Anixter
Interoperability Laboratory

PRODUCT EVALUATION

IBM 8265
Nways ATM SWITCH V3.3.5

155 Mbps CLIENT/SERVER TEST

September 1998

Table of Contents

Introduction

ATM switching is being deployed in many LAN and WAN backbones to address various network requirements. One such requirement is the consolidation of different traffic types such as voice, video, and data into a single digital-transport stream. This increases the aggregate available bandwidth in the backbone, which potentially translates into increased productivity and cost savings. ATM can also be installed on many different cabling plants, thus making implementation easier than with other backbone technologies.

Since data delivery alone is no longer the deciding concern for many companies in the midst of re-evaluating and retooling their backbone infrastructures, ATM offers an excellent option for voice convergence. With ATM's native quality-of-service features and cell size (53 bytes), it handles time-sensitive applications better than frame-based technologies. Prioritization schemes in the frame world have not yet been solidified and will remain tentative for some time. In light of this, many companies have elected to upgrade their backbones to ATM and even adopt it as the wide area technology of choice. This can dramatically reduce the operating costs associated with maintaining multiple parallel networks. Service providers continue to attract more customers with cost-effective alternatives that compete directly with older, less accommodating technologies. Together, these benefits provide customers with a degree of future-proofing that can take them comfortably into the next millenium.

Figure 1: Migration from frame-based to cell-based backbone

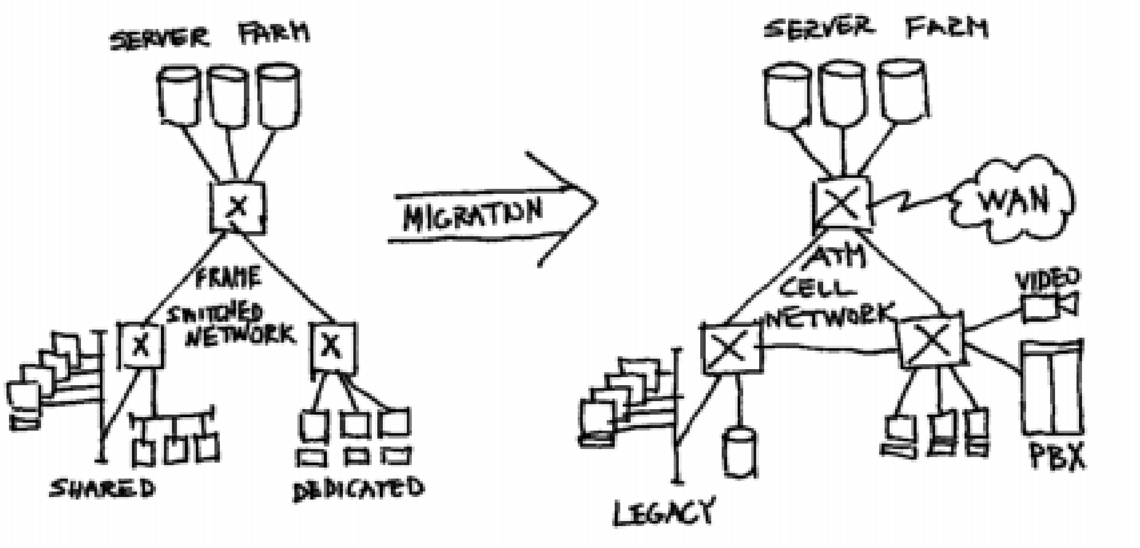

One of our goals at Anixter Networking is to share our technical knowledge with our customers and the networking industry as a whole. Much of our knowledge comes from the work we do in our Interoperability Laboratory. Our interoperability engineers use this multimillion-dollar facility to simulate, test, and troubleshoot customer networks. Another key function of the lab is to evaluate new technologies and product interoperability while solving connectivity issues.

These test results allow us to help customers make informed networking decisions, identify and correct potential interoperability problems, and expedite network implementation. With networking products from the industry's top manufacturers as well as multiple network-management platforms, no other network-centric lab is as large or offers the variety and scope of network-management and hardware-configuration options. Moreover, Anixter is a Principal member of the ATM Forum. We help to define the specifications at technical committee meetings by leveraging our knowledge with the needs of our customers.

In this paper, we will evaluate the IBM 8265 ATM Switch in our lab environment. After presenting a brief product overview, we describe the test methodology, outline the performance results, and discuss some of the product's user functions. This paper is intended primarily to help network infrastructure managers and network operations and planning professionals to make informed decisions about their networks.

Product Overview

The IBM 8265 Nways ATM Switch is a 17-slot chassis used in enterprise LAN and WAN backbone networks. The custom-developed ASIC, known as Prizma or the switch-on-a-chip, delivers an aggregate throughput of 12.8 Gbps (full duplex) on a 25 Gbps backplane. It is a single-stage nonblocking crossbar-switching fabric made up of two parallel Prizma chips on a dual-slotted module known as the CPSW module (control point switch). The CPSW module must occupy slots 9 and 10 where specially-fitted backplane connections exist; a redundant CPSW may be positioned in slots 11 and 12. The CPSW module is outfitted with a PCMCIA flash memory card for the microcode image. An intelligent controller (single-slot mini-module) is responsible for bus clocking, power management, and inventory control. It is a required component seated in special slot 18 (and slot 19 for redundancy).

The 8265 is standards-based and adheres to many ATM Forum specifications as well as some IETF ones. It supports all ATM service categories including CBR, VBR, (real time and non-rt), ABR, and UBR. Both traffic shaping and policing are provided when contention occurs on the outbound ports. A pool of output buffers then helps maintain the traffic contracts. Some of the relevant specifications are IISP, PNNI 1.0, Signaling 4.0, Traffic Management 4.0, and LAN Emulation Client 1.0 from the ATM Forum. The IETF provides relevant specifications with RFCs 1483,1577, and 2233, Multi Protocol Encapsulation over ATM, Classical IP over ATM, and Next Hop Resolution Protocol, respectively.

Figure 2: External view of the IBM 8265 Nways ATM Switch (front and back)

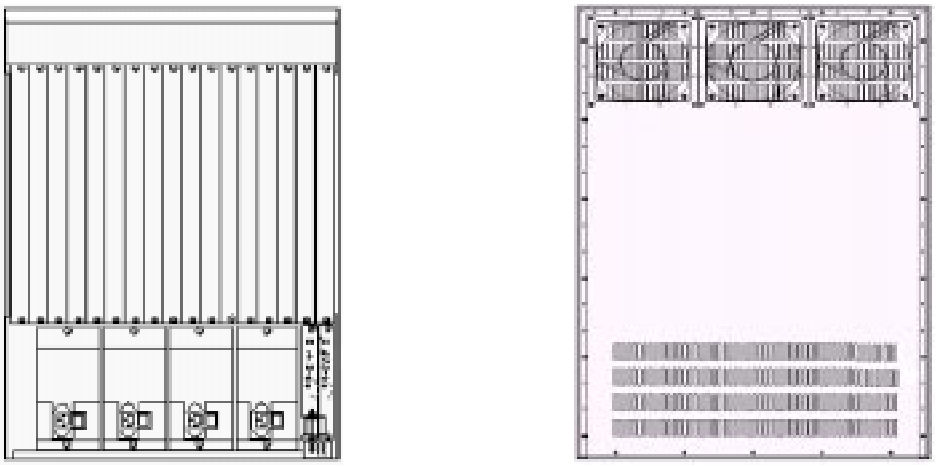

One of the distinct advantages of the 8265 Nways ATM Switch is that it can be integrated with the multiprotocol switched services (MSS) module. Some of the MSS functions include routing services for IP, IPX, and AppleTalk; LANE, MPOA, NHRP server functionality and authentication, to name a few.

