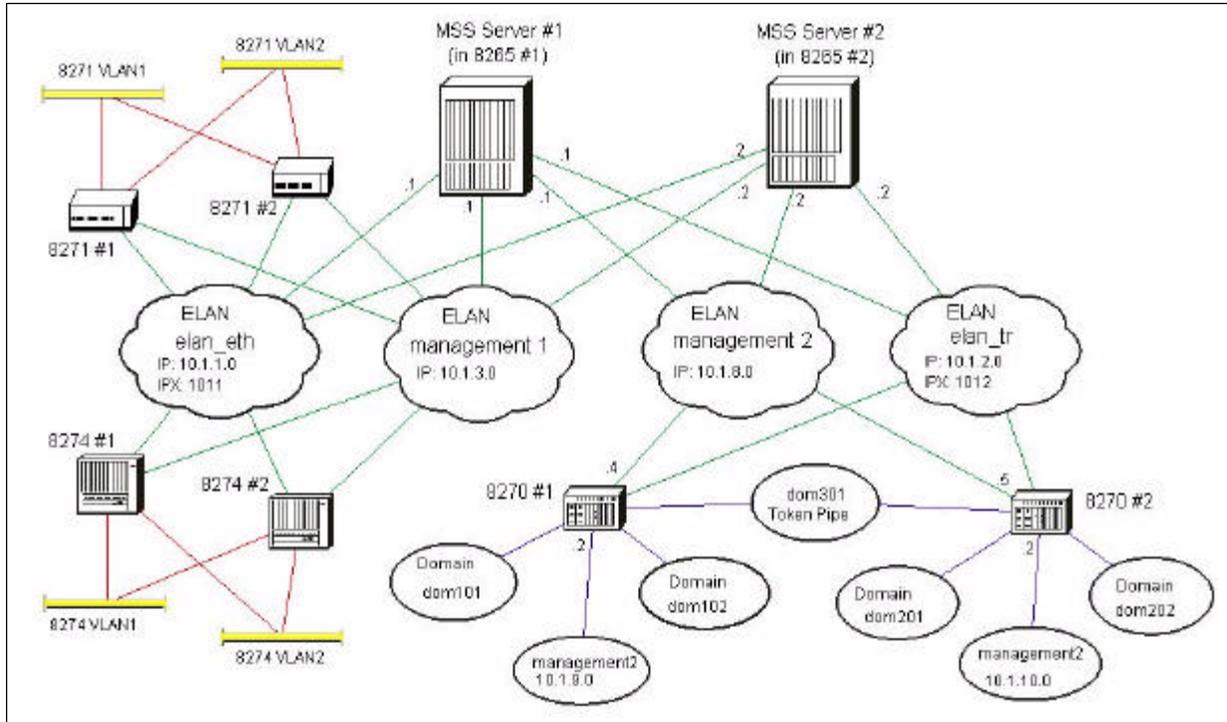


Figure 3. Logical Network Diagram

The Token-Ring switches "source-route" user data (IP, IPX, SNA, and NetBIOS) within the directly connected Token-Ring segments and also to the ATM backbone. However, MSS Clients route the SNMP management data to and from the Token-Ring devices. We purposely separated user traffic from SNMP traffic for efficiency, manageability, and network security. The two management ELANs (shown in Figure 3) are an Ethernet ELAN for the Ethernet devices (named *management1*) and a Token-Ring ELAN for the Token-Ring devices (named *management2*). We also configured the MSS Clients for MPOA support, so that IP shortcuts within a switch (local) and IP shortcuts between switches reduce the amount of IP traffic through the MSS Server.

Similarly, Ethernet switches (8271 and 8274) bridge the user data among their LAN ports, which were grouped in virtual LANs (VLANs). We used port-based VLANs for both 8271 and 8274 switches. Furthermore, each VLAN in these switches was bound (bridged) to an appropriate ELAN in the backbone (*elan_eth* in our configuration) through the LAN emulation client (LEC) proxy function of these switches. Again, the MSS Server routes or bridges the data pertaining to these devices depending on the data type. The SNMP management data from these devices is also bridged to the *management1* ELAN.

To provide more support for redundancy and load-balancing for the IP protocol, we employed the redundant default gateway function of MSS as shown in Figure 4.

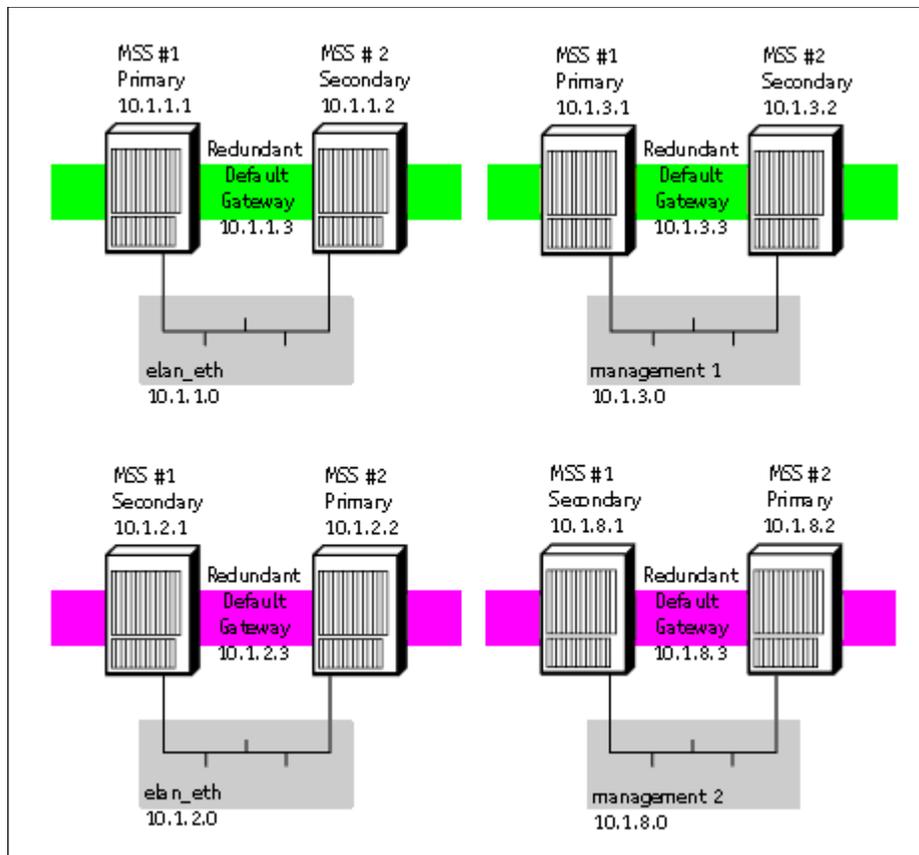


Figure 4. Redundant Default Gateway Setup

In a redundant default gateway configuration for each IP subnet, the end stations will not need to reconfigure their default IP gateway if the primary default gateway fails.

To support backbone redundancy, we used two MSS modules when one would have been sufficient for the number of stations connected to our network. This redundancy enables LANE servers and services to be available on a hot standby basis in the backup MSS if the primary MSS fails. As shown in Figure 5, MSS #1 acts as the primary LANE server for the *elan_eth* and *management1* ELANs while at the same time acting as the backup for the *elan_tr* and *management2* ELANs. MSS #2 acts in a similar way by providing primary LANE services for the *elan_tr* and *management2* ELANs and backup LANE services for the *elan_eth* and *management1* ELANs. This approach provides a fault-tolerant, load-balanced and distributed configuration.

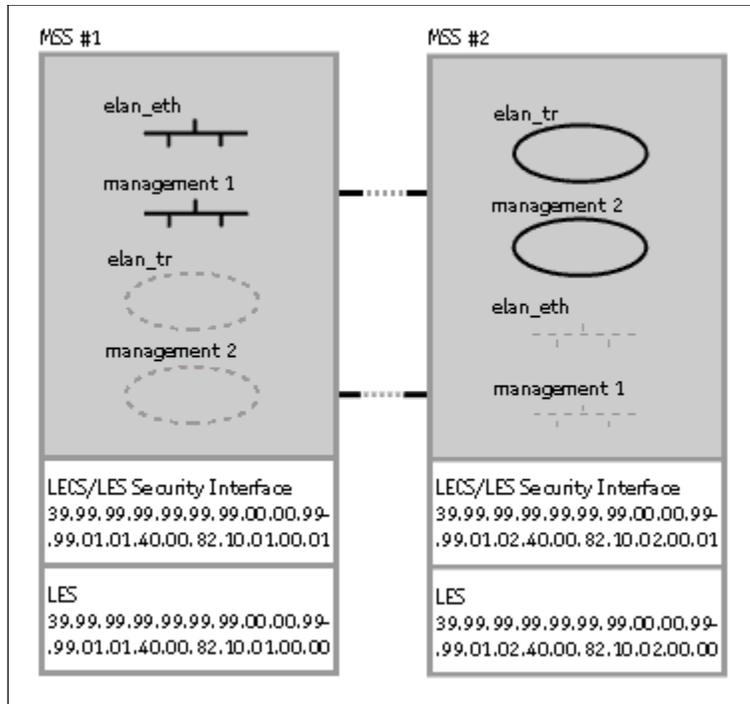


Figure 5. Logical Configuration of MSS Modules

We also mapped an IP subnet to each ELAN. This is a likely scenario in most customer sites that currently plan to use IP for all of their data communications. We chose a private (unregistered) 10.x.x.x IP network and subnetted it into further networks by using 255.255.255.0 for the subnet mask. We also chose RIP as our routing protocol.

The 8270s in Building B are attached to four Token-Ring LAN segments (Rings 101 and 102 to Switch#1 and Rings 201 and 202 to Switch#2), which support the majority of users. When we configured the switches we named the domains to correspond to the ring numbers, for example, domain dom101 for Ring 101. We also used Ring 301 to represent the TokenPipe connection between the switches in both switches. The TokenPipe provides a redundant link for domains dom101, dom102, dom201 and dom202 in both switches. The ATM LEC on each of the 8270s is attached to Ring 103. Ring 103 refers to the ATM network in our design and is seen as Ring 103 by all of the Token-Ring LAN switches and devices. We enabled the SRB function in the MSS client in both switches to bridge between their respective LAN domains (dom101 and dom102 in switch #1), the TokenPipe (dom301) and the ATM network(Ring103).

Each 8270 and MSS Client require a separate IP address for management. On each MSS Client, we configured an IP address for the LEC on the *management2* ELAN (10.1.8.4 for Switch#1 and 10.1.8.5 for Switch#2). We also configured an additional Token-Ring interface on the MSS Client and assigned it an IP address (10.1.9.1 for Switch#1 and 10.1.10.1 for Switch#2). We defined a domain in each 8270 switch that corresponds to the Token-Ring interface on its MSS Client, and we assigned an IP address to that domain (10.1.9.2 in Switch#1 and 10.1.10.2 in Switch#2). We named the domain in each switch *management2* as shown in Figure 6. Our network management stations resided off these domains. Therefore, all management traffic was routed through the MSS Clients to the *management2* ELAN.

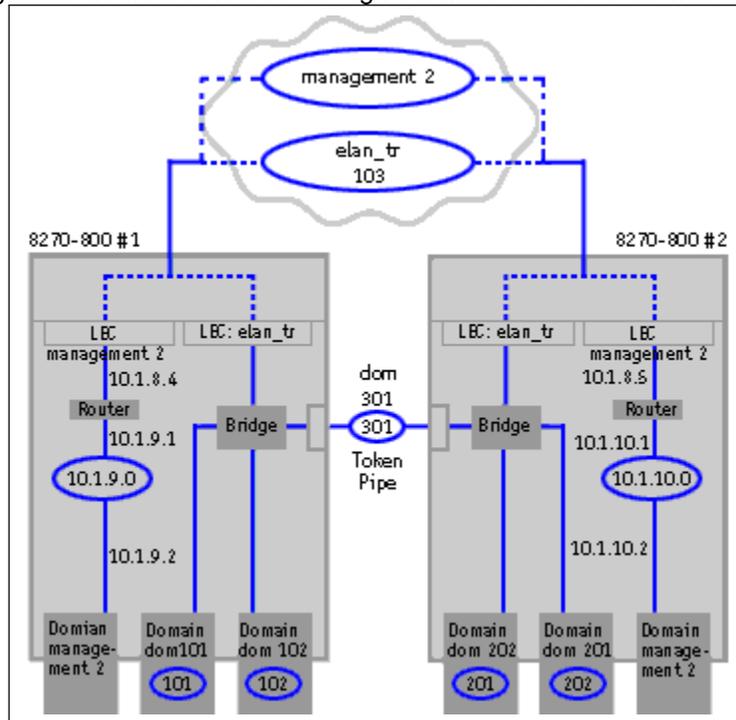


Figure 6. Logical Configuration of MSS Client

The MSS Clients offer additional features not available on the 8270s. MPOA is enabled by default. This support allows IP shortcuts to be made both within the switch and between switches, which reduces the amount of traffic through the MSS Server. This filtering reduces the amount of broadcast traffic on each interface by sending out broadcasts only for learned IP subnets on that interface.

