NASA/CR—1999-209412

PNNI Performance Validation Test Report

Robert P. Dimond
GTE Technology Organization, Cleveland, Ohio

December 1999

The NASA STI Program Office . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the Lead Center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA's counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or cosponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that complement the STI Program Office's diverse offerings include creating custom thesauri, building customized data bases, organizing and publishing research results . . . even providing videos.

For more information about the NASA STI Program Office, see the following:

- Access the NASA STI Program Home Page at ***http://www.sti.nasa.gov***

- E-mail your question via the Internet to help@sti.nasa.gov

- Fax your question to the NASA Access Help Desk at (301) 621-0134

- Telephone the NASA Access Help Desk at (301) 621-0390

- Write to:
  NASA Access Help Desk
  NASA Center for AeroSpace Information
  7121 Standard Drive
  Hanover, MD 21076

NASA/CR—1999-209412

PNNI Performance Validation Test Report

Robert P. Dimond
GTE Technology Organization, Cleveland, Ohio

Prepared under Contract NAS3–99175

National Aeronautics and
Space Administration

Glenn Research Center

December 1999

Available from

NASA Center for Aerospace Information
7121 Standard Drive
Hanover, MD 21076
Price Code: A03

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22100
Price Code: A03

PNNI Performance Validation Test Report

Robert P. Dimond
GTE Technology Organization
Cleveland, Ohio 44135

Abstract

Two PNNI neighboring peers were monitored with a protocol analyzer to understand and document how PNNI works with regards to initialization and recovery processes. With the processes documented, pertinent events were found and measured to determine the protocols behavior in several environments, which consisted of congestion and/or delay. Subsequent testing of the protocol in these environments was conducted to determine the protocol's suitability for use in satellite-terrestrial network architectures.

Introduction: Evolution to the PNNI Protocol

Since its inception in the mid 1980's, Asynchronous Transfer Mode (ATM) has positioned itself as the lead networking architecture to offer a reliable method of accommodating the quality of service (QOS) needs of voice, video and data traffic types. While the ability to accommodate real time and non real time network applications was significant, ATM suffered from drawbacks due to administrative overhead and a lack of interoperability between equipment vendors. Because ATM is a connection-oriented technology, virtual circuits or VCs have to be established between every port and switch utilized in the communications path between two endpoints. In the absence of common signaling/routing protocols, this communications path would be made by manually creating Permanent Virtual Circuits or PVC's on each border node in the path, making setup and QOS changes of a circuit less attractive across large ATM networks comprised of multiple equipment vendors.

To address these shortcomings the ATM Forum Technical Committee released document af-pnni-0026.000, the specifications for the Interim Inter-Switch Protocol or IISP. IISP allows the routing of Switched Virtual Circuit (SVC) setup requests from an originating switch to a destination switch, allowing for the setup of SVCs across ATM networks comprised of multiple equipment vendors. The IISP signaling protocol is based upon User Network Interface version 3.1 (UNI 3.1), with a next-hop routing protocol that uses a static table to map an ATM address prefix to a specific switch port. Since there is no method of propagating these table entries, the information must be configured manually on all switches comprising an ATM network. In addition, IISP did not allow for the ability to select a route based upon QOS parameters. The limitations of the IISP protocols resulted in SVC setup calls which would proceed through an ATM network until it either reached the destination switch, or was rejected due to QOS requirements, routing failures, or link state failures. In the event of any failure, the call would result in a RELEASE COMPLETE, in which the call in progress would be torn down at every hop preceding the error, all the way back to the call originator.

In March of 1996, the ATM Forum Technical committee released document af-pnni-0055.000, The Private Network-Network Interface Specification Version 1.0 (PNNI 1.0). This serves as the basis for implementation of the PNNI signaling protocol, a method of establishing a Switched Virtual Circuits across an ATM network. In addition to a signaling protocol, based upon UNI 3.1, the PNNI specification also details the implementation of the PNNI routing protocol. Based loosely upon OSPF, the PNNI routing protocol is the method of distributing the topology information, across switches comprising an ATM network, for use by the PNNI signaling protocol. The routing and switching protocols comprising PNNI offer these important advantages over their predecessors:

- **Supports hierarchical routing** – PNNI Routing domains can be configured as parent or child groups much like the use of IP network subnetting or supernetting Nodes configured for a common hierarchy level and identifier are referred to as members of a peer group. In hierarchical PNNI networks, lower level peer groups in the routing domain are represented logically at higher levels, or parent Peer Groups, when determining routes across a network. This type of representation also greatly reduces the amount of routing information, which must be flooded across the ATM network. The methods of routing summarization and propagation allow the PNNI protocols to scale to very large networks.

- **Supports QOS** – Although derived from a link state protocol OSPF, the PNNI protocol provides QOS resource information for administrative costs, cell delays, cell losses, and available cell rates for each of the five ATM QOS categories. This allows PNNI to make intelligent routing decisions based upon the QOS the connection needs. This information is distributed in PNNI Topology State Elements (PTSEs) bundled in PNNI Topology State Packets or PTSPs. These elements are propagated through the PNNI routing domain at a predetermined length of time, or whenever a significant event has occurred.