The Nways 8265 supports the following features and modules:

- Maximum configuration of 56 OC-3 (155 Mbps) or 14 OC-12 (622 Mbps) ports
- SNMP management and configuration with GUI, Web, and CLI interfaces
- NEBS (Network Equipment-Building System) certified from Bell Communications Research (Bellcore)
- MSS integrated-routing support services
- 8260 ATM (slots 1, 3, 5, and 7) – 12-port 25 Mbps UTP/STP, 4-port 100 Mbps MMF, 2-port 155 Flex, 3-port 155 Flex, 8-port video (MPEG) distribution, 8281 LAN bridge (2-slot), 8271 Ethernet switch (2- and 3-slot), 8272 Token Ring switch (2- and 3-slot)
- 2-port ATM WAN for E3, DS3, STM-1, and OC-3
- 4-port ATM WAN for E1, T1, and J1
- 4-port OC-3 flexible for any combo of MMF, SMF (short and long range), and UTP/STP
- 4-port OC-3 MMF
- 1-port OC-12 MMF
- 1-port OC-12 SMF (short and long range)
- FiberCom ATM circuit emulation
- Environmental controller (slots 18 and 19)
- Up to four load-sharing power supplies (295 and 415 watt capacities)

Test Methodology

As mentioned earlier, many large companies consider ATM to be the only viable transport mechanism for consolidating different traffic types with various quality-of-service requirements. By placing ATM switching in the core of their networks, these companies ensure some immediate benefits in the form of scalable aggregate bandwidth and a degree of future-proofing that is not attainable with other current LAN technologies. Perhaps the widespread use of integrated voice and video applications has yet to impact most networks, but this still looms ahead. In the meantime, replacing frame- based internetworking devices in the core of the network with ATM cell-based switching can help many of these overburdened networks to avoid bottlenecks.

The purpose of this product evaluation is to compare and contrast client/server-oriented communication in both a frame-switched and cell-switched environments. Similar network designs are employed and throughput performance results are given. LANE Services V1.0 are used in the ATM design to mimic the functionality of the Layer 2 frame-switched design. This provides for investment protection, the continued use of current client/server based applications, and the capability to integrate frame-based (Ethernet and Token Ring) devices with ATM direct-attached ones in the same network.

We will also describe several other incremental steps taken on the cell side during this integration process. For instance, the intersubnet communications and the operation of dynamic Private Node-Node Interface V1.0 (PNNI) fail-over are examined in terms of their performance impact on client/server and LANE communications.

The test suite involves four different scenarios.
Scenario 1—The traditional frame-based network consists of 8271 Ethernet switches in the core as our baseline configuration.
Scenario 2—Phase One of the ATM network integration involves placing 8265 cell switches in the core and repositions the 8271s as edge devices. The MSS is introduced to support LANE Services for a single ELAN; IISP is used between the switches.
Scenario 3—Phase Two of the ATM integration introduces another 8271 and a second ELAN group is created. Both ELANs will communicate through the routed services of the MSS. Instead of IISP, PNNI is implemented between the switches.
Scenario 4—Phase Three of the ATM network integration involves communications between two ELAN groups, one with direct-attached OC-3 clients and the other group connected via the 8271.

Each of these scenarios will be presented separately and in detail; the configurations, test methods, and results are inclusive. For Scenarios 1 and 2, the Netcom Smartbits was connected to the 8271s to generate traffic. Ganymede's Chariot tool was also used in all of the scenarios to generate application-oriented traffic (referred to here as client/server communications) from the endpoint to the servers. For a more elaborate explanation of how both of these test tools operate, please refer to the Appendix.

Scenario 1: Frame-based network (baseline)

This baseline scenario uses two 8271 (model 712) Ethernet switches back-to-back via a 100-Mbps Fast Ethernet full-duplex connection. Both switches used hardware version 5, upgrade software 3.22a, and boot software 3.10. Two servers, S1 and S2 with IBM 10/100 EtherJet PCI adapters (configured for full-duplex 100 Mbps), are attached to one switch. Six client workstations, WS1–WS6 with the same EtherJet PC1 adapters, are attached to the other switch (configured for half-duplex 10 Mbps). All eight endpoints are Compaq P133s loaded with WinNT Workstation. In this particular scenario, all physical connections for all devices—including the Smartbits—were made with Category 5 cabling.

Figure 3: Frame-switched backbone

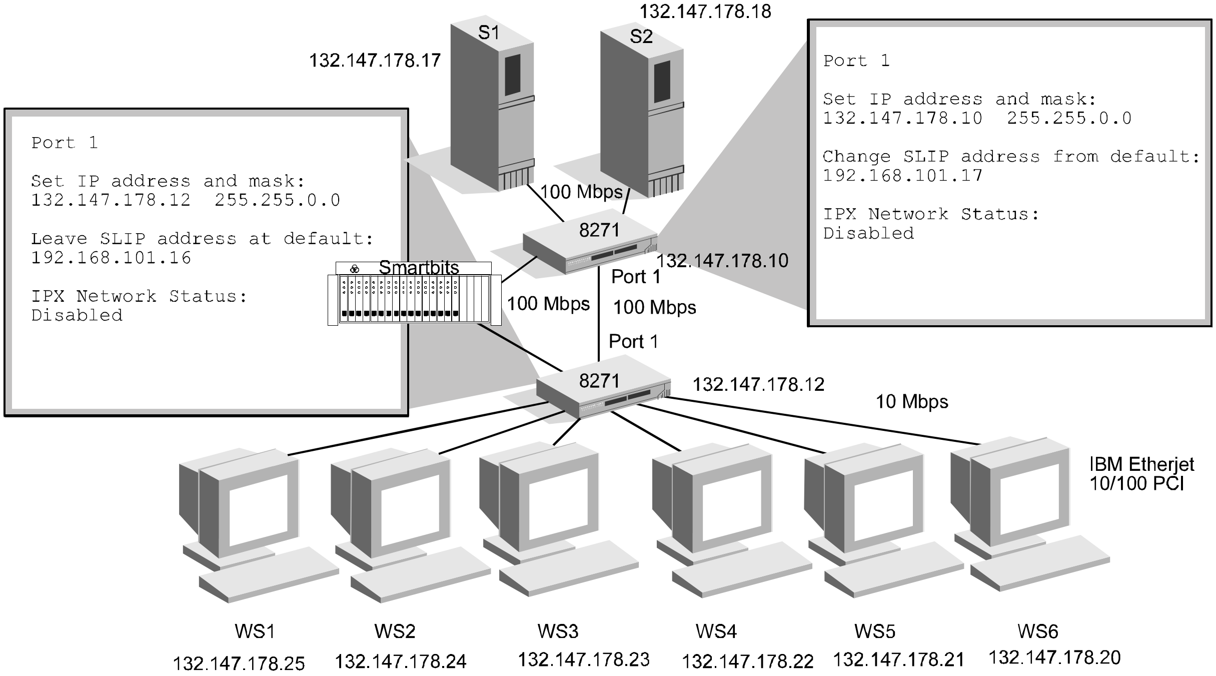

```
Port 1

Set IP address and mask:
132.147.178.12  255.255.0.0

Leave SLIP address at default:
192.168.101.16

IPX Network Status:
Disabled
```

```
Port 1

Set IP address and mask:
132.147.178.10  255.255.0.0

Change SLIP address from default:
192.168.101.17

IPX Network Status:
Disabled
```

Chariot endpoint software v2.1 was also loaded on the servers and the clients while another Compaq desktop, attached to the server switch, was used for the Chariot console.

Table 1: Client/Server corresponding TCP sessions for Scenarios 1, 2 and 3

Pair #	**Source**	**Destination**	**Pair #**	**Source**	**Destination**
1	WS1	S1	2	WS1	S2
3	WS2	S1	4	WS2	S2
5	WS3	S1	6	WS3	S2
7	WS4	S1	8	WS4	S2
9	WS5	S1	10	WS5	S2
11	WS6	S1	12	WS6	S2

The Chariot endpoints in each workstation had two TCP sessions (connection-oriented) attached to each of the servers. For example, pair 1 refers to a session from WS1 to S1, pair 2 refers to a session from WS1 to S2, and so on. The Chariot script "send file long" was used to establish a total of 12 sessions: This file was 100 Kbytes. All of the scenarios used the same corresponding workstation-to-Chariot pairings except for Scenario 4, which will be presented later.

The results are below in Table 2. For all pairs the aggregate throughput was about 47 Mbps. Look at the individual pair 10, which represents WS5 communicating to S2. Its throughput was about 4 Mbps. These test results take into consideration both the Smartbits, operating at 49 percent load, and the Chariot test that ran simultaneously. The frame size of the Smartbits traffic was 1500 bytes.