ELAN Size

One major decision when you design an ATM solution is the number of stations in an ELAN. There are two different stations to consider: ATM-attached and legacy-attached stations. For ATM-attached stations, a single MSS can theoretically support a maximum of 1000 stations. Regardless of the number of ELANs for which the MSS is providing LES/BUS functions, the total number of ATM-attached stations should not exceed 1000.

The legacy-attached stations will be connected through a proxy device to the MSS LES/BUS servers. Proxy devices, such as the 8271 and 8274 are providing access for the legacy-attached stations into an ATM environment. The proxy devices use a table of LANE ARP entries to keep track of MAC-address to ATM-address mappings. The proxy devices use the mappings to establish sessions with those MAC addresses. The size of the LE ARP cache potentially limits proxy devices, especially transparent-bridging devices such as the 8271 or 8274. The LE ARP cache size for the 8271 Models 524, 612, 624 and 712 is 1024. The LE ARP cache size for the 8274 is 4096.

Therefore, any given ELAN for Ethernet should not exceed 1024 legacy-attached stations when you use the 8271 and 4096 legacy-attached stations when you use the 8274. The MSS Servers' LECs' LE ARP table sizes and transparent bridge cache sizes must be able to support the number of downstream Ethernet MAC addresses.

To design a fine-tuned and balanced ATM network other factors such as the current traffic profile, expected performance, availability and type of applications in use should also be taken into account as well as the above theoretical limits.

General Configuration

The 8271s are easy to configure. If your ELANs are configured to allow devices to join by ELAN type (such as Ethernet), just connecting the 8271s to the ATM network gives you complete connectivity. We also made the following changes:

Gave the device an IP address, subnet mask, and default gateway.

Set the community name and gave the IP address of the management station that will receive traps.

Placed ports in correct VLANs (because we were using VLANs).

Set 8271 to 8271 ports as VLAN Trunk (VLT) ports (because we were using VLANs and having a resilient pair for the ATM port).

Set the VPI bits on the ATM port to 1 (so that the VCI bits increase to 10 and match the 8265).

Set the ELAN name that each VLAN will connect to.

Tip: Spanning Tree is not supported in the 8271 switches when ATM or resilient links are used. Be careful not to create loops in your network. Consider this when setting up resilient links.

Tip: The *8271 User's Guide* states that the VPI and VCI bits are negotiated using the Integrated Local Management Interface (ILMI) with the ATM switch. This negotiation does not currently work with the 8265 so they should be set to match.

The 8274s were easy to configure. We did the following:

Defined a port-based VLAN for PCs and servers.

Defined another port-based VLAN for management and assigned an IP to it.

Defined ATM parameters, for example, the UNI version, LECS, and ELAN name.

Mapped each VLAN to its appropriate ELAN.

Defined SNMP parameters for network management.

The 8270s were easy to configure. We did the following:

Gave the management domain an IP address, subnet mask, and default gateway.

Set the community name and gave the IP address of the management station that will receive traps.

Configured ports in the correct domains.

Configured the MSS Client to join the ELAN name (*elan_tr*).

Tip: On an 8270, you need to clear the configuration when switching from a configuration with an ATM UFC to a configuration with MSS Client.

Tip: When you download code to an 8270, make sure that you configure not only the correct address of your TFTP server but also the correct domain as well.

The MSS Client is a little more complicated to configure, partly because of the added function. In general we did the following:

Added the Token-Ring domains.

Created and defined a LEC to join *elan_tr*.

Created and defined a LEC to join *management2* and configured an IP address for management purposes.

Configured an SNMP read/write community name and provided the IP address of the management stations that will receive traps.

MPOA is enabled by default.

Configured source-route bridging.

Configured IP routing for management purposes.

The 8265 ATM Switches were also easy to configure. We did the following

Gave a unique ATM address to each switch within a peer group.

Connected all of the ATM modules to the ATM backplane.

Enabled as UNI the ports that connect to Ethernet switches, MSS Clients, or servers.

Configured the correct UNI version for the end device ports: 3.0, 3.1, 4.0 or auto.

Enabled as PNNI the ports that connect the 8265s together.

Created a LEC assigning it a MAC address, IP address, subnet mask and default gateway.

Configured a read/write community name and provided the IP address of the management station that will receive traps.

Tip: To determine that you have the correct working UNI version for the end device ports, check it on the 8265 and in the MSS LES/BUS. Issue the *show reachable address slot#.port#* command on the 8265 hub. You should see the expected ATM address. Determine whether the device has joined the LES/BUS, by issuing the following Talk 5 commands:

```
NET 0  
LE-S  
WOR
```

Then select the correct LES and issue the DATABASE LIST ALL LEC command and check whether the expected ATM address exists in the database.

Tip: Do not use the network portion of the ATM address that is default on the 8265. If you do, and an 8265 that has not had its ATM address changed is connected to the network, PNNI routing problems will occur.

Tip: In a very large network with a lot of broadcasts, it is best to create a "Management ELAN" that has only LECs from the network devices in it. This way they do not have to process all of the other broadcasts.

The MSS Server can often seem very confusing to configure. However, once the different concepts (such as LAN Emulation and routing) are understood, it becomes quite manageable. In general, we did the following:

Created and defined four ELANs (enabling BUS Monitor, BCM, LES/BUS redundancy support).

Created and defined a LEC for each ELAN.

Configured each LEC with IP addresses each in different subnets and configured RIP routing.

Configured an SNMP read/write community name and provided the IP addresses of the management stations that will receive traps.

Configured adaptive source-routing transparent (ASRT) bridging.

Configured IPX routing.

MPOA was enabled by default.

For more detailed step-by-step notes about how we configured all of these devices, see **Detailed Configuration Steps** at:

wwwidd.raleigh.ibm.com/netsolut/atm/mm/steps/detailedsteps.html

Network Management

We tested two network management environments in this solution: IBM AIX and Microsoft Windows NT. The AIX management suite included Tivoli NetView and IBM Nways Campus Manager for AIX. The Windows NT environment used IBM Nways Workgroup Manager for Windows NT. This section provides feature and installation information for these products.

The network management stations were Token-Ring-attached to an IBM 8270 Nways LAN Switch Model 800 in the network. Figure 2 in the "Solution Design and Configuration" section shows where the network management stations were connected.

Because of the size of typical networks implementing this solution, Windows NT would probably be the preferred choice of management environments. The AIX solution adds rich function to ATM campus management and is included in this section. However, the cost differential and availability of expertise in smaller environments makes Workgroup Manager for Windows NT the more attractive option.

Network Management on Windows NT Platform

Nways Workgroup Manager for Windows NT Version 1.1.3 provided all the network management on Windows NT in our test environment. It became available in December 1998 and added support for the 8265. Version 1.1.2 supports all the other devices in the network, except the 8265.

The following products made up the network management suite on Windows NT:

- Microsoft Windows NT Version 4.0
- Nways Workgroup Manager for Windows NT Version 1.1.3
- IBM DB2® Universal Database Version 5.0

Configuration on Windows NT

Nways Workgroup Manager provides IP topology services and autodiscovers IP devices in the network configuration. It does not provide any configuration management for VLANs.

Installing Nways Workgroup Manager for Windows NT

We installed Version 1.1.3 from the CD-ROM and did not encounter any problems. Follow the instructions in the README file for setting up DB2 with the Nways Workgroup Manager to store performance data for viewing and reporting.

Tip: You will be given an opportunity to configure your autodiscovery options during the installation. Set the autodiscovery IP address to the beginning address for the discovery process and the subnet mask appropriately. An error here might cause the autodiscovery process to take an excessive amount of time or to hang.

Setting Up Nways Workgroup Manager for Windows NT

Figure 7 is a topology view of the test network. We created it by dragging subsystems from the autodiscovery view, adding subnets (with the Edit -> Create menu option) and creating links by clicking on and connecting icons.

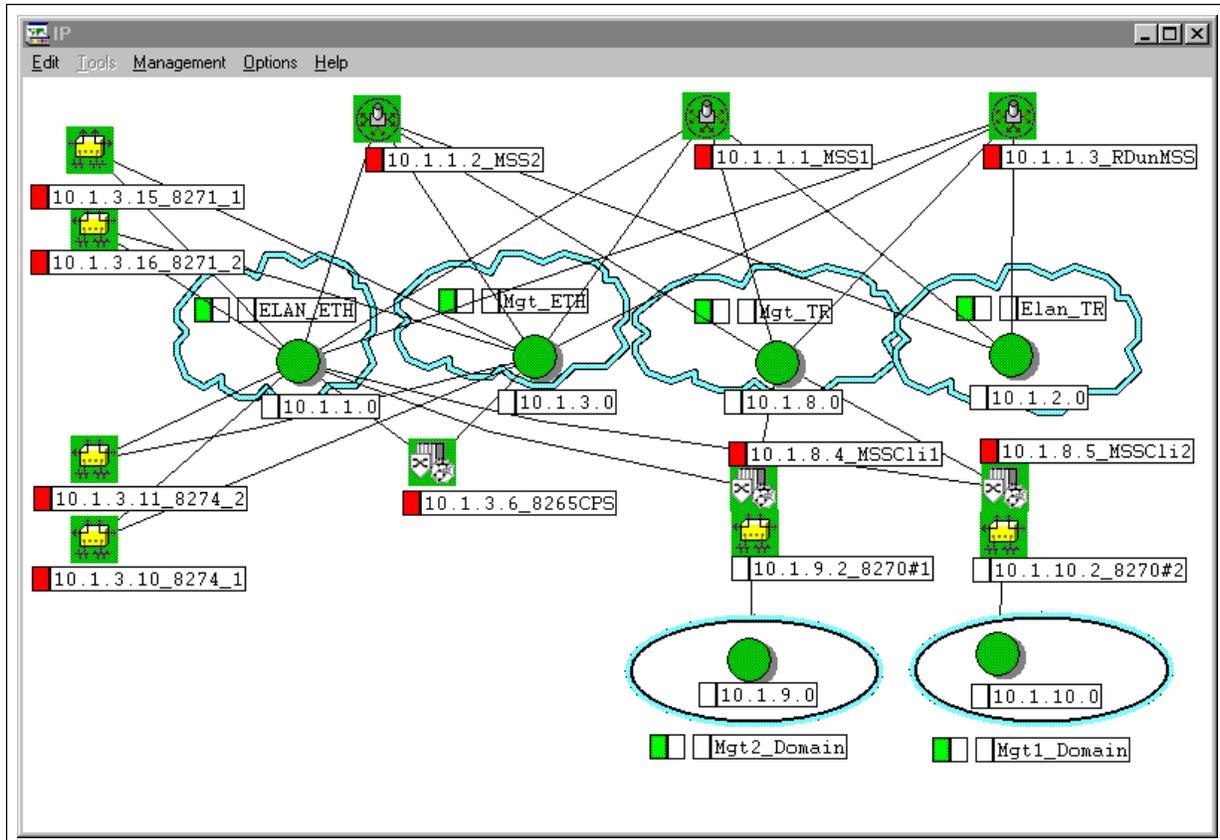


Figure 7. Topology View

Topology views enable you to monitor an entire network or a portion of a network and provide an entry point to the device management applications such as the Java™ Management Applications. Icons represent the devices and their status in the network. For example, a green icon indicates that a device is in a normal state while a red icon shows that the device is in a critical state. A flashing icon indicates a resource status change caused by a poll for status or a device trap. Note that the topology view shown in Figure 7 does not include workstations. You can add a separate view showing connected workstations if you want to monitor them. Non-SNMP devices (that is, workstations with no SNMP support) are monitored but not managed.

If the Netfinity Manager™ is installed along with the Nways Workgroup Manager, it can be set up to initiate a remote session with a server or workstation in the network. The device must have the Netfinity® agent running. When the setup is complete, you can click on the server or workstation in the Nways Workgroup Manager to initiate the remote session.

Device Management

Device management is provided by IBM's Product-Specific Modules (PSMs) and the Java Management Applications (JMAs).