Table 2: Chariot results from the frame-based test (Scenario 1)

Group/ Pair	Throughput Average (Mbits/sec)	Throughput Minimum (Mbits/sec)	Throughput Maximum (Mbits/sec)	95% Confidence Interval	Measured Time (secs)	Relative Precision
All Pairs	**47.416**	**2.981**	**4.459**			
1	3.954	2.981	4.233	0.009	598.897	0.231
2	3.958	3.891	4.396	0.004	598.311	0.103
3	3.946	3.023	4.180	0.008	598.143	0.212
4	3.948	3.863	4.440	0.004	599.793	0.097
5	3.952	3.203	4.459	0.007	599.172	0.181
6	3.949	3.153	4.346	0.007	599.719	0.187
7	3.946	3.086	4.152	0.008	598.074	0.196
8	3.951	3.895	4.303	0.004	599.364	0.093
9	3.955	3.901	4.211	0.003	598.734	0.071
10	**3.953**	**3.884**	**4.167**	**0.003**	**599.053**	**0.074**
11	3.951	3.426	4.219	0.005	599.360	0.124
12	3.954	3.906	4.437	0.004	598.894	0.098
Totals:	47.416	2.981	4.459			

We then ran the Smartbits by itself to determine latency and frame loss. By increasing the load output by the Smartbits—starting at 49 percent at increments of 10—we saw that at 79 percent, the 8271s began to drop frames. The average latency across the 8271s (send and receive) was about 23.4 microseconds. The frame size of the Smartbits traffic used in these tests was 64 bytes, the typical accepted value used when performing latency tests. It is realized that most network traffic uses much larger frame size rather than this minimum. For the specific Smartbits configuration and results, please refer to the Appendix at the back of this paper.

Scenario 2: ATM network (integration of Phase One)

The first step in the ATM integration process includes three 8265s running version 3.3.5 boot and flash EEPROMS. The 8265s are configured in a delta design with IISP. The 8271s and an MSS are attached via 155-Mbps ATM connections to two of the 8265s. The MSS provides LAN Emulation Services, among other things, and includes the LECS, LES, and BUS functionalities. This is a bridged network with a single ELAN and no routing. The ATM 155-Mbps connections between the 8271s and the 8265s are MMF.

Figure 4: ATM backbone in bridged mode (one ELAN) and IISP

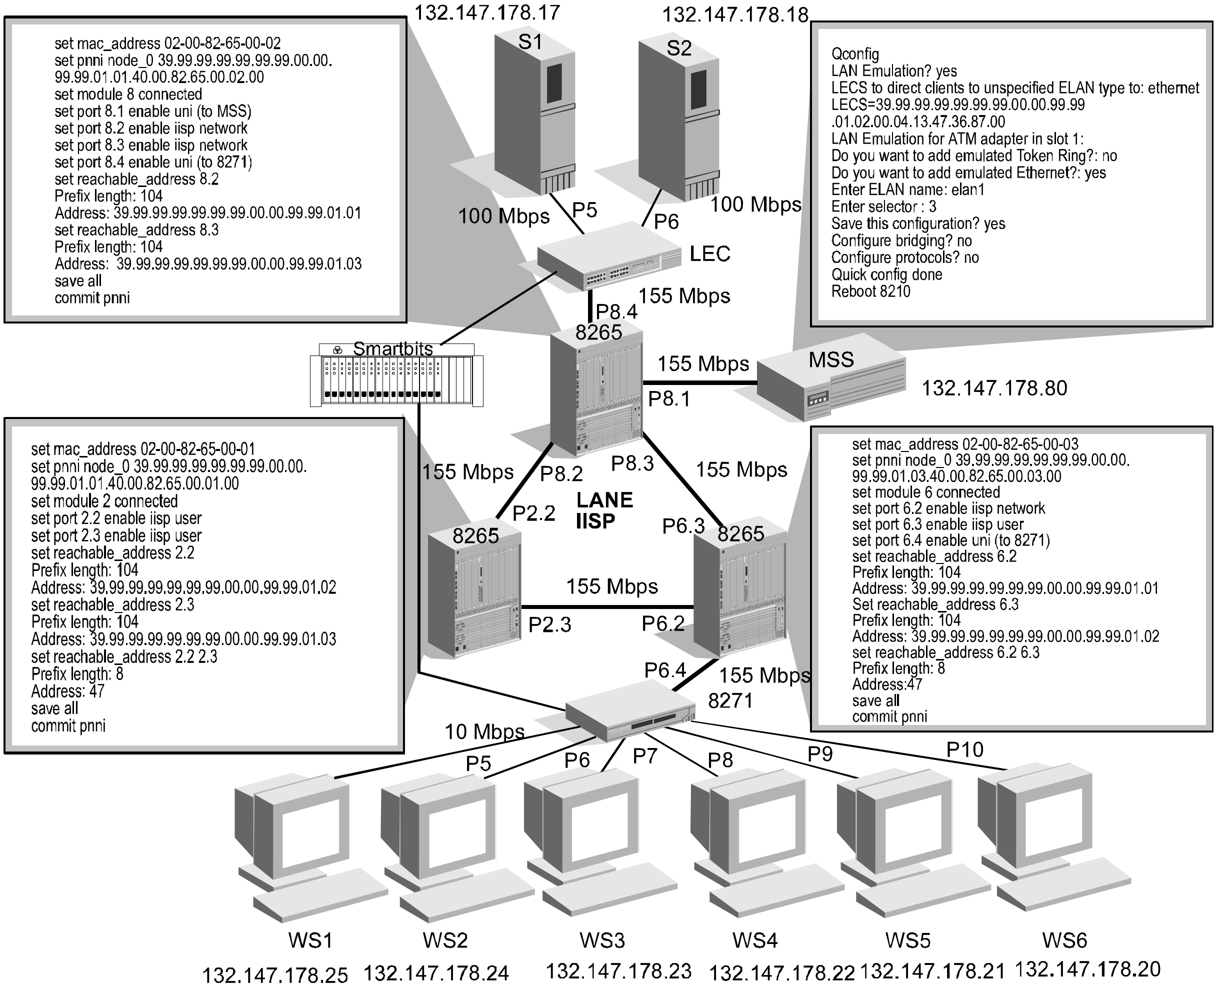

```
set mac_address 02-00-82-65-00-02
set pnni node_0 39.99.99.99.99.99.99.00.00.
99.99.01.01.40.00.82.65.00.02.00
set module 8 connected
set port 8.1 enable uni (to MSS)
set port 8.2 enable iisp network
set port 8.3 enable iisp network
set port 8.4 enable uni (to 8271)
set reachable_address 8.2
Prefix length: 104
Address: 39.99.99.99.99.99.99.00.00.99.99.01.01
set reachable_address 8.3
Prefix length: 104
Address:  39.99.99.99.99.99.99.00.00.99.99.01.03
save all
commit pnni
```

```
Qconfig
LAN Emulation? yes
LECS to direct clients to unspecified ELAN type to: ethernet
LECS=39.99.99.99.99.99.99.00.00.99.99
.01.02.00.04.13.47.36.87.00
LAN Emulation for ATM adapter in slot 1:
Do you want to add emulated Token Ring?: no
Do you want to add emulated Ethernet?: yes
Enter ELAN name: elan1
Enter selector : 3
Save this configuration? yes
Configure bridging? no
Configure protocols? no
Quick config done
Reboot 8210
```

```
set mac_address 02-00-82-65-00-01
set pnni node_0 39.99.99.99.99.99.99.00.00.
99.99.01.01.40.00.82.65.00.01.00
set module 2 connected
set port 2.2 enable iisp user
set port 2.3 enable iisp user
set reachable_address 2.2
Prefix length: 104
Address: 39.99.99.99.99.99.99.00.00.99.99.01.02
set reachable_address 2.3
Prefix length: 104
Address: 39.99.99.99.99.99.99.00.00.99.99.01.03
set reachable_address 2.2 2.3
Prefix length: 8
Address: 47
save all
commit pnni
```

```
set mac_address 02-00-82-65-00-03
set pnni node_0 39.99.99.99.99.99.99.00.00.
99.99.01.03.40.00.82.65.00.03.00
set module 6 connected
set port 6.2 enable iisp network
set port 6.3 enable iisp user
set port 6.4 enable uni (to 8271)
set reachable_address 6.2
Prefix length: 104
Address: 39.99.99.99.99.99.99.00.00.99.99.01.01
Set reachable_address 6.3
Prefix length: 104
Address: 39.99.99.99.99.99.99.00.00.99.99.01.02
set reachable_address 6.2 6.3
Prefix length: 8
Address:47
save all
commit pnni
```

The rest of the network remained the same; the clients and the servers were configured as they were previously. Both the Smartbits and Chariot test tools also remained configured and attached like before.