The Generic Java Management Application manages the 8270, 8271, and 8265 switches. (The RouteSwitch Manager manages the 8274 but was not tested in this solution.) The MSS Java

Management Application provides management for the MSS Server and the MSS Client. Each of these programs displays a realistic backplane view of the managed device and provides access to standard and private MIB information that the device supports.

The Java Management Applications provide some configuration management for the devices in the test network, but the ability to configure them through SNMP varies. We performed most of the configuration using the configuration program or on the devices themselves (through Telnet or console attachment).

You can display a JMA device panel by double-clicking on a device icon in a topology view. Figure 8 shows the JMA device panel for the 8265. Port status is displayed and updated dynamically by the traps. Pointing to the ports with the mouse displays context information about the port. Clicking the right mouse button on a port provides status. The Navigation tree on the left is used to interact with the device and the center area displays the results. This is the basic format for all Java device panels.



Figure 8. Java Management Application Device Panel for the 8265



Figure 10. MSS Client Panel

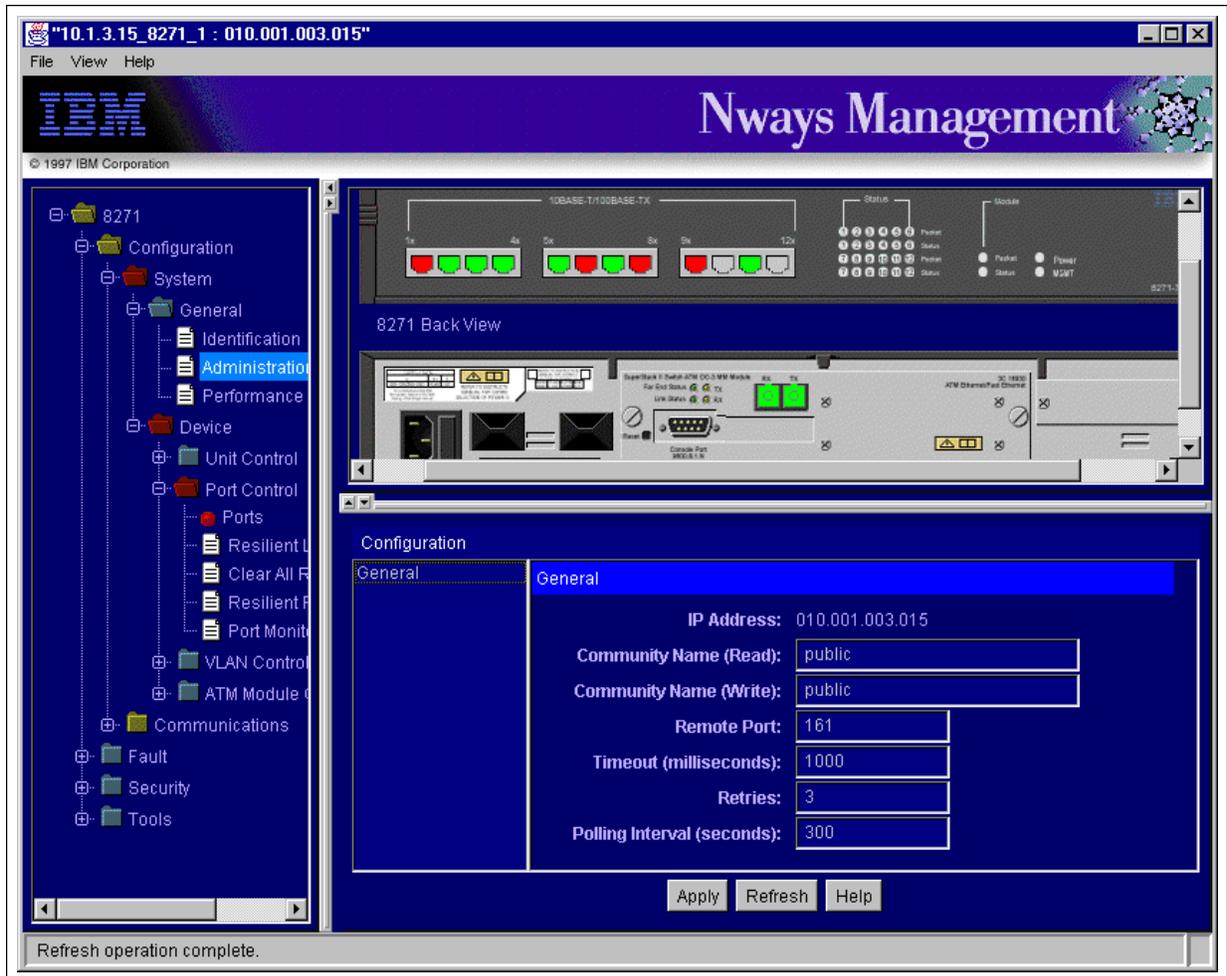


Figure 11. 8271 Device Panel

Performance Management on Windows NT

The Java Management Applications of Nways Workgroup Manager provided all performance management.

We used the Java Performance Manager (JPM) feature of Nways Workgroup Manager to gather and display historical performance data (such as interface traffic, interface utilization and memory usage) for the devices. The reporting feature of JPM prepares historical data reports that are available remotely with a Web browser. JPM is an integrated function within the Nways Java Management Applications. It collects basic performance information about interfaces and protocols by default using a 20-minute polling interval when a JMA is launched for a particular device. You can customize JPM by selecting different MIB objects to be polled, changing the polling frequency, or changing the grouping of objects on graphs.

Figure 12 shows a JMA performance panel for the MSS gateway. It shows a performance graph of data retrieved from the DB2 database. Green circle icons located throughout the navigation tree represent performance data options and produce charts or graphs.

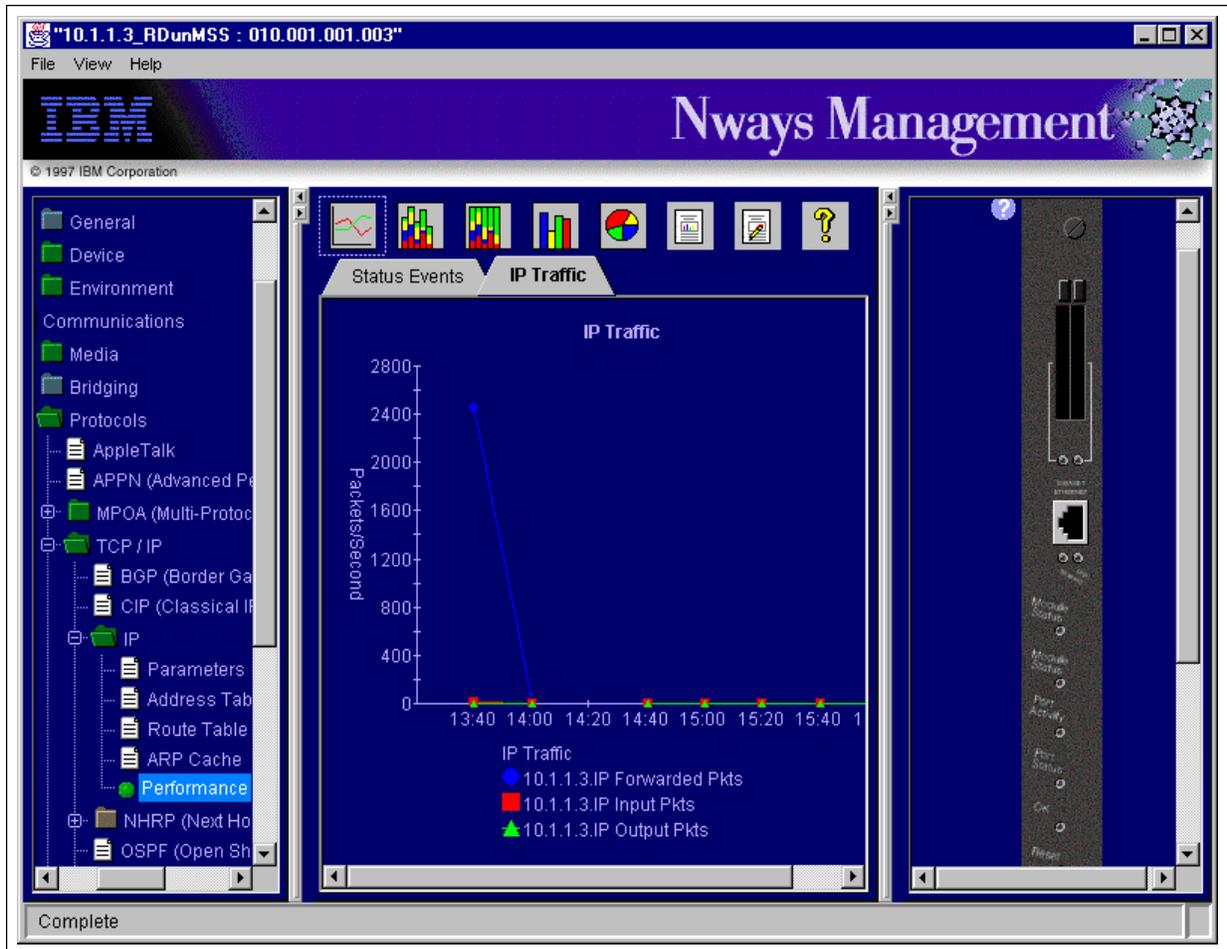


Figure 12. Performance Graph for the MSS Gateway

The **Windows NT Peer Web Services** provided the Web server function on the management station and provided access to JMAs and JPM reports to Web browsers on other workstations. The Java Performance Manager stored collected data in the **DB2 Universal Database**, although you can use other JDBC-compliant databases. Additional information about using the Java Performance Manager can be found in "A Quick Guide to Java Performance Manager" under "Documentation" at:

www.networking.ibm.com/cma/cmasolut.html

The 8270 and 8271 support the RMON standard MIBs, which provide a wide variety of performance statistics. The IBM Nways Workgroup Manager ReMon product can display information from these MIBs in graphical and tabular formats. This product was not tested as part of this configuration.

Besides Nways Workgroup Manager, a **World Wide Web Interface** is provided with the MSS Server, MSS Client, and the 8265 to enable you to configure and monitor these products with a Web browser on any operating system platform. To access the home page for a device, point your browser to the URL:

http://<IP-address of the device>

You might be asked to enter a userid and password, if it is required to configure the box. Figures 13 and 14 show home pages for the 8265 and the MSS Server respectively.