The aggregate throughput for all pairs and for pair 10 specifically (in Table 3) is approximately 55 Mbps and 4.6 Mbps, respectively. This is an increase of around 18 percent from the previous scenario. Again, these results take into consideration both the Smartbits, operating at 49 percent load, and the Chariot test, which ran simultaneously. The frame size of the Smartbits traffic was 1500 bytes.

Table 3: Chariot results from the ATM bridged test (Scenario 2)

```
             Throughput  Throughput  Throughput        95% Measured
Group/        Average     Minimum     Maximum Confidence     Time  Relative
Pair      (Mbits/sec) (Mbits/sec) (Mbits/sec)   Interval   (secs) Precision
--------- ----------- ----------- ----------- ---------- -------- ---------
All Pairs      54.933       3.276       6.993
 1              4.581       3.300       6.584      0.014  599.037     0.312
 2              4.575       3.310       6.706      0.014  599.741     0.313
 3              4.574       3.513       4.665      0.010  598.229     0.212
 4              4.585       3.935       6.488      0.010  598.539     0.208
 5              4.580       3.571       6.074      0.011  599.141     0.239
 6              4.577       3.537       5.785      0.011  599.472     0.237
 7              4.579       3.276       5.442      0.014  599.240     0.301
 8              4.576       3.902       6.993      0.013  599.713     0.279
 9              4.589       3.899       5.319      0.008  599.642     0.169
 10             4.562       3.564       4.944      0.010  599.734     0.219
 11             4.577       3.941       4.654      0.007  599.471     0.153
 12             4.578       3.904       5.109      0.007  599.420     0.153
Totals:        54.933       3.276       6.993
```

The Smartbits was also run individually to determine the latency and frame loss. The average latency incurred in Scenario 2 was about 90 microseconds; however, this needs to be put into the proper perspective. This measurement includes the latency previously reported as 23.4 microseconds across two 8271s, two SAR operations, two ATM switch hops, and LANE encapsulation. At 79 percent load, the 8271s again began to drop frames. At the time of testing, Netcom Systems was unable to provide us with stable 155-Mbps ATM modules; therefore, we determined not to use them in the tests. They have since rectified the problem, but that is why we strapped the Smartbits to the 8271s in Scenario 2. It also enabled us to compare the two scenarios. For the specific Smartbits configuration and results, please refer to the Appendix at the back of this paper.

Scenario 3: ATM network (integration of Phase Two)

Instead of the servers connecting at 100 Mbps to the 8271, they are now direct-attached at 155 Mbps (using Interphase adapters) to the 8265 ATM switch. The 8271 is connected to one of the 8265s and supports three 10 Mbps clients, which constitutes ELAN 1 on subnet 132.147.0.0. The other 8271 is connected to the last 8265 and supports three 10-Mbps clients, which constitutes ELAN 2 on subnet 132.1.0.0. Server S1 is the server for ELAN 1, and server S2 is for ELAN 2. All cabling remains the same except for the servers, the 8265s, and the MSS, which are now connected with MMF.

Figure 5: ATM backbone in routed mode (two ELANs) and PNNI

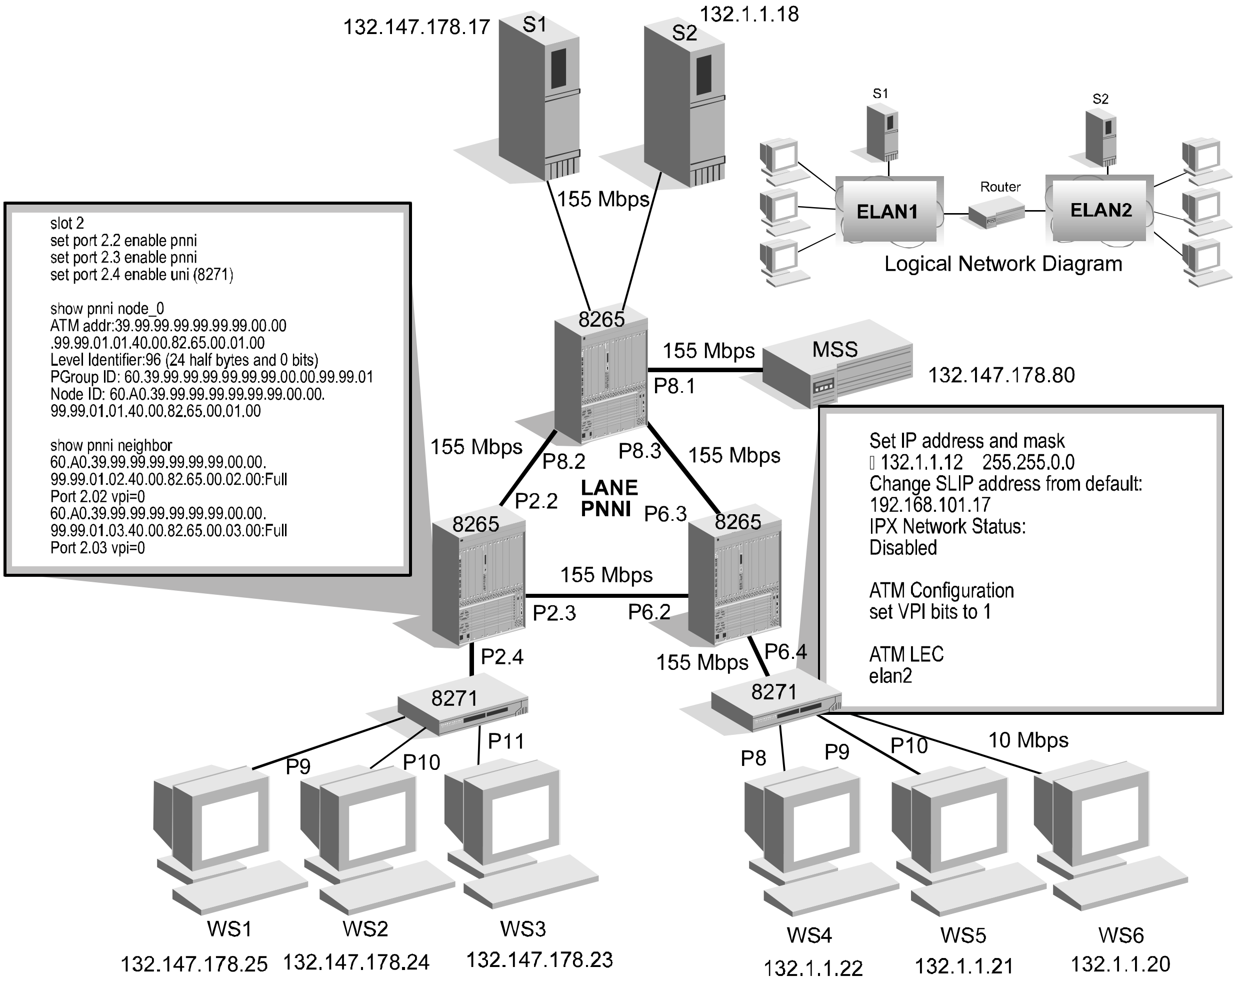

```
slot 2
set port 2.2 enable pnni
set port 2.3 enable pnni
set port 2.4 enable uni (8271)

show pnni node_0
ATM addr:39.99.99.99.99.99.99.00.00
.99.99.01.01.40.00.82.65.00.01.00
Level Identifier:96 (24 half bytes and 0 bits)
PGroup ID: 60.39.99.99.99.99.99.99.00.00.99.99.01
Node ID: 60.A0.39.99.99.99.99.99.99.00.00.
99.99.01.01.40.00.82.65.00.01.00

show pnni neighbor
60.A0.39.99.99.99.99.99.99.00.00.
99.99.01.02.40.00.82.65.00.02.00:Full
Port 2.02 vpi=0
60.A0.39.99.99.99.99.99.99.00.00.
99.99.01.03.40.00.82.65.00.03.00:Full
Port 2.03 vpi=0
```

```
Set IP address and mask
▯ 132.1.1.12   255.255.0.0
Change SLIP address from default:
192.168.101.17
IPX Network Status:
Disabled

ATM Configuration
set VPI bits to 1

ATM LEC
elan2
```

Communications between these different ELANs will be accomplished through the MSS routing function. Also, the IISP signaling has been replaced by PNNI between the ATM switches.