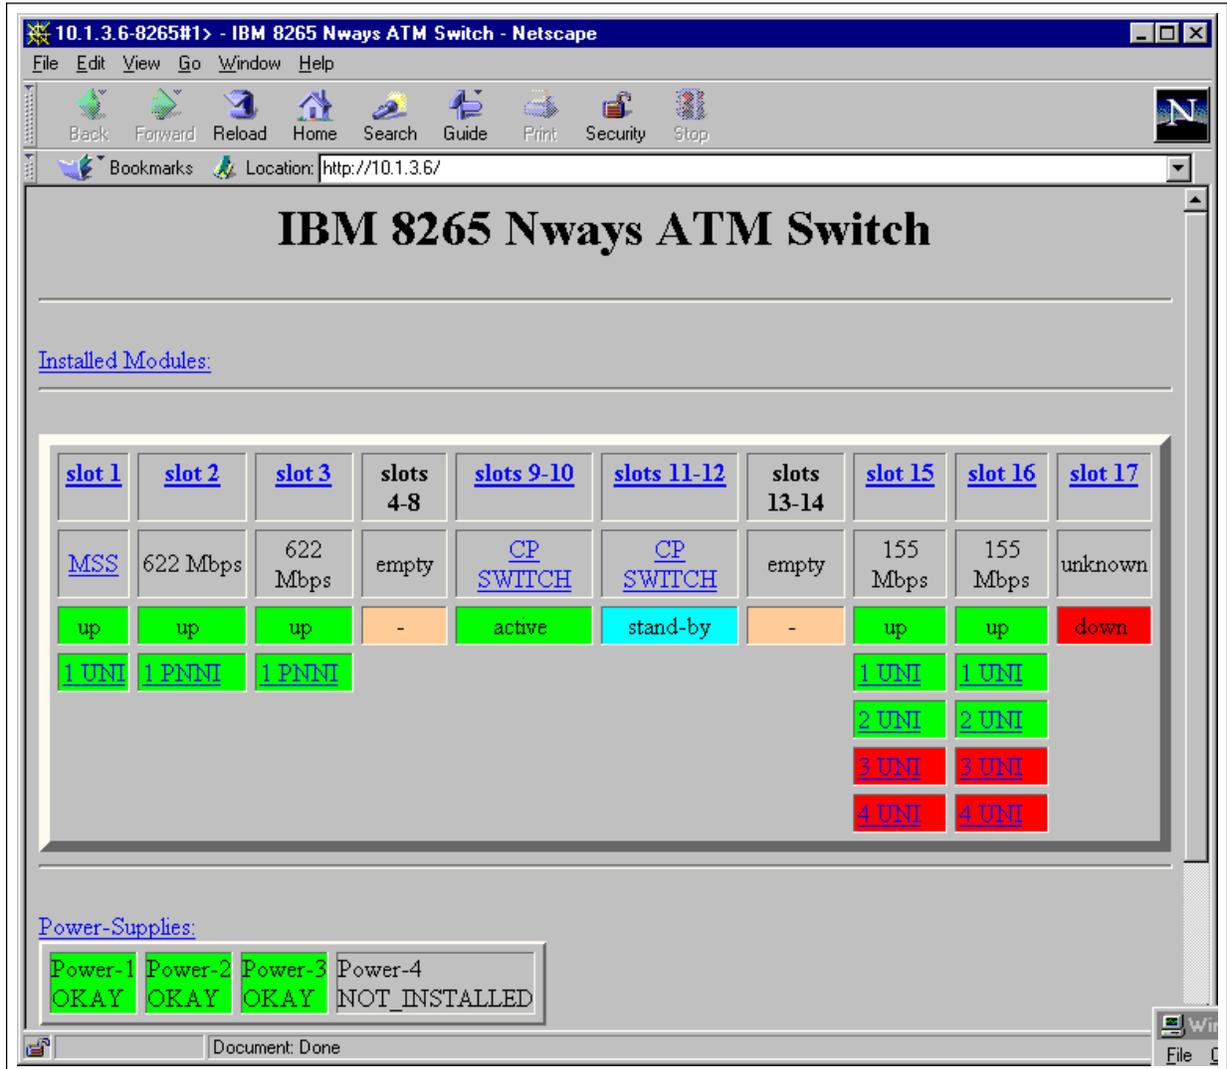


Figure 13. 8265 Home Page

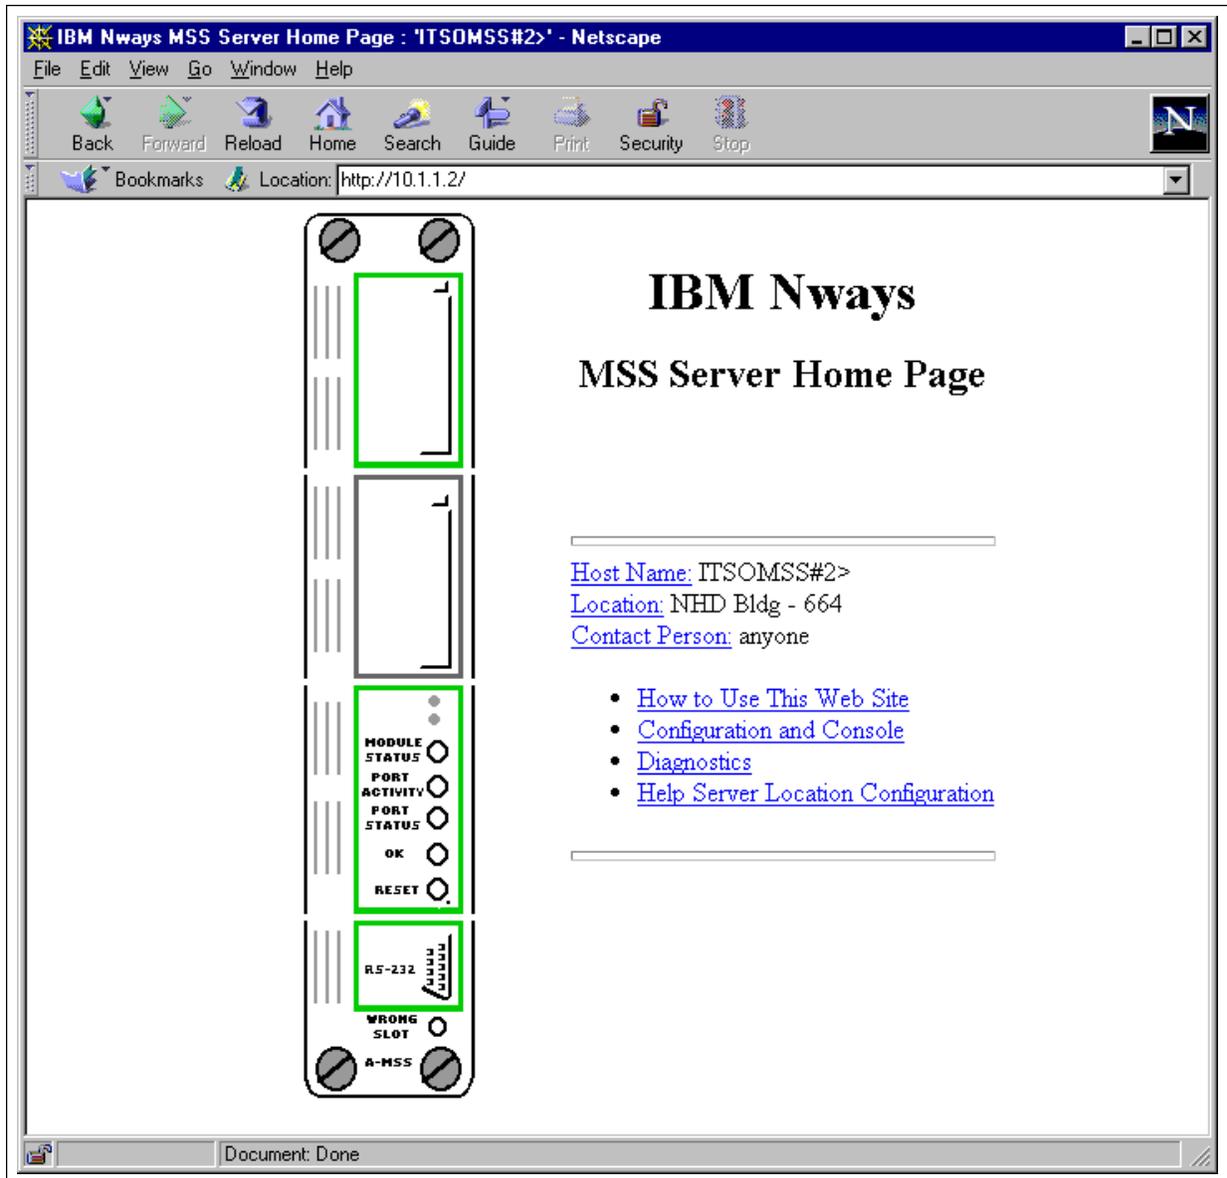


Figure 14. MSS Server Home Page

Detailed Network Management Installation Information for AIX

Installing NetView for AIX

Before starting the NetView install:

1. Install the following AIX file sets:

- X11.compat.fnt.pc -- available on AIX install media
- bos.compat.links -- available on AIX install media
- bos.loc.pc_compat.En_US -- available from:
ftp://aix.boulder.ibm.com/aix/fixes/v4/os

2. Consider the appropriate application file system sizes.

The following table lists the applications and the approximate amount of space they use in the /usr file system. These sizes will vary depending on the options you select when installing the applications.

Application	Approximate Size	Directory
ObjectStore	60408 KB	(/usr/lpp/ODI)
DB2	62808 KB	(/usr/lpp/db2_05_00)
Tivoli	77848 KB	(/usr/local/Tivoli)
NetView	160228 KB	(/usr/OV)
DynaText/Books	69788 KB	(/usr/ebt)
Nways Manager	300168 KB	(/usr/CML)

You might want to create additional file systems so that these applications are not all sharing the same file system.

Tip: DB2 stores its historical data in the /home file system. The DB2 database can grow to more than 200 MB in a few weeks, depending on the amount and frequency of traffic collected. To prevent the /home file system from filling up, it is recommended that you monitor the size of the files under the home directory of the DB2 instance user that owns the IBMNMPDB database. You can purge data from the IBMNMPDB database using the database maintenance tool.

The Java Management Applications log files are in the /log directory, which will usually be in the /(root) file system. These files take a maximum of 5 MB.

3. Consider paging space. Running NetView, Nways Campus Manager and DB2 simultaneously will require at least 256 MB of paging space.

4. After installing NetView, be sure to source /etc/Tivoli/setup_env.sh in .profile

Installing Nways Campus Manager

See the *IBM Nways Manager for AIX Version 1.2.1 Installation Guide* for information about installing this product. Here are some other relevant tips:

Do not install Router and Bridge Manager using the Java front end. It tries to install both the client and the server image, which results in certain files being owned by more than one LPP. If you want RABM, install rabm.obj separately using SMIT.

Remember to run /usr/opt/ifor/ls/conf/i4cfg after completing the install. The Nways Campus Manager books are not put in the .ebtrc file that DynaText uses, if you installed DynaText with NetView. Copy the COLLECTION lines from /usr/ebt/data/.ebtrc to /usr/local/Tivoli/bin/dyantext/data/.ebtrc.

You must modify the /usr/CML/JMA/bin/dpadmin script (as noted in Installation Guide) to run on AIX.

If you want to run a Java Management Application remotely in a Web browser, the documentation instructs you to copy the file ClientClasses.jar to a directory on the client (Web browser) workstation to improve performance. This file (which is a "zipped" collection of class files) is missing several class files. You must either download a new version of ClientClasses.jar (available at www.networking.ibm.com/support/) or zip the following files into ClientClasses.jar:

- ibm.nways.ras.Traces
- ibm.nways.ras.ErrorLogServerHandle
- ibm.nways.ras.ErrorLogServerImpl_stub
- ibm.nways.jdm.modelgen.TableFilter

This step is recommended only for experienced users of zip or WinZip.

Network Management on the AIX Platform

The following IBM products made up the network management suite on AIX:

- AIX Version 4.2.1
- Tivoli NetView Version 5.0
- Nways Campus Manager Version 1.2 (LAN and ATM components)
- DB2 Universal Database Version 5.0
- Internet Connection Secure Server Version 4.2.1 for AIX

The features of each product used in this solution will be discussed in terms of the network management disciplines of Configuration, Performance, and Fault-Monitoring. Detailed information regarding the installation and configuration of the network management products can be found in the "Detailed Network Management Installation Information for AIX" section.

Configuration on AIX

Tivoli NetView provides configuration management from a TCP/IP network perspective. The autodiscovery mechanism of NetView (netmon daemon) will discover most devices and workstations in the network when you use an autodiscovery "seed" file containing the IP address of the MSS. Workstations that are not actively sending traffic might not be discovered.

Tip: Sending a ping from the management station to workstations that are not actively sending traffic causes netmon to discover them and place them on the IP submap.

The VLAN application of **Nways Campus Manager** provides configuration management of the network from a Emulated LAN (ELAN) perspective. This application can be used to monitor the LAN Emulation Configuration Server (LECS) and Broadcast and Unknown Server (BUS) functions in the MSS Server and LAN Emulation Clients (LECs) participating in the ELAN. It can also be used to create and modify the ELAN configuration. Details about using the VLAN application can be found in the IBM *ELAN Management Using Nways Campus Manager for AIX* Redbook at:

publib.boulder.ibm.com/pubs/pdfs/redbooks/sg244821.pdf

Tip: To use the VLAN application to fully configure ELANs managed by the MSS Server, you must configure the MSS Server with at least two community names. Configure one community name (such as “public”) with read-only access. Configure a second community name with read/write access, and use this community name to access the MSS Server from Nways Campus Manager.

Figure 15 is a VLAN application panel showing a summary view of an ELAN.



Figure 15. Summary View of an ELAN

Figure 16 shows a detailed view of an ELAN.