Table 4 below shows the aggregate throughput for all pairs and pair 10, which were 54 Mbps and 4.6 Mbps, respectively. When compared to the bridged test in Scenario 2, there was virtually no performance penalty when communicating between different ELANs through the MSS router. The Smartbits was not used to test for latency or frame loss in Scenarios 3 and 4.

Table 4: Results from the ATM routed test (Scenario 3)

```
             Throughput  Throughput  Throughput        95% Measured
Group/          Average     Minimum     Maximum Confidence      Time  Relative
Pair        (Mbits/sec) (Mbits/sec) (Mbits/sec)   Interval    (secs) Precision
--------- ----------- ----------- ----------- ---------- -------- ---------
All Pairs      54.602       2.187       5.878
 1              4.565       4.278       4.988      0.006  599.338     0.122
 2              4.536       3.203       4.657      0.012  599.679     0.262
 3              4.567       4.022       4.703      0.006  599.087     0.141
 4              4.535       3.768       4.711      0.007  598.086     0.160
 5              4.568       4.122       4.703      0.006  599.018     0.129
 6              4.533       3.730       4.659      0.008  598.341     0.168
 7              4.536       4.202       5.878      0.010  599.641     0.217
 8              4.561       2.187       4.831      0.030  598.173     0.655
 9              4.522       3.423       5.181      0.011  599.697     0.240
 10             4.577       3.939       4.720      0.007  599.514     0.162
 11             4.522       3.808       4.825      0.007  599.681     0.161
 12             4.581       4.175       4.698      0.005  599.019     0.119
Totals:        54.602       2.187       5.878
```

During this scenario, we tested the fail-over operation of the ping links. A continuous ping was sent in equal distribution from the workstations to the servers at the same time as the Chariot "send file long" test. The MMF cable between the top and bottom left 8265s was disconnected. The Chariot reporting program captured this time interval as approximately 12 seconds. The ATM system took about 12 to realize the failure and route around it. This was quicker than other switches we have tested.

Scenario 4: ATM network (integration of Phase Three)

In this final scenario, one of the 8271s that was connected to an 8265 was removed from the design. The three clients that were attached to it have been reconfigured as direct-attached 155 Mbps (using Interphase adapters) to the 8265 ATM switch. These clients still remain part of ELAN 1 and are cabled with MMF connections. Logically, S1 and S2 both have instances of ELANs 1 and 2. More specifically, ELAN 1 on S1 is associated with subnet 132.147.178.17, and ELAN 2 on S1a is associated with 132.1.1.17. ELAN 1 on S2 is associated to subnet 132.147.178.18, and ELAN 2 on S2a is associated with subnet 132.1.1.18. Communications between ELANs require the routing functions of the MSS. There are dual PNNI links between two of the 8265s to test the effects of link failure on the client/server communications. Unfortunately, we were unable to properly test the PNNI dual-link operation due to time constraints.

Figure 6: ATM backbone in routed mode (two ELANs) with redundant PNNI link

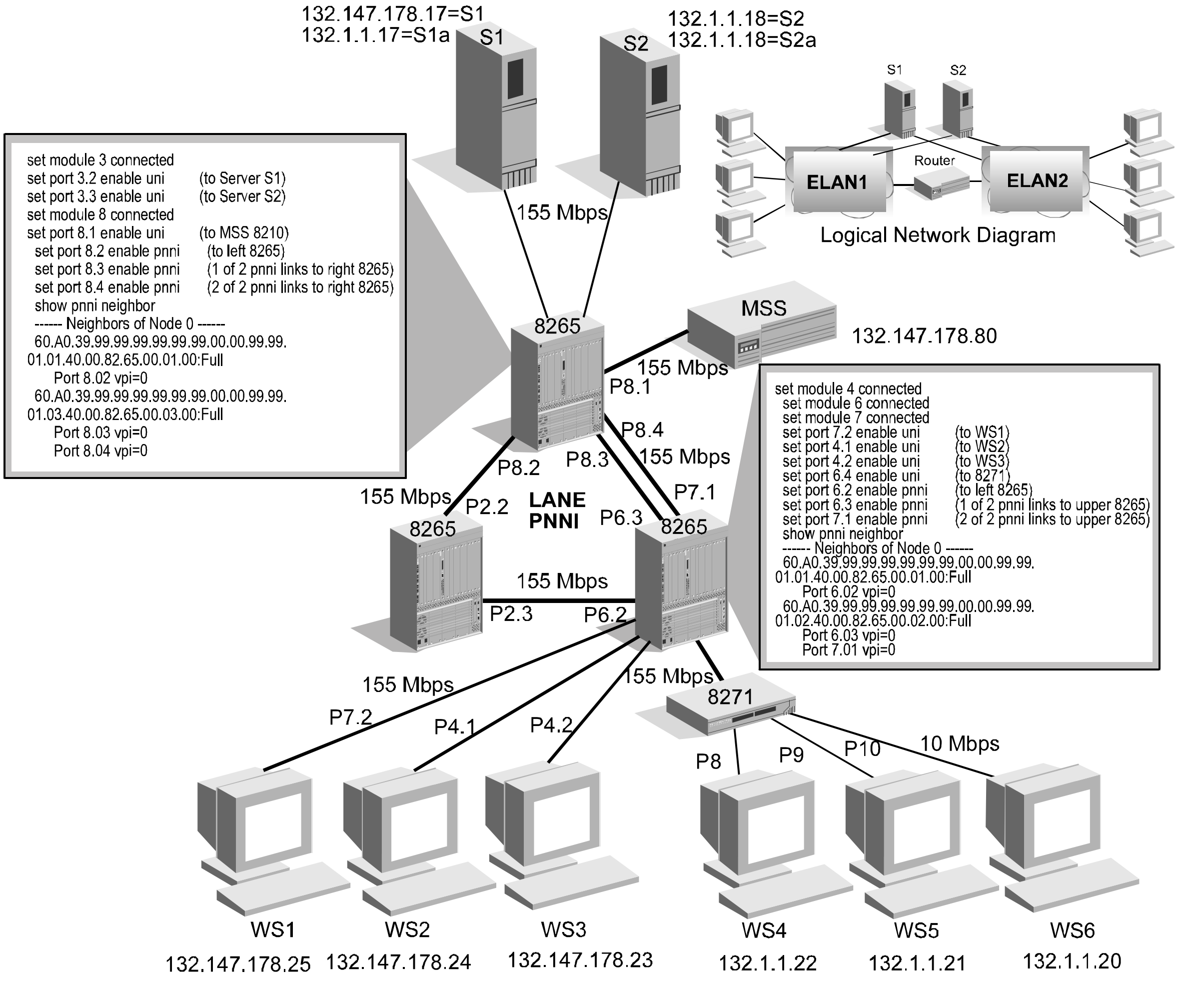

```
set module 3 connected
set port 3.2 enable uni        (to Server S1)
set port 3.3 enable uni        (to Server S2)
set module 8 connected
set port 8.1 enable uni        (to MSS 8210)
  set port 8.2 enable pnni      (to left 8265)
  set port 8.3 enable pnni      (1 of 2 pnni links to right 8265)
  set port 8.4 enable pnni      (2 of 2 pnni links to right 8265)
  show pnni neighbor
  ------ Neighbors of Node 0 ------
  60.A0.39.99.99.99.99.99.99.00.00.99.99.
01.01.40.00.82.65.00.01.00:Full
       Port 8.02 vpi=0
  60.A0.39.99.99.99.99.99.99.00.00.99.99.
01.03.40.00.82.65.00.03.00:Full
       Port 8.03 vpi=0
       Port 8.04 vpi=0
```

```
set module 4 connected
  set module 6 connected
  set module 7 connected
  set port 7.2 enable uni       (to WS1)
  set port 4.1 enable uni       (to WS2)
  set port 4.2 enable uni       (to WS3)
  set port 6.4 enable uni       (to 8271)
  set port 6.2 enable pnni      (to left 8265)
  set port 6.3 enable pnni      (1 of 2 pnni links to upper 8265)
  set port 7.1 enable pnni      (2 of 2 pnni links to upper 8265)
  show pnni neighbor
  ------ Neighbors of Node 0 ------
  60.A0.39.99.99.99.99.99.99.00.00.99.99.
01.01.40.00.82.65.00.01.00:Full
       Port 6.02 vpi=0
  60.A0.39.99.99.99.99.99.99.00.00.99.99.
01.02.40.00.82.65.00.02.00:Full
       Port 6.03 vpi=0
       Port 7.01 vpi=0
```

ATM networks can now use PNNI's signaling and dynamic routing protocol, which compares to Spanning Tree Protocol in the frame-based bridged or switched environment.

The Virtual Router Redundancy Protocol (VRRP) is another standard currently referenced by the IETF's RFC 2338. It is being used in frame-based routed networks to provide a standardized means for fault tolerance and load balancing in a mixed vendor- router environment. We used a single peer group here: PNNI V1.0 allows for link backup and load balancing between different PNNI peer groups.

Table 5: Client/Server corresponding TCP sessions for Scenario 4

Pair #	**Source**	**Destination**	**Pair #**	**Source**	**Destination**
1	WS1	S1	2	WS1	S2a
3	WS2	S1	4	WS2	S2a
5	WS3	S1	6	WS3	S2a
7	WS4	S1a	8	WS4	S2
9	WS5	S1a	10	WS5	S2
11	WS6	S1a	12	WS6	S2

Table 6: Results from the ATM routed test (Scenario 4)

Group/ Pair	Throughput Average (Mbits/sec)	Throughput Minimum (Mbits/sec)	Throughput Maximum (Mbits/sec)	95% Confidence Interval	Measured Time (secs)	Relative Precision
All Pairs	**69.362**	**0.092**	**18.433**			
1	8.569	0.636	17.506	0.702	599.377	8.196
2	8.176	0.528	17.544	0.724	595.877	8.856
3	8.024	0.565	16.394	0.703	599.234	8.766
4	8.938	0.782	16.461	0.546	598.809	6.109
5	**8.548**	**0.318**	**17.621**	**1.036**	**598.055**	**12.124**
6	8.103	0.447	18.433	0.910	599.290	11.228
7	3.360	0.454	4.706	0.183	597.685	5.448
8	3.265	0.429	5.517	0.227	597.782	6.943
9	3.294	0.222	4.384	0.378	597.514	11.476
10	**3.034**	**0.140**	**4.776**	**0.621**	**598.557**	**20.476**
11	3.003	0.092	4.782	0.842	599.417	28.038
12	3.049	0.184	5.022	0.510	598.223	16.742
Totals:	69.362	0.092	18.433			

Table 6 shows the combined throughput for all pairs of nearly 70 Mbps. Pair 5, which is WS3 (an ATM 155 Mbps direct-attached workstation) that communicates to S1, had an average throughput of around 9 Mbps. Pair 10, which is WS5 (a 10 Mbps HD workstation) that communicates to S1, averaged only 3 Mbps. This dropped almost 1.5 Mbps from the previous scenario and is lower than average throughput for WS4, 5, and 6 when compared with Scenarios 1, 2, and 3. This is attributable to the PCs being used as servers and reaching the limit of their processing power.

Performance Results Summary

Given all of the hoopla about the cell tax of the ATM header—almost 11 percent—and the overhead associated with LANE encapsulation, the throughput performance results show an increase from Scenario 1 to Scenario 2. The segmentation and re-assembly (SAR) function that converts frames to cells and vice versa did not have that much of an impact on latency either; however, a 30 percent increase in average latency was measured in Scenario 2 in comparison to Scenario 1. When considering the addition of the ATM switches, there was more going on in Scenario 2. The cell-based scenario did not have as high an average percentage of dropped frames, although this is not exactly revealing since the frame switches were the weak links and began dropping frames at 79 percent load. (See Scenario 1.) Please refer to the Appendix for more detailed latency and frame loss results.

When comparing results from Chariot between Scenarios 2 and 3 (bridged versus routed ATM ELANs), the overall throughputs were almost identical at about 55 Mbps. This seems to show the performance benefit of routing between ELANs using the MSS. Strict testing measurements were not completed on the MSS since the focus of this evaluation was on the 8265 ATM switch. As we mentioned earlier, latency and frame loss tests using Smartbits were not done on Scenarios 3 and 4 since the ATM modules for the Smartbits were not stable enough to do so at the time of testing. We did feel compelled, however, to use this limitation to our advantage and serve as a reference for comparison between the frame- and cell-based environments.

Table 7: Bar graph of Scenarios 1–4 showing aggregate throughput performance

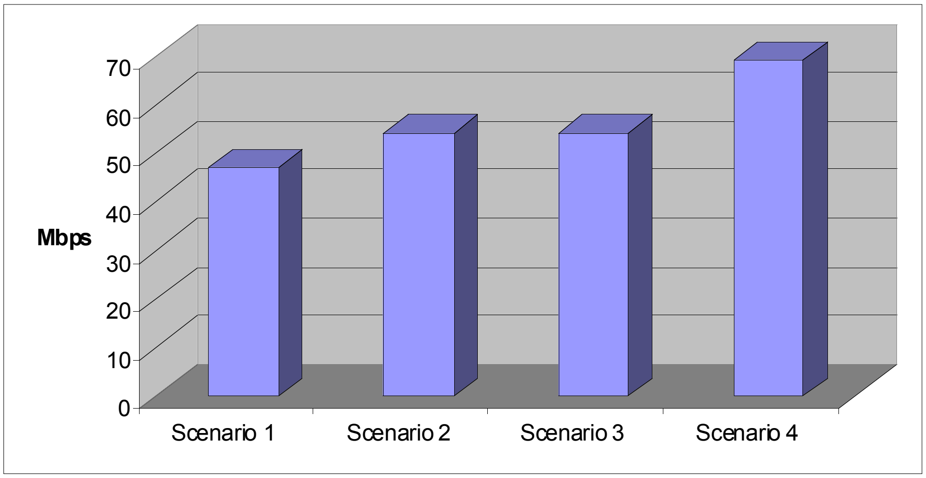

Scenario 4 shows the increase in aggregate throughput; it increased 25 percent from previous ATM scenarios. This is predominately a result of adding three direct-attached 155 Mbps workstations to the design. Again, routing between ELANs with the MSS had minimal negative impact on the throughput numbers.

User Functions

Configuration

Several configuration options exist, including a web server component for browser access and a CLI. A serial port connection allows a local console for service operations. Configuration code upload/download operations are done via TFTP (in-band or out-of-band). Remote access is accomplished via Telnet sessions both in-band (IP over ATM, IP over LANE) or out-of-band. There is an RJ-45 auxiliary port for connecting an Ethernet management station.

Documentation

Complete documentation is included on a web browser-based CD-ROM. Bound manuals and PNNI software can both be ordered separately as line items. A document that ships with the PCMCIA media has an ID and password that is needed to download the latest version of PNNI code for the CPSW from IBM's support web site. The most current operational, boot, and FPGA code for all modules is available free of charge as well. In fact, IBM recommends that all code be upgraded to the latest available versions prior to switch installation.