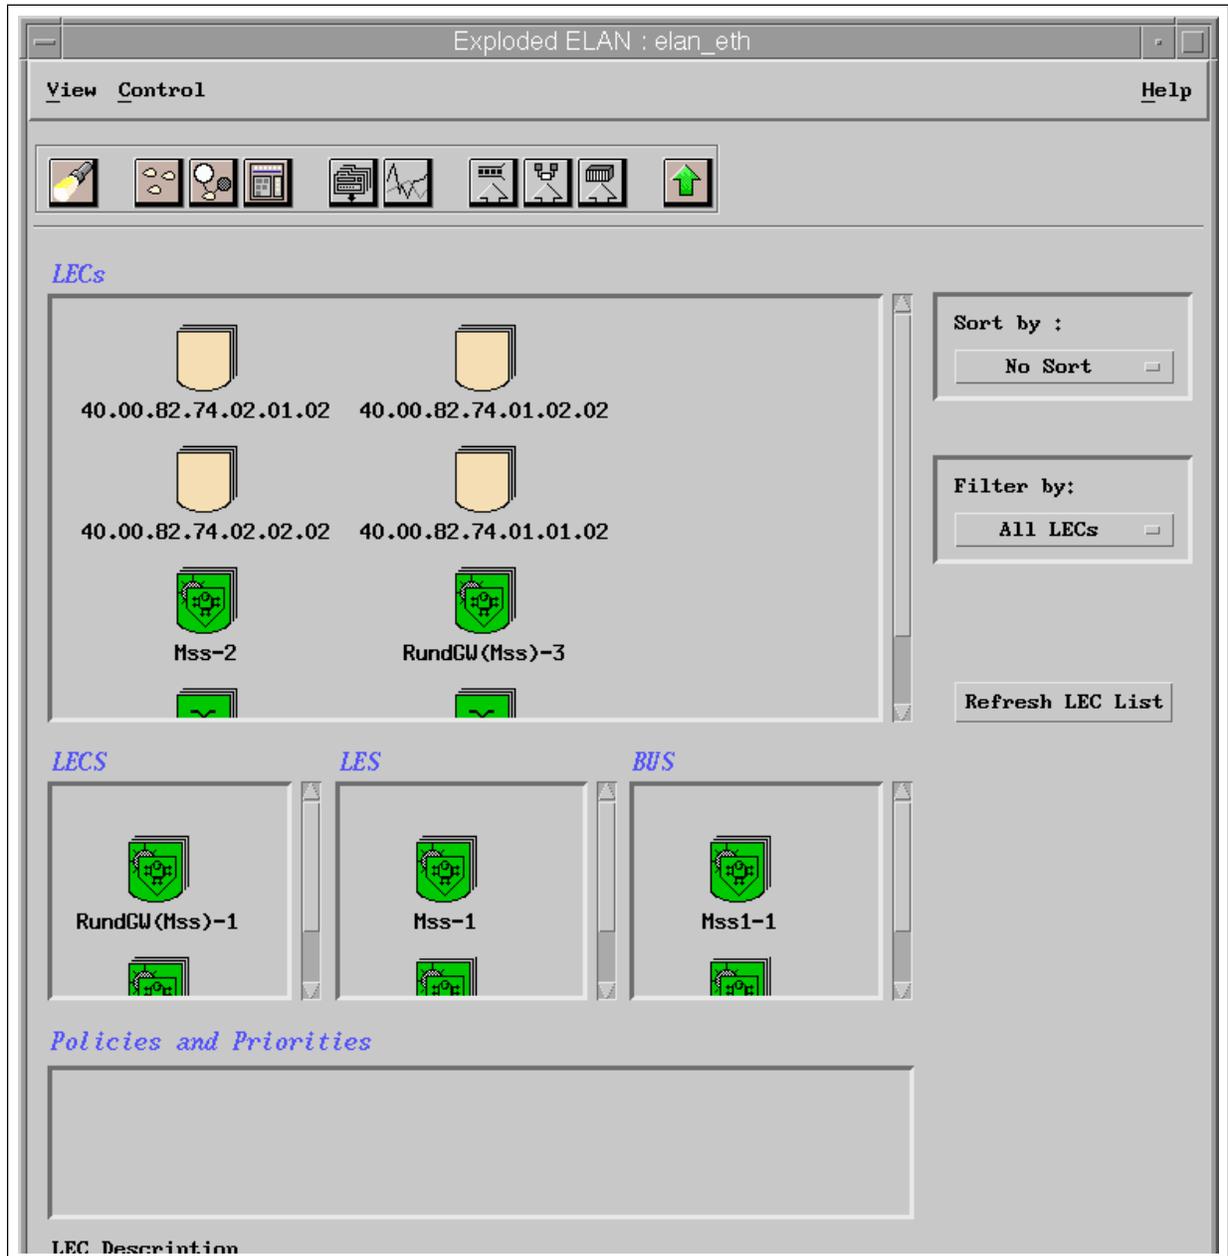


Figure 16. Detailed View of an ELAN

The ATM application of Nways Campus Manager provides management of the ATM physical and logical topologies. The functions of the ATM application include:

- Monitoring the status and configuration of ATM devices and interfaces
- Displaying the PNNI configuration of ATM devices and interfaces
- Monitoring the status of ATM connections
- Displaying and configuring switched virtual circuits (SVCs) and private virtual circuits (PVCs)
- Displaying virtual connections and calls logged over virtual connections

Figure 17 is an ATM application panel showing the ATM view of an 8265 switch. Each of the five-sided icons in the figure represents an interface on the switch; each diamond-shaped icon represents the device that is attached to that interface.

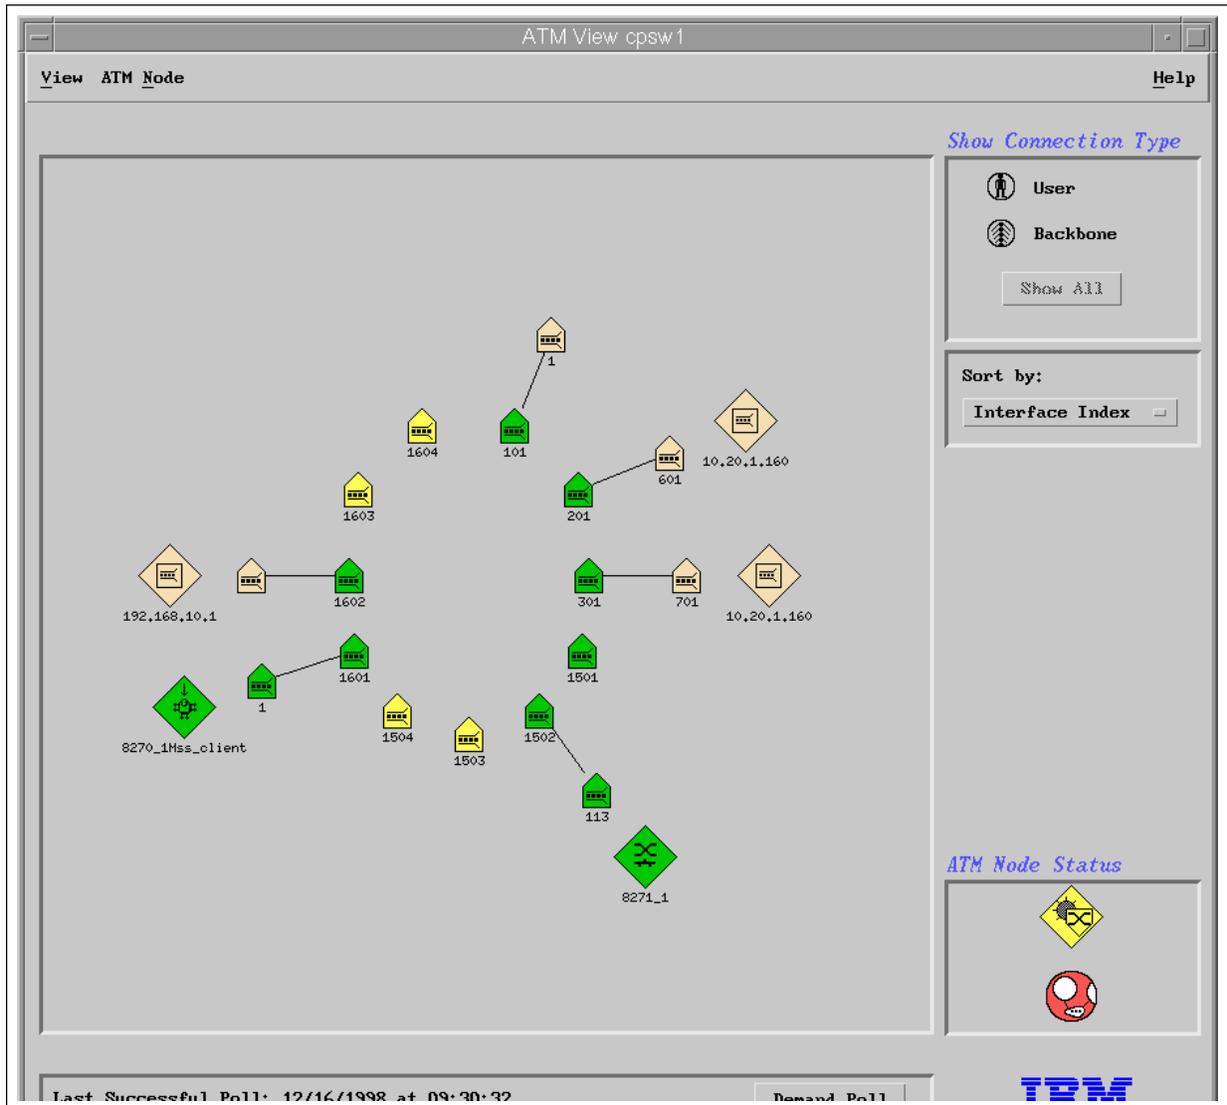


Figure 17. ATM View

Figure 18 shows the connection tracking view of a multicast switched virtual circuit (SVC). Using this feature, you can trace a virtual circuit from end-to-end through the ATM network.

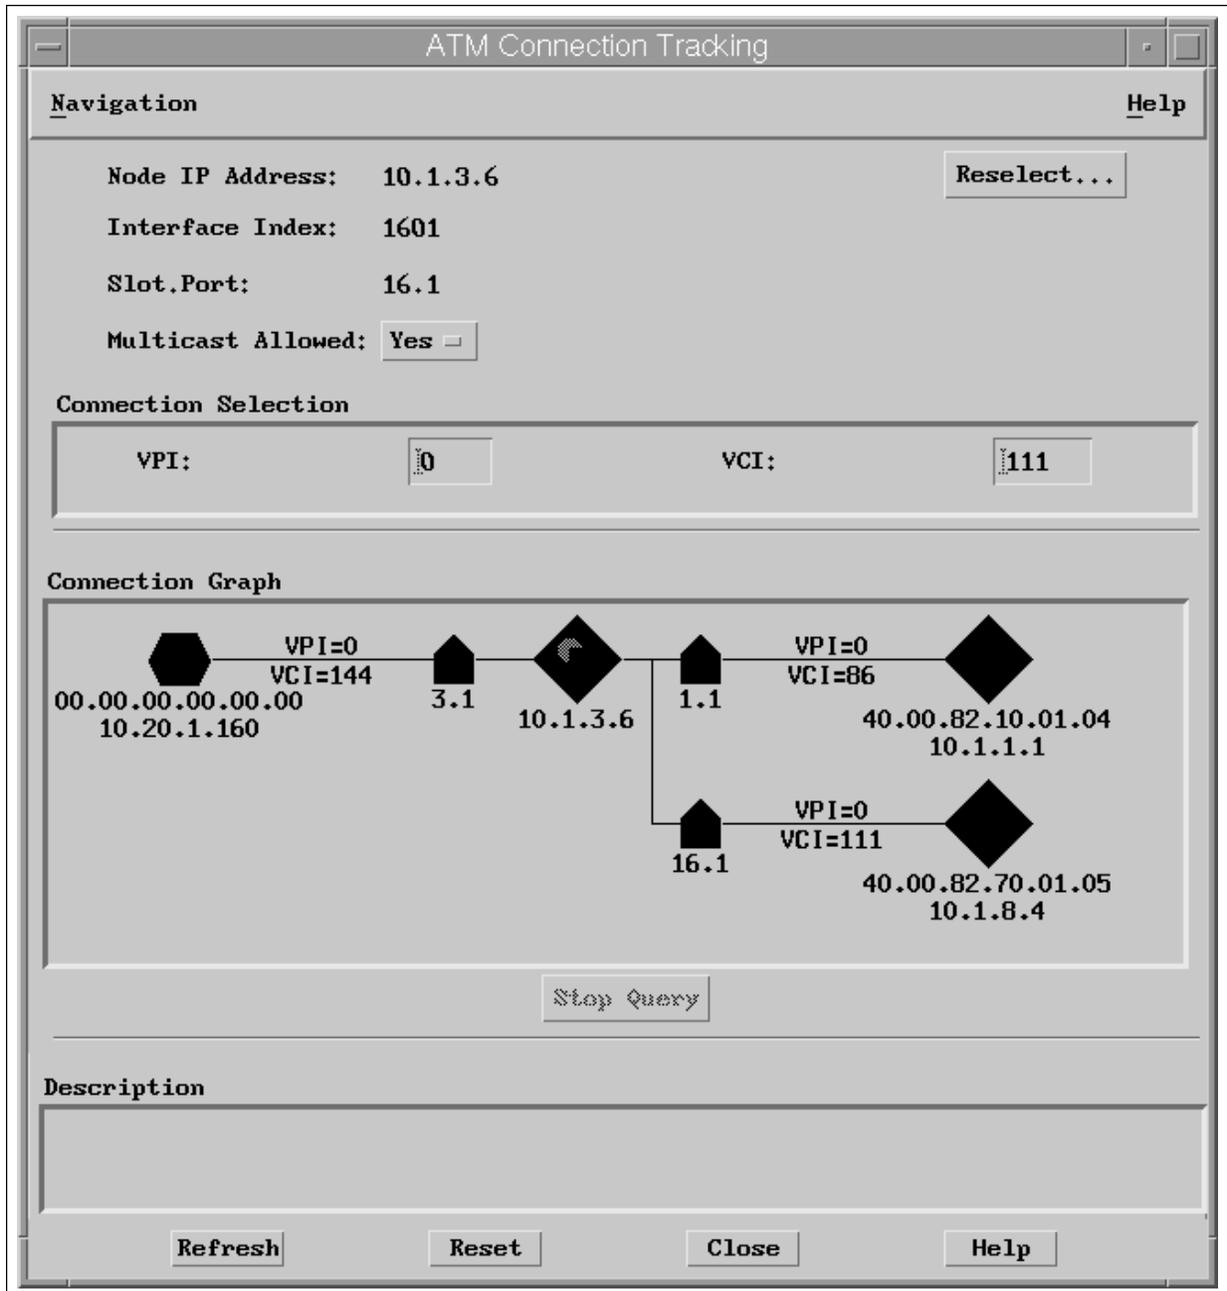


Figure 18. Connection Tracking View of a Multicast SVC

The device management applications of Nways Campus Manager provide some configuration capability for individual devices in the network (see the "Performance Management on AIX" section). The ability to configure these devices through SNMP varies, however. Most of the configuration for this test was performed using the configuration program for the MSS Server and Client or directly on the devices (through Telnet or console attachment).

Performance Management on AIX

Because these applications are written in Java, identical function is available on Windows AIX as on Windows NT. As on Windows NT, **DB2 Universal Database Version 5.0** stored historical performance data.

The device management applications of the Nways Campus Manager LAN component of Nways Manager for AIX provide basic health management (primarily status monitoring) of the specific devices in the network. The Hub Manager function of Nways Campus Manager LAN component manages the 8265 switch. The Generic Java Management Application manages the 8270 and 8271 switches. The MSS Java Management Application provides management for the MSS blade and the MSS Client. This function is identical to the Nways Workgroup Manager Platform. Please refer to that section for details.

For the 8270 and 8271 switches, the backplane picture is static. For the other devices, the backplane picture is dynamic, with a color-coded diagram to represent the status of the various ports on the device. A color-coded SNMP navigation tree indicates areas of failure as detected by the device's SNMP agent.

The Java Performance Manager (JPM) feature of Nways Campus Manager was used to gather and display historical performance data (such as interface traffic, interface utilization and memory usage) for the devices in the network. As in the Nways Workgroup Manager for Windows NT, the reporting feature of Java Performance Manager prepares historical data reports that are available remotely with a Web browser. JPM is an integrated function within the Nways Java Management Applications. JPM collects basic performance information about interfaces and protocols by default using a 20-minute polling interval when a JMA is launched for a particular device. You can customize JPM by selecting different MIB objects to be polled, changing the polling frequency, or changing object groups on graphs.

The **IBM Internet Connection Secure Server** provided the Web server function on the management station and provided access to JMAs and JPM reports to Web browsers on other workstations. The Java Performance Manager stored collected data in the **DB2 Universal Database**, although you can use other JDBC-compliant databases. Additional information about using the Java Performance Manager can be found in "Quick Guide to Java Performance Manager" under "Documentation" at:

www.networking.ibm.com/cma/cmasolut.html

The 8270 and 8271 support the RMON standard MIBs, which provide a wide variety of performance statistics. The IBM Nways Manager ReMon product can display information from these MIBs in graphical and tabular formats. This product was not tested as part of this configuration.

Fault Monitoring on AIX

The Events program of NetView monitors faults associated with the networking devices in the network.

For NetView to correctly format certain traps from the devices in this configuration, run the following shell scripts, which are located in the directory /usr/CML/bin:

```
8271.addtraps  
mss.addtraps
```

Tip: Nways Workgroup Manager for Windows NT does not automatically display traps received from devices that are managed by the Generic Java Management Application (the 8270 and 8271 switches in this configuration). To display traps received from these devices, select the device on the IP topology, and select **Management -> Performance -> Trap Management** from the menu bar. From the Trap Management window that is displayed, configure which traps you wish to receive and save this configuration to a file. You can use this saved configuration for other instances of the same device.

In this configuration, we configured the Trap Management window for 8271 switches to display the traps listed in the following table.

Trap Name	Enterprise	Specific Trap ID	Description
resResilienceSwitch	a3com (1.3.6.1.4.1.43)	43	Generated when a change of state of one of the ports in a resilient pair results in a switch of active port.
resStateChange	a3com (1.3.6.1.4.1.43)	44	Generated when a change of state of one of the ports in a resilient pair does not result in a switch of active port.
localManagementUpdate	a3com (1.3.6.1.4.1.43)	14	Indicates that the configuration of the device has been modified via the ASCII agent.

Fault-Tolerance Testing

This network environment was designed to be fault tolerant with no single points of failure. This section describes the redundancies that we tested for 8270, 8271, and 8274 switches, MSS Servers, MSS Clients, and 8265 switches.

Also included in this section are the test results and recovery times when we deliberately created the types of failures for which the redundant devices and links are designed. The tests were conducted using the following simple scenarios between multiple Windows 95 and Windows NT workstations:

- Performing multiple pings in different directions with IP
- Transferring files with IPX clients and servers (Netware V4.1)
- Having 3270 sessions active with SNA host

8270 Token-Ring Switch Failures

At the edge of the network, the 8270 Token-Ring Switches were configured such that if one switch fails, the other switch will become the active path to the network. We also defined a Token-Pipe connection between the two 8270 switches which provides a redundant link between the switches in the event of an ATM uplink failure.

Test Results

◆ TCP/IP

ATM uplink: When we removed the ATM uplink, the workstations experienced an outage and they all recovered within 2 minutes 30 seconds to 4 minutes 30 seconds. The longer recovery times tended to be for Windows NT workstations ([see Tip below](#)). Reinserting the ATM uplink was non-disruptive and produced no outages on the network.

Tip: The longer times for Windows NT workstations are due to the ARP cache timeout value configured in the workstations. A workstation must re-ARP to find the new path through the network. Shorter recovery times, can be achieved by reducing the ARP cache timeout value in the workstation. Windows NT 4.0 does not allow the end user to change the ARP cache timeout values. However, we were able to use the fix that is available to allow this capability.

TokenPipe: When we disengaged the TokenPipe connections, active TCP/IP connections terminated between workstations attached to either 8270. This caused outages of between 5 to 10 minutes depending on the ARP cache settings. Connectivity can be reestablished immediately by manually clearing the workstation's ARP cache and reconnecting via the new path through the network.

◆ SNA

The 3172 Host connection in our network was directly attached to 8270#1. Therefore both Token-Ring and Ethernet workstations connected to the host through Token-Ring.

ATM uplink: When we removed the ATM uplink, workstations attached to the 8270 switches experienced no SNA session interrupts. Ethernet attached workstations experienced session outages of around 3 minutes at which point a manual disconnect and reconnect of the 3270 session restored connectivity. Reinserting the ATM uplink was non-disruptive to any sessions.

TokenPipe: Removing the TokenPipe connection resulted in SNA session outages only for those workstations attached to the 8270 switch #2. A manual disconnect and reconnect of the 3270 session restored connectivity.

Redundant Power Supply: The 8270 Model 800s that we used in our test environment have redundant power supplies installed. When we disengaged the active power supply, the redundant power supply became active and no outage occurred on the switch.

◆ IPX

ATM uplink: When we removed the ATM uplink, workstations attached to the 8270 switches generally recovered within 30 to 45 seconds. After reinserting the ATM uplink, workstations recovered within 10 seconds.

TokenPipe: Removing the TokenPipe connection did not affect the IPX traffic as all traffic was going over the ATM interfaces.

Redundant Power Supply: The 8270 Model 800s that we used in our test environment have redundant power supplies installed. When we disengaged the active power supply, the redundant power supply became active and no outage occurred on the switch.

When a failing switch recovers from an outage, existing traffic on the network is not affected and experiences no outage upon the switch's recovery. In general, in a Token-Ring environment, it is the responsibility of the end stations to find the path through the network. The recovery times may also vary depending on the operating system and protocol.

8274 LAN Switch Failures

At the edge of the network, the 8274 switches were configured with redundant ATM uplinks. If one of the fibers breaks or the ATM port fails the redundant ATM uplink will take over and reroute traffic through the switch. The Spanning Tree function configured on the switch manages the path selection according to the configured path cost values. The 8274 switches were also equipped with a redundant Management Processor Module (MPM) and a redundant power supply.

Test Results

◆ TCP/IP

Redundant ATM uplink: When we removed the ATM uplink from the switch, workstations experienced an outage of 30 to 50 seconds, which correlates with the default Spanning Tree recalculation time. Reinserting the uplink resulted in another 30-second outage before all the workstations recovered.

Redundant MPM: When we removed the primary MPM, an outage of from 1 minute 30 seconds to 2 minutes resulted before the backup MPM took over and all the workstations reestablished connectivity. Reinserting the MPM is non-disruptive as it becomes the redundant MPM.

◆ SNA

Redundant ATM uplink: When we removed the ATM uplink from the switch, 8274-attached workstations experienced session hangs of around 30 seconds. When we reinserting the uplink, the workstations experienced another 30-second session hang before they recovered.

Redundant MPM: When we removed the primary MPM, 8274-attached workstations experienced an outage after 1 minute. A manual disconnect and reconnect of the 3270 session restored connectivity. Reinserting the MPM is non-disruptive because it becomes the redundant MPM.

Redundant Power Supply: The 8274 Models W53 and W93 that we used in our test environment have redundant power supplies installed. Removing power from the active power supply invokes the redundant power supply and produces no outage on the switch.

◆ IPX

Redundant ATM uplink: When we removed the ATM uplink from the switch, 8274-attached workstations recovered within 40 seconds. When we reinserting the uplink, the workstations again recovered within 40 seconds.

Redundant MPM: When we removed the primary MPM, 8274-attached workstations experienced an outage of less than 1 minute and were then able to reattach to their server. Reinserting the MPM is non-disruptive because it becomes the redundant MPM.

Redundant Power Supply: The 8274 Models W53 and W93 that we used in our test environment have redundant power supplies installed. Removing power from the active power supply invokes the redundant power supply and produces no outages on the switch.

8271 LAN Switch Failures

At the edge of the network, the two 8271 Model 712s support resilient links. If a primary link fails, the resilient link will become active, as shown in Figure 2 in the "Solution Design and Configuration" section. The 8271 Ethernet Switches in our network had Fast Ethernet connections set up between them acting as resilient links, should the ATM uplinks fail. The status of an ATM connection is based on light passing through the fiber seen from the ATM switch. If a fiber breaks or a switch port fails, the 8271 uses the resilient link. If the ATM connection recovers and no loss of link is detected for 4 to 5 minutes, the switch disables the resilient link and begins using the ATM connection again. Note that a resilient link is not used if the port on the 8265 ATM switch is disabled, because the light on the fiber going to the 8271 is not lost. For more details about how resilient links work with the 8271, see the *8271 Nways Ethernet LAN Switch User's Guides* at:

www.networking.ibm.com/did/8271bks.html.

Tip: As explained in the preceding paragraph, an 8271 with an ATM module will use the configured resilient link when a loss of light occurs on the fiber to the ATM switch. This recovery does not occur if the switch loses contact with the LES/BUS. Therefore, use redundant MSS Servers and redundant paths throughout the ATM network to handle that possible failure.

Test Results

◆ TCP/IP

ATM uplink: When we removed the ATM uplink and dynamically switched to the resilient link, active IP sessions experienced an outage of 2 to 3 seconds. When the ATM uplink was reinserted there was no immediate effect upon active IP sessions. It took around 5 minutes before the switch disabled the resilient link and returned to the ATM uplink. At this point, the active IP sessions experienced slightly longer outages of 20 to 30 seconds. Any devices attempting to establish a new IP session to another device might however experience outages in the range of 3 to 5 minutes.