Security Options

Network access to the 8265 ATM network is provided for all types of ATM applications. When an ATM station connects to the 8265, it must register its address through ILMI. The network administrator can specify which ATM stations are allowed access. The MSS and not the 8265 switch institute many additional security features that can be implemented in an IBM ATM network.

Network Management

The following IBM SNMP management applications are used for the 8265:
- Nways Campus Manager for AIX (LAN, ATM, Suite)
- Nways Manager for Windows
- Nways Campus Manager LAN for HP
- Nways Campus Manager ATM for HP

Nways Campus Manager for Windows NT will be shipping around the first of the year in 1999.

MIB 2 support includes IETF A-to-M MIB, ATM Forum PNNI, and LANE MIBs.

Figure 7: Screen capture of Nways Campus Manger for AIX

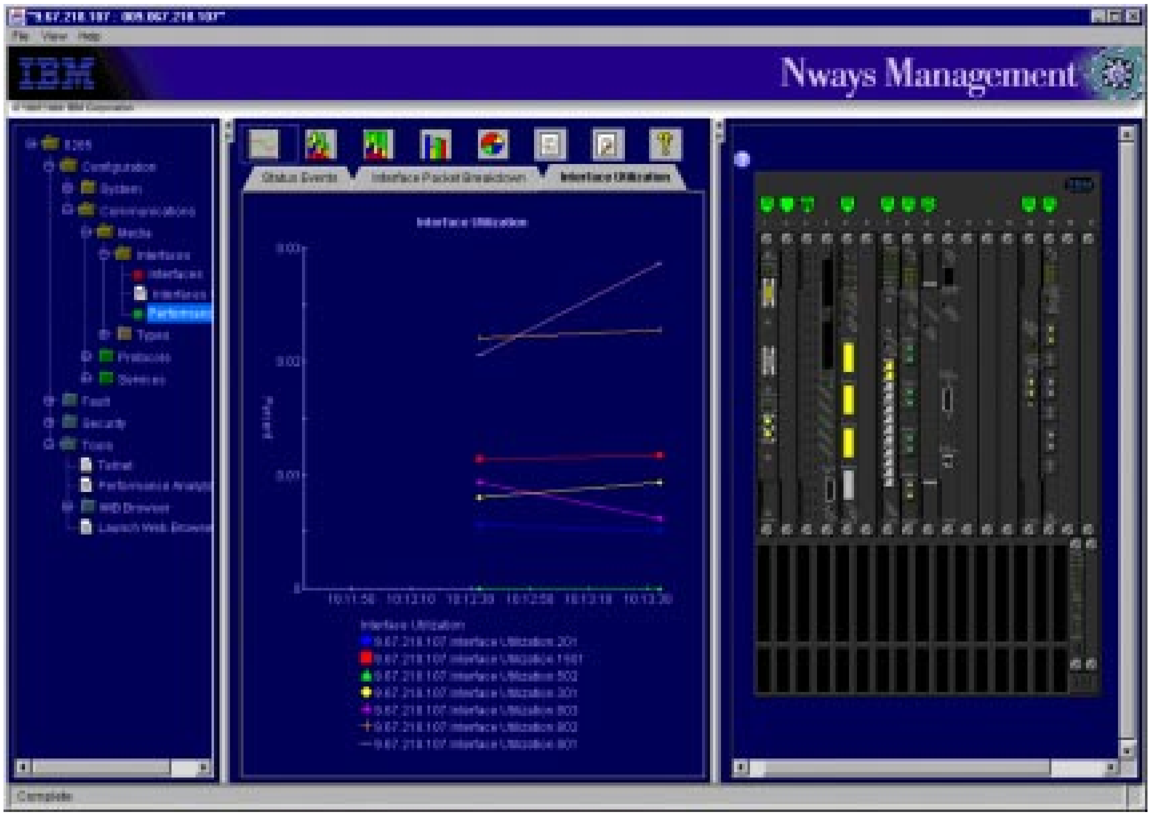

One feature, the environmental chassis-monitoring system, allows network administrators to monitor certain parameters of the 8265 such as power supplies, temperature, and inventory from the ATM CPSW module. Administrators get local access to all environmental parameters via the local console port or remote access via Telnet. The system triggers SNMP traps upon major events such as power supply failure or if the temperature threshold is exceeded.

Analysis Tools

Port mirroring can be configured on any of the ATM ports. Java applets are used for the display of the configuration, the PNNI topology, and trace/dump utility. The 8265 includes an error log and internal tracing capability for use by product support engineers. A recent innovation that was added after our product evaluation is IBM's ATM ping; it is more than a connectivity test between ATM switches. ATM ping allows you to check the signaling and PNNI protocol status. This is similar to the trace-route function of IP ping, and is useful for quickly identifying physical and logical problems in the ATM network.

Fault-tolerant Options

Three separate fan units are standard for temperature regulation. When a single fan unit fails, the 8265 is able to continue operation. Four 415-watt hot-swappable load-sharing power supplies (n + 1) are optional. Operation of two of the power supplies is sufficient to keep a loaded 8265 going.

The environmental-controller module functionality, which supplies power and temperature management, can be made redundant by adding a second controller module. (This is done in its own slot and leaves all host port slots available.)

Traffic Management

The 8265 has some advanced features that help to provide true QoS. There are four priority queues based on the five ATM service categories: CBR, VBR, real time and non-rt, ABR, and UBR. Each module has both input and output buffers that—along with per-VC queuing—help maintain service level agreements (SLAs). Traffic can also be regulated through the use of a relative rate of operation for ABR traffic. Other features include early and/or partial packet discard for any kind of traffic type, VC policing for congestion control, and instant viewing of counters per connection, port, and module. PNNI V1.0 uses automatic call rerouting at peer-group boundary by applying crankback.

Control Point Switch Module

The CPSW supports SVC signaling, both point-to-point and point-to-multipoint, according to ATM Forum signaling specifications V3.0, 3.1, and 4.0; interworking is available between the first two versions. Different signaling types such as UNI, IISP, and PNNI can also be implemented on the same physical interface.

Virtual path assignment per QoS allows traffic with different QoS requirements to be split over various VPs sharing the same ATM physical interface. This is a key consideration when connecting to a WAN-ATM carrier service as this will help maintain lower costs. Traffic shaping is done at the virtual path level. VP tunneling supports the interconnection of ATM campus switches using permanent virtual paths that allow SVCs to be transparently passed through the WAN. It also supports soft VPs in the carrier network that permit signaling to take place between the ingress and egress switches within the cloud.

Conclusion

The 8265 Nways ATM Switch was a solid performer with impressive features and functions. Some modules that were not available at the time of testing, like the ATM Inverse Multiplexer (AIM) and the ESCON channel attachment, are now obtainable. A new WinNT application, aimed at enterprise network management, is also currently shipping. The latest Control Point Switch V4 supports a dual boot load for versatile distributed-software operation.

A comparison between the bridged network designs in the frame- and cell-based environments was the foundation upon which we built several steps of ATM integration. The performance tests show a 17 percent increase in the aggregate throughput from Scenario 1 to Scenario 2. While the latency did increase from an average of 23 microseconds to an average of 90 microseconds between Scenarios 1 and 2, we were actually surprised that it was not more. This is in light of the fact that there were additional variables in Scenario 2—like three ATM switches, two SAR functions, and the LANE encapsulation—included in this measurement. In the frame-loss test, the 8271s did drop frames at about 79 percent load; however, the frame size of 64 bytes is not what you can generally expect in most production networks. We would not expect to see the 8271 drop frames at larger frame sizes like 1500 bytes.

We expected to see better overall throughput numbers in the ATM network because we added more ATM switches and ATM adapters to the servers and workstations—no surprise there. But the individual throughput numbers for ATM direct attached workstations wasn't as high as we expected; this was due to the workstations running out of processor power. The routing between ELANs did not show an appreciable degradation in throughput between the bridged and routed ATM scenarios.

Network managers will appreciate the straightforward configuration of the 8265. Additional functionality provided by the CPSW software and, in particular, the quality-of-service features like traffic management, make it an enterprise-class ATM switch. It is unfortunate that we were not able to test these benefits as they relate to converging different traffic types into a single transport stream. We believe this is one of the true advantages that ATM has over frame-based technologies.

One thing is for certain: The 8265 should be a top performer in any large ATM network because it delivers low latency and predictable behavior. This is just what the doctor ordered—stability for your networks.

Appendix

Ganymede Chariot

Chariot is a performance-test tool that can emulate a wide variety of network applications and system configurations. It is used in this test environment to help evaluate product performance. This software, which can run remote multiprotocol tests among many different operating systems, is controlled from a console located anywhere on the enterprise network. Chariot runs as a set of distributed applications over LANs, WANs, protocol stacks, adapters, routers, and virtually any device on the network. End-users can use Chariot to test the effects of adding new applications, measure response times, tune the network, and guide help-desk personnel in troubleshooting efforts.

Chariot includes two software components: the console and the endpoints. Tests are created and controlled from the console; the endpoints are special agents that run on desktops and generate traffic for the tests. Instructions to the console determine which endpoints will be communicating with each other and which protocol to use (HTTP, FTP, Telnet, etc.,).

Once a test is built, the console begins running it. Let's say you are running application scripts between endpoints 1 and 3, and between endpoints 2 and 4. The console begins by sending the appropriate scripts to endpoints 1 and 3. Endpoint 1 keeps its half of the script and sends the other half to endpoint 3. Endpoint 2 does the same thing with its partner. All of the endpoints wait for the console to tell them when to start. Once the console knows that the endpoints are ready, it tells them to begin executing their test scripts. When an endpoint runs an application script such as a file transfer or database query, it produces the same network flows that the application would have created. Endpoints keep track of the performance timing and forward this information back to the console. The console summarizes and displays the results that can be exported to text or HTML.

Netcom Systems' Smartbits

Smartbits can generate packets from any port or group of ports of fixed or variable length. Packet contents and the interpacket gaps are monitored while the rate errors, triggers, and latency can be output to a scope or logic analyzer. Smartbits can perform hub and router tests such as throughput, packet-loss rate, back-to-back performance, and latency. Smartbits can also monitor multiple networks simultaneously, perform latency tests between networks, and do network load simulation. It can also monitor error generation and the evaluation of overall network before and after the introduction of software applications. The interface modules available are 10 Mbps Ethernet, 10/100 Mbps Ethernet, 4/16 Mbps Token Ring, 100 Mbps FDDI, 155 Mbps ATM, as well as others.

Figure 8: Smartbits configurations

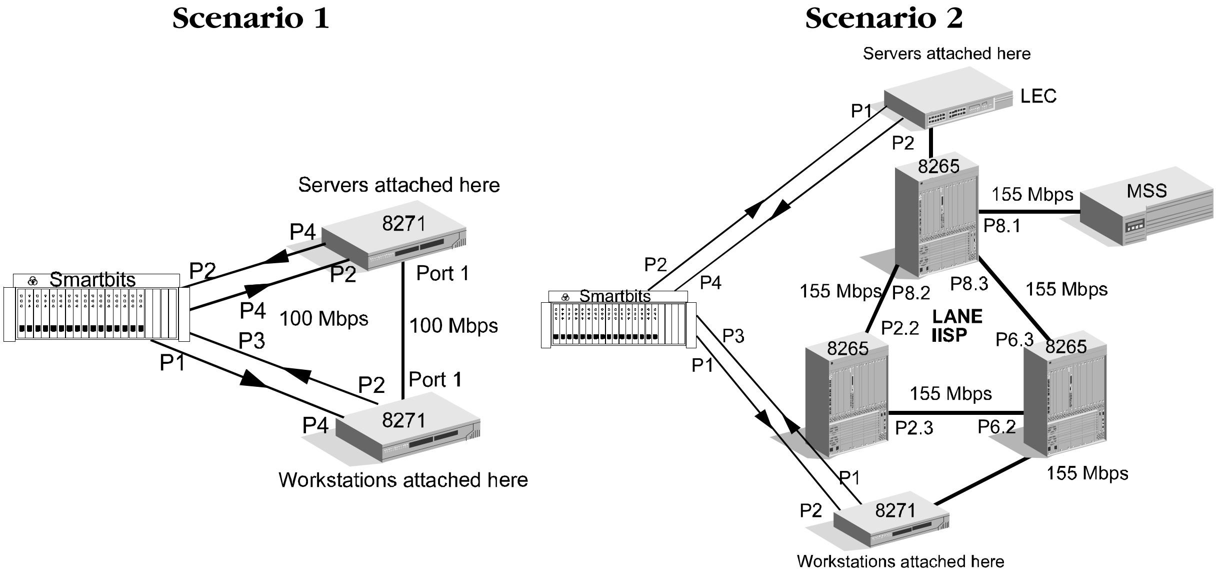

Scenario 1

Smarbits was used to generate a 50 percent load on the Ethernet ports using a packet size of 64 bytes, and was configured as follows: Port 1 transmitted 100 Mbps full-duplex traffic into the client switch and received it back on Port 2. Port 4 transmitted 100 Mbps full-duplex traffic into the server switch and received it back on Port 3. The first half of Table 8 below shows the latency from the Smartbits. At 49 percent load, the 22.90 microseconds of latency was measured between ports 1 and 2, and 23.83 microseconds between ports 4 and 3. At 79 percent load, the 8271 started to drop packets at a rate of one percent between ports 1 and 2, and eight percent between ports 4 and 3. These latency numbers at 79 percent load are not relevant since they are an artifact of the Smarbits' reporting program. This occurs when packets are lost and Smartbits begins to report invalid numbers.

Table 8: Results from the frame-based test (Smartbits)

Initial Rate(%)	1 to 2 (us)	4 to 3 (us)	Average
	100M -100M	100M -100M	
49%	22.90	23.85	23.375
59%	23.50	24.75	24.125
69%	22.45	24.95	23.700
79%	5972644.50	17331.30	2994987.900

Initial Rate(%)	1 to 2 (%)	4 to 3 (%)	Average
	100M -100M	100M -100M	
49%	0.000	0.000	0.000
59%	0.000	0.000	0.000
69%	0.000	0.000	0.000
79%	1.017	8.025	4.521

Scenario 2

The first part of Table 9 below shows the latency from the Smartbits. At 49 percent load, the 88.80 microseconds of latency were measured between ports 1 and 2, while 91.05 microseconds were measured between ports 4 and 3. As we mentioned above, the latency numbers at 79 percent load do not indicate a latency time, but rather that frames were dropped. At 79 percent load, the 8271s started to drop packets at a rate of five percent between ports 1 and 2 and none between ports 4 and 3.

Table 9: Results from the cell-based test (Smartbits)

Initial Rate(%)	1 to 2 (us)	4 to 3 (us)	Average
	100M-100M	100M-100M	
49%	88.80	91.05	89.925
59%	93.30	94.00	93.650
69%	98.75	91.80	95.275
79%	16330.20	8515.55	12422.875

Initial Rate(%)	1 to 2 (%)	4 to 3 (%)	Average
	100M-100M	100M-100M	
49%	0.000	0.000	0.000
59%	0.000	0.000	0.000
69%	0.000	0.000	0.000
79%	5.283	0.000	2.641

Glossary

ABR – Available Bit Rate

ASIC – Application Specific Integrated Circuit

BUS – Broadcast Unknown Server

CBR – Constant Bit Rate

CLIP – Classical Internet Protocol

CPSW – Control Point Switch

ELAN – Emulated Local Area Network

ESCON – Enterprise Systems CONnection

FPGA – Field Programmable Gate Array

IETF – Internet Engineering Task Force

IISP – Interim Interswitch Signaling Protocol

LANE – LAN Emulation

LEC – LAN Emulation Client

LECS – LAN Emulation Configuration Server

MIB – Management Information Base

MPOA – Multi-Protocol Over ATM

NHRP – Next Hop Resolution Protocol

PCMCIA – Personal Computer Memory Card International Association

PNNI – Private Node–Node Interface

PVC – Permanent Virtual Circuit

QoS – Quality of Service

SVC – Switched Virtual Circuit

UBR – Unspecified Bit Rate

VBR – Variable Bit Rate

VP – Virtual Path

VRRP – Virtual Router Redundancy Protocol