◆ SNA

ATM uplink: When we removed the ATM uplink, workstations recovered within 20 seconds. When the ATM uplink was reinserted there was no immediate effect upon any sessions. It took around 3 minutes before the switch disabled the resilient link and returned to the ATM uplink. At this point, workstations recovered within 1 minute. Some workstations lost their connection to the server but manual reconnect restored connectivity.

◆ IPX

ATM uplink: When we removed the ATM uplink, dynamic switching to the resilient link was non-disruptive. When the ATM uplink was reinserted there was no immediate effect upon any sessions. It took around 3 minutes before the switch disabled the resilient link and returned to the ATM uplink. At this point, workstations attached to the recovering 8271 switch experienced a session outage. A manual disconnect and reconnect of the 3270 session restored connectivity.

MSS Server Failures

The most critical redundancy components in the network are multiple MSS Servers. We used two MSS Servers. The first MSS Server (on MSS#1) is the primary LAN Emulation Server (LES) and Broadcast and Unknown Server (BUS) for the Ethernet ELANs (*elan_eth*) and (*management1*). It is also the backup LES/BUS for the Token-Ring ELANs (*elan_tr*) and (*management2*). The second MSS Server (MSS#2) is the primary LAN Emulation Server (LES) and Broadcast and Unknown Server (BUS) for the Token-Ring ELANs and the backup LES/BUS for the Ethernet ELANs.

The MSS Server performs the routing between subnets on the backbone ELANs. The MSS Server also supports a function called redundant default IP gateway. If the MSS Server that is acting as the default gateway fails, the same MAC address and IP address that all of the end stations use for the default gateway become active on the other MSS Server. All other protocols except IPX are bridged throughout the network. For more details about the redundancy features of the MSS Server, see the *MSS Quick Guides* on the **IBM Technical Reports** page at:

www.widd.raleigh.ibm.com/tr2/tr2over.html

Test Results

We tested MSS Server redundancy to determine the recovery time for a LES/BUS failure.

- ◆ **TCP/IP:** When we removed the first MSS Server from the ATM network, some workstations experienced an outage of from 16 seconds to 1 minute 20 seconds. When the MSS Server was reinserted back into the chassis it took around 3 minutes to reboot. When the reboot completed, some workstations experienced another outage of from 7 to 45 seconds before the network environment stabilized.
- ◆ **SNA:** When we removed the first MSS Server from the ATM network, some workstations experienced an outage after 30 seconds. A manual disconnect and reconnect of the 3270 sessions restored connectivity again within 45 seconds.
- ◆ **IPX:** When we removed the first MSS Server from the ATM network, workstations recovered within 10 seconds. When the MSS Server was reinserted into the network, workstations recovered within 10 seconds.

8265 Failures

The IBM 8265 Nways ATM Switches have multiple paths available. Because the switches are using compliant PNNI routing, you can use all of the paths between the switches. If any path fails, the switches route future virtual channel connections (VCCs) around it. You can also configure the 8265 with a redundant Control Point Switch (CPSW) module. This redundancy satisfies a requirement of most networks, no single point of failure. If the active CPSW module fails, the redundant CPSW module takes over. This recovery however is disruptive. It breaks all of the VCCs and resets other modules within the chassis, including the MSS Server. Other additional redundancy features which exist in the 8265 include power supplies and controller modules. For more information about its redundancy capabilities, see the **8265 Nways ATM Switch Product Description** at:

www.networking.ibm.com/did/8265bks.html

Test Results

The components within the 8265 that we tested for redundancy were the PNNI links, CPSW modules and Redundant Controller (RCTL) modules. Our final test was to power off each 8265 and use the other as the backup.

◆ TCP/IP

PNNI Links: When we removed either of the PNNI links, workstations experienced 5-second outages. Reinserting the fiber caused no disruption to the network.

CPSW: When we removed the active CPSW module from the ATM network, some workstations experienced an outage of from 3 to 6 minutes due to the timeout values of the ARP cache in the end stations. Reinserting the CPSW module is non-disruptive because this module will be in redundant mode as the standby CPSW.

RCTL: Removing the RCTL module from the ATM network also causes the ATM subsystem to reset. This resulted in some workstations experiencing an initial outage of from 20 seconds to 1

minute 30 seconds. Another small outage of from 2 to 5 seconds occurs when the MSS reboot completes.

8265s: When we powered off each 8265, workstations experienced outages from 20 seconds to 5 minutes while the other 8265 took over. Most workstations recovered within 1 to 2 minutes of the power-off. Powering the 8265 back on resulted in another outage with recovery times ranging from 5 seconds to 6 minutes before all workstations recovered.

◆ SNA

PNNI Links: When we removed either of the PNNI links, some workstations experienced session outages after 40 to 50 seconds. A manual disconnect and reconnect of the 3270 sessions restored connectivity. Reinserting the fiber caused no disruption to the network.

CPSW : When we removed the active CPSW module from the ATM network, some workstations again experienced session outages after 40 to 50 seconds. A manual disconnect and reconnect of the 3270 sessions will restore connectivity. Reinserting the CPSW module is non-disruptive because this module will be in redundant mode as the standby CPSW.

RCTL: Removing the RCTL module from the ATM network also causes the ATM subsystem to reset. Some workstations experienced an initial outage after around 40 seconds. A manual disconnect and reconnect of the 3270 sessions will restore connectivity.

8265s: Powering off each 8265 produced the similar failure and recovery results as the previous two tests on the CPSW and RCTL modules.

◆ IPX

PNNI Links: When we removed either of the PNNI links, some workstations experienced session outages of less than 40 seconds. Those which lost connection to the server were again able to reconnect. Reinserting the fiber caused no disruption to the network.

CPSW: When we removed the active CPSW module from the ATM network, some workstations again experienced session outages but recovered within 40 seconds. Removal of the primary CPSW causes a reset of the ATM subsystem, which caused a reboot of the MSS Server. When the server returns, traffic recovers within 10 seconds. Reinserting the CPSW module is non-disruptive because this module will be in redundant mode as the standby CPSW.

RCTL: Removing the primary RCTL module from the ATM network also causes the ATM subsystem to reset. Some workstations experienced an initial outage but recovered within 40 seconds. When the MSS server returns, traffic hangs but recovers within 10 seconds.

8265s: Powering off each 8265 produced the similar failure and recovery results as the previous two tests on the CPSW and RCTL modules.

Performance Testing

This section summarizes the performance testing conducted to characterize the throughput capacity of this solution. First, we describe the test configuration and then present the results. Note that the results in your network might vary depending on your traffic configuration and other factors.

Test Configuration

The test configuration is shown in Figure 19. We used SmartBits from NetCom Systems to measure both throughput and CPU utilization of the MSS Server. In this configuration, packets between Ethernet and Token-Ring devices are routed through the MSS Server.

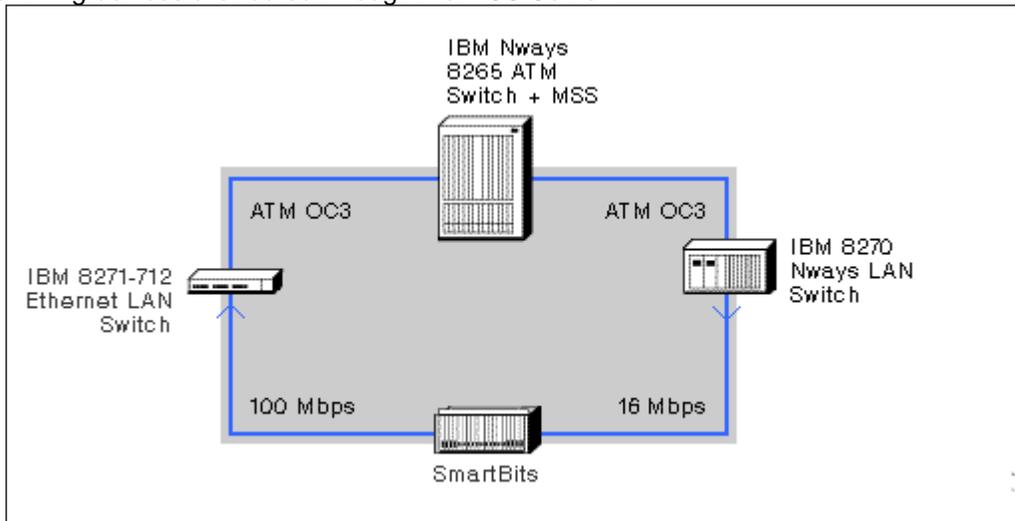


Figure 19: Test Configuration

Test Results

This section provides test results for the throughput capacity and CPU utilization of the MSS Server.

Throughput Capacity

Figure 20 shows the end-to-end throughput capacity for routing IP packets between Ethernet and Token-Ring devices over an ATM backbone. This throughput was achieved using three 100-Mbps Ethernet streams. As expected, as the frame size increases, the throughput approaches the maximum theoretical throughput of the 155-Mbps ATM link (accounting for ATM cell headers). For small (64-byte) frames the throughput is around 43 K packets per second, partly because every 64-byte frame will consume two ATM cells (a total of 106 bytes), which causes 60% efficiency. This throughput was achieved for unidirectional (one-way) traffic.

Note that this throughput was measured using MSS with software level V2R2. Previous software levels experienced lower throughput.

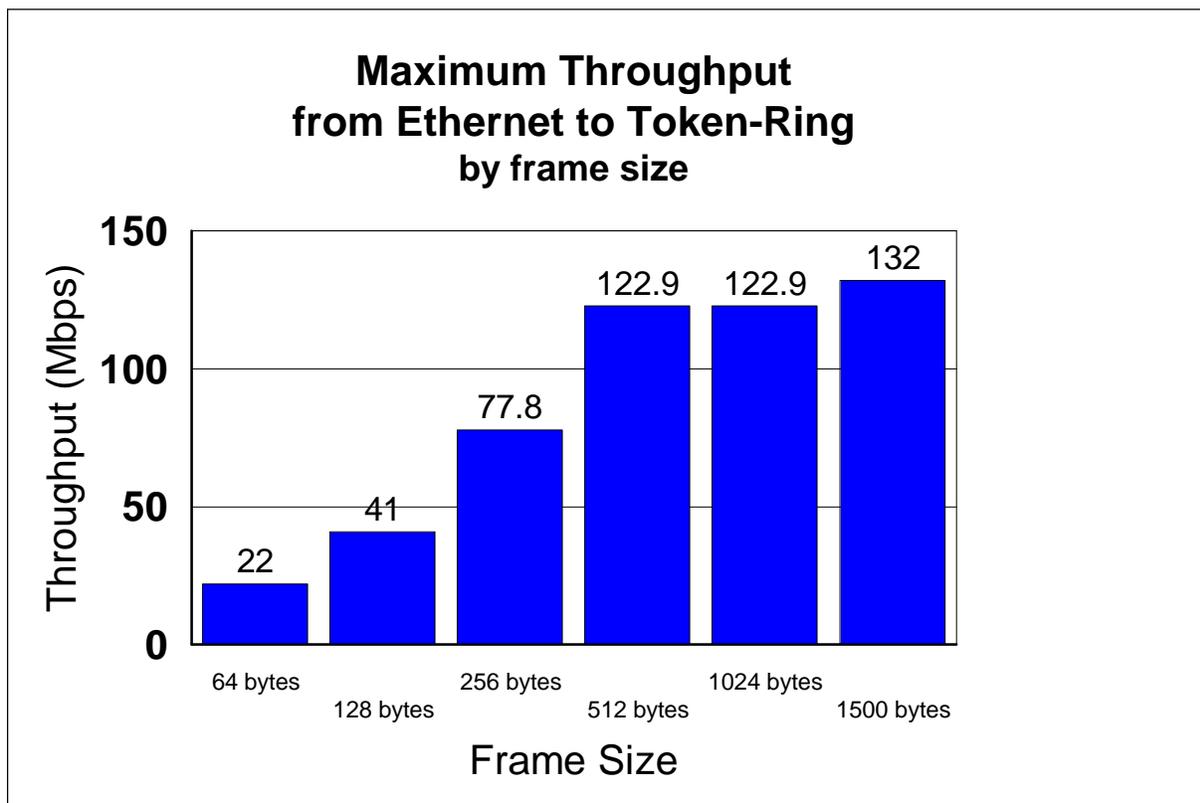


Figure 20: End-to-End Throughput Capacity

MSS CPU Utilization

Figure 21 shows CPU utilization of the MSS Server as a function of routed IP packets. We show utilization for 64-byte and 1500-byte packets, which are the smallest and largest Ethernet frame sizes respectively.

For 64-byte frames, CPU utilization is a linear function of the traffic load in packets per second and it approaches 100% at around 43 K packets per second. For 1500-byte packets, throughput is about 35% for 10 K packets per second and higher. Note that at 10 K packets per second throughput is 120 Mbps, which is very close to the throughput capacity of the ATM uplink.

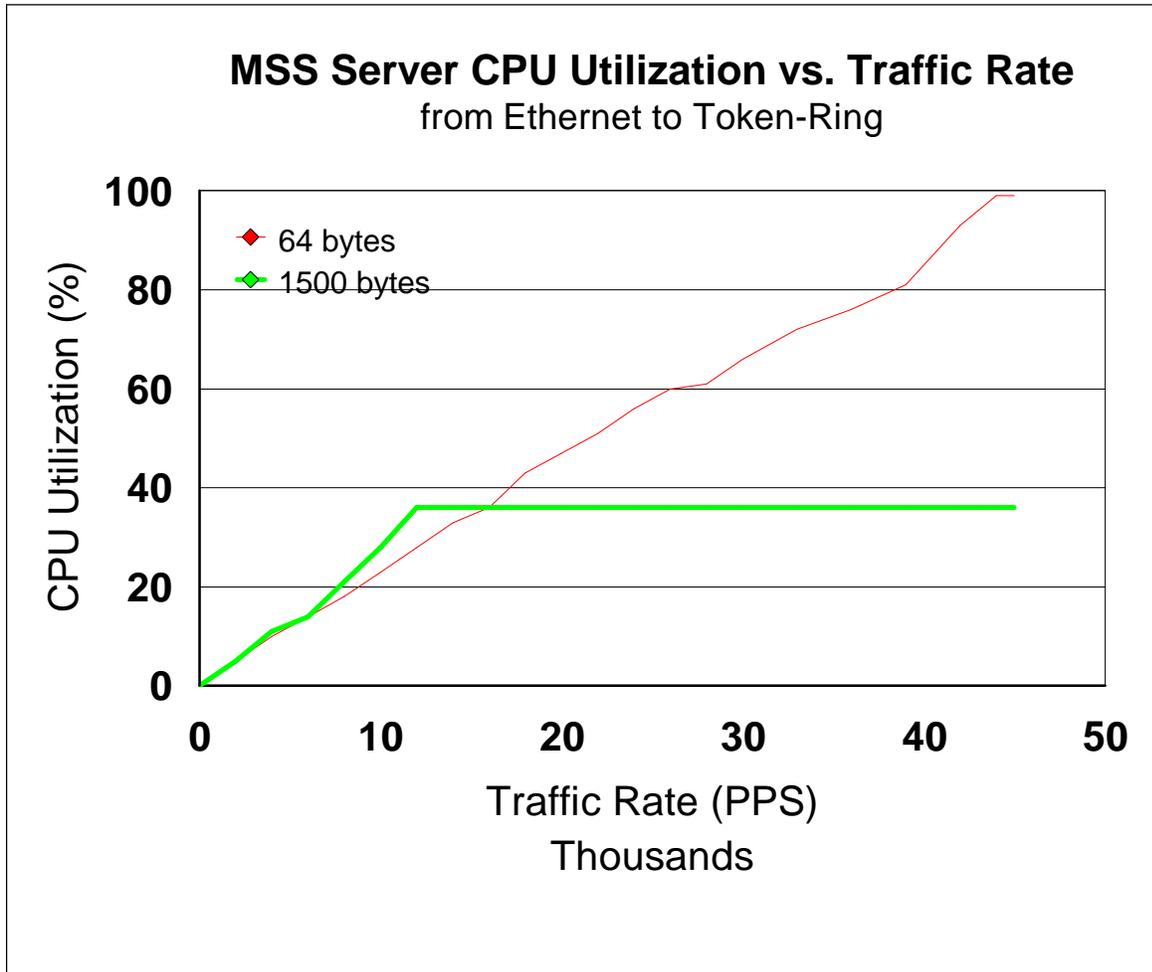


Figure 21. MSS CPU Utilization

Tip: We highly recommended that you upgrade your MSS software to level V2R2 to take advantage of performance enhancements.

Conclusion

Ethernet and Token-Ring switching with an ATM LAN Emulation backbone is a high-performance, scalable and manageable solution. One of its many benefits is high-speed access to servers. It enables you to connect servers with Fast Ethernet and/or 155-Mbps links, and as required move them to Gigabit Ethernet and/or 622-Mbps ATM connections.

We hope that our experience helps you to understand some of the details of the solution and avoid any potential delays or obstacles to a smooth implementation. Overall, we found this to be an easy network to configure and use. Our solution was designed for an easy migration to ATM from existing Ethernet and Token-Ring networks. For multiprotocol (IP/IPX, SNA, NetBIOS) traffic, we used MSS to do the routing and bridging. IP traffic was routed while all other protocols were bridged between the Ethernet and/or Token-Ring switches. We used the MSS Server to route IP and IPX traffic within the backbone network and the 8274 to route IP and IPX traffic in the edge Ethernet network. This distributed routing solution can provide more routing performance when you expand your network by adding more 8274s.

We provided for reliability and fault tolerance in our network. Besides backup for key device components, such as controllers and power supplies, we implemented MSS Server redundancy for reliability of LANE services (LECS and LES/BUS). The redundant default IP gateway function provided redundancy for IP routing. We used Spanning Tree algorithms to enable backup for ATM uplinks in the 8274s and 8270 MSS Clients. When we broke an ATM link, we found that the Spanning Tree function provided good recovery of ATM services on the 8274. Recovery of an ATM connection by an 8271 resilient link was quick for all protocols, although you need to be aware of a possible delay when the ATM link returns. The TokenPipe connection between 8270 switches can provide backup for the ATM links, although the recovery time was a little longer.

We would like to end this document with some general guidelines about implementing this design for your network. A general guideline when designing this type of solution is to keep it simple. Configure multiple Emulated LANs (ELANs) only if there is a valid reason, for example, for security. The IP addressing scheme might also require them. Most networks have multiple IP subnets and a variety of other routed and bridged protocols. As IP becomes more and more predominant, it often dictates the structure of the network.

Definitions

Broadcast and Unknown Server (BUS)

A LAN Emulation Service component responsible for the delivery of multicast and unknown unicast frames.

BUS

See *Broadcast and Unknown Server*.

domain

A logical grouping of machines in a network for the purpose of common management.

ELAN

See *Emulated LAN*.

Emulated LAN (ELAN)

A specific implementation of a virtual LAN, as it relates to LAN Emulation in ATM networks. An ELAN consists of one or more LAN Emulation clients (LECs) that share the same LAN Emulation Server and Broadcast and Unknown Server (LES/BUS). Broadcasts by any member of the ELAN are contained within the boundaries of that ELAN, and ELAN membership can be assigned based on configurable policies.

LAN Emulation Configuration Server (LECS)

A LAN Emulation Service component that centralizes and disseminates configuration data.

LAN Emulation Server (LES)

A LAN Emulation Service component that resolves LAN Destinations to ATM addresses.

LAN segment

- (1) Any portion of a LAN (for example, a single bus or ring) that can operate independently but is connected to other parts of the establishment network via bridges.
- (2) An entire ring or bus network without bridges.

LECS

See *LAN Emulation Configuration Server*.

LES

See *LAN Emulation Server*.

virtual LAN

A logical grouping of hosts, independent of physical location in the network, that defines what stations can communicate to each other. A VLAN can be based on different policies, such as protocol or network address. Each end station joining a VLAN will share a broadcast domain with other end stations in the same VLAN.

VLAN

See *virtual LAN*.

Address Table

This table provides the IP and ATM addresses for the devices in this solution:

Description	IP Address	ATM Address
MSS Servers		
MSS#1 LECS	n/a	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.01.00.00
MSS#1 LECS Interface	n/a	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.01.00.01
LES/BUS for <i>elan_eth</i>	n/a	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.01.00.10
LES/BUS for <i>management1</i>	n/a	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.01.00.11
Backup LES/BUS in MSS#1 for <i>elan_tr</i>	n/a	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.01.00.12
Backup LES/BUS in MSS#1 for <i>management2</i>	n/a	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.01.00.22
MSS#1 LEC for <i>elan_eth</i>	10.1.1.1	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.01.01.01
MSS#1 LEC for <i>management1</i>	10.1.3.1	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.01.03.03
MSS#1 LEC for <i>elan_tr</i>		39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.01.02.02
MSS#1 LEC for <i>management2</i>	10.1.2.1	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.01.04.04
MSS#2 LECS	n/a	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.02.00.00
MSS#2 LECS Interface	n/a	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.02.00.01
LES/BUS for <i>elan_tr</i>	n/a	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.02.00.10
LES/BUS for <i>management2</i>	n/a	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.02.00.11
Backup LES/BUS in MSS#2 for <i>elan_eth</i>	n/a	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.02.00.12
Backup LES/BUS in MSS#2 for <i>management1</i>	n/a	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.02.00.22
MSS#2 LEC for <i>elan_tr</i>	10.1.2.2	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.02.01.01
MSS#2 LEC for <i>management2</i>	10.1.8.2	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.02.03.03
MSS#2 LEC for <i>elan_eth</i>	10.1.1.2	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.02.02.02
MSS#2 LEC for <i>management1</i>	10.1.3.2	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.10.02.04.04
8265s		
8265#1 CPSW LEC for <i>management1</i> ELAN	10.1.3.6	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.65.01.00.00
8265#2 CPSW LEC for <i>management2</i> ELAN	10.1.8.6	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.65.02.00.01

Description	IP Address	ATM Address
8271s		
8271-712#1 LEC for <i>management1</i> ELAN	10.1.3.15	39.99.99.99.99.99.00.00.99.99.01.01.08.00.4E.36.78.79.00
8271-712#2 LEC for <i>management1</i> ELAN	10.1.3.16	39.99.99.99.99.99.00.00.99.99.01.02.08.00.4E.37.4D.41.02
8274s		
8274-W53#1 LEC for <i>elan_eth</i> ELAN	n/a	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.74.01.01.02
8274-W53#1 backup LEC for <i>elan_eth</i> ELAN	n/a	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.74.01.02.02
8274-W53#1 LEC for <i>management1</i> ELAN	10.1.3.10	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.74.01.02.02
8274-W53#1 backup LEC for <i>management1</i> ELAN	10.1.3.10	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.74.01.01.02
8274-W93#1 LEC for <i>elan_eth</i> ELAN	n/a	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.74.02.01.02
8274-W93#1 backup LEC for <i>elan_eth</i> ELAN	n/a	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.74.02.02.02
8274-W93#1 LEC for <i>management1</i> ELAN	10.1.3.11	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.74.02.02.03
8274-W93#1 backup LEC for <i>management1</i> ELAN	10.1.3.11	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.74.02.01.03
8270s		
8270-800#1 LEC for <i>elan_tr</i> ELAN	n/a	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.70.01.04.04
8270-800#2 LEC for <i>elan_tr</i> ELAN	n/a	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.70.02.04.04
8270-800#1 LEC for <i>management2</i> ELAN	10.1.9.1 or 10.1.8.4	39.99.99.99.99.99.00.00.99.99.01.01.40.00.82.70.01.05.05
8270-800#2 LEC for <i>management2</i> ELAN	10.1.10.1 or 10.1.8.5	39.99.99.99.99.99.00.00.99.99.01.02.40.00.82.70.02.05.05
8270-800#1 LAN Switch	10.1.9.2	n/a
8270-800#2 LAN Switch	10.1.10.2	n/a

Related Web Sites

These are the Web sites we referenced in this document.

IBM 8270 Nways LAN Switch documentation	www.networking.ibm.com/did/8270bks.html
IBM 8265 Nways ATM Switch documentation	www.networking.ibm.com/did/8265bks.html
IBM Nways Multiprotocol Switched Services (MSS) Server documentation	www.networking.ibm.com/did/8210bks.html
IBM Nways Manager for AIX Version 2.1 Installation Guide	www.networking.ibm.com/cma/instweb.html
ELAN Management Using Nways Campus Manager for AIX Redbook	publib.boulder.ibm.com/pubs/pdfs/redbooks/sg244821.pdf
A Quick Guide to Java Performance Manager	www.networking.ibm.com/cma/cmasolut.html (select the documentation link from here)
IBM Networking Support Page	www.networking.ibm.com/support

Trademarks

AIX, DB2, IBM, Netfinity, Netfinity Manager, and Nways are trademarks of International Business Machines Corporation in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks of Microsoft Corporation.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Lotus Notes is a trademark of Lotus Development Corporation in the United States or other countries or both.

NetView and Tivoli are trademarks of Tivoli Systems, Inc.

Other company, product, and service names may be trademarks or service marks of others.