- **Utilizes Source Routing** – By having the source select the route through the network, PNNI avoids problems found in next-hop routing protocols such as routing loops due to inconsistencies routing databases, or incompatible routing algorithms.

- **Dynamic Route Alteration –** As a call traverses a PNNI routing domain its originator creates a Designated Transit List or DTL to determine the necessary nodes it must traverse to reach its destination. From the originator's DTL, these nodes can consist of physical nodes to be traversed along a single lower level peer group or logical nodes representing entire peer groups residing within a common parent peer group. When a call is traversing multiple peer groups, the node where the call enters into the peer group creates the DTL to traverse that peer group. Should a failure occur because of a link state change or resource availability change, the PNNI protocol can crankback the call setup to the originator of the last issued DTL. The originator can then create another DTL using an alternate path, or crankback the call further until one of the originators of a previous DTL can successfully reroute the call. This avoids the pitfalls of having to release a call all the way back to the originator if other routes can fulfill the call's requirements.

Because of virtues such as, the ability to propagate routing and QOS information throughout an ATM network, the ability to make routing decisions based upon QOS requirements, and the ability to reroute a call without releasing it back to the originator, the PNNI protocol offers great potential for use in ATM communications architecture comprised of satellite and terrestrial nodes where range and atmospheric conditions can directly impact a link's quality of service.

Purpose

Using the PNNI version 1.0 specification, our goal was to examine the PNNI protocols at the transaction level, learn how the protocols performed those transactions, determine a method of testing the protocols performance, and performing the tests to determine if PNNI will work correctly in a satellite or hybrid network environment, consisting of high speeds and long delays.

Experiment Goals for all PNNI Tests

Essentially, these experiments are setup to witness and document the behavior of the PNNI protocols and their transactions for neighboring PNNI peers (nodes residing in the same peer group, directly connected) under the following conditions:

- Control Conditions with no Delay – Test consist of using two ATM switches running the PNNI protocol with no external network connections. This test will serve to document the PNNI initialization process in ideal conditions.

- Control Conditions with 250ms Delay – Once again two ATM switches running PNNI will be connected exclusively to each other. However, their connection will pass through a delay simulator. This will examine the performance of the protocol over a satellite link.

- Live Network Conditions with no Delay – This consist of connecting an ATM switch running the PNNI protocol to a switch this is an active participant on a large ATM network. This will be to determine the performance of the protocol in real world conditions.

- Live Network Conditions with 250ms Delay – This consist of connecting an ATM switch running the PNNI protocol to a switch this is an active participant on a large ATM network, through a delay simulator, to simulate real world performance over a satellite link.

In each environment, we will be looking for specific events in the PNNI initialization process. According to the PNNI v1.0 specification, PNNI routing initialization consist of the following transactions, which will be observed and measured:

- **The Hello Process** – In this process a both nodes send information such its Node ID, Peer Group ID, Remote Node ID and Remote Peer Group ID (the remote fields are initially set to zero). This information is used to determine whether a neighboring node is in the same peer group and whether or not subsequent steps such as database synchronization and PTSE exchange will occur. The three resulting states are:

    1. **Two-way inside** – Results from both nodes confirming that they belong to the same peer group and that their hello information has been implicitly acknowledged because they find their information in the remote fields of the received hello packets. In this case, Database Synchronization and PTSE Exchange will take place, with the information being propagated to other neighbors in the group.

    2. **Two-way outside** – Results form both nodes confirming that they DO NOT belong to the same peer group and that their hello information has been implicitly acknowledged because they find their information in the remote fields of the received hello packets. In this case the node failed to find a common level in their routing hierarchy. Except for review of subsequent hellos, no further action is taken.

    3. **Common Outside** – Results form both nodes confirming that they DO NOT belong to the same peer group and that their hello information has been implicitly acknowledged because they find their information in the remote fields of the received hello packets. Unlike the Two-way outside scenario, the nodes can advertise to their upnode that the link exists.

- **The Database Synchronization Process** – Is lock-step process in which the ID and originator for all PTSEs are exchanged between the neighboring peers. This is done through a Master/Slave relationship in which the master is determined by the greater Node ID. Here is a simplification of the steps:

    1. **Master sends a database init packet** – packet contains specific sequence number with the "master" "more" and "initialization" bits set to one. This indicates that the slave that the DB Summary process has started.

    2. **Slave sends PTSE information to master** – The slave will acknowledge the masters database initialization requests by sending its PTSE summaries in a database summary packet to the master, with the same sequence number as the masters init packet. Should the summaries require more than one packet, subsequent series of data will be sent after the master sends a packet with an incremented sequence number.

3. **Master sends PTSE information to slave** – The master will acknowledge the slaves data by incrementing its sequence number by 1 in the packets containing data for the slave. Should there be more data to send, the master will set the more bit to one, and wait for the slave to acknowledge the outstanding packet by echoing the masters last used sequence number, before sending subsequent series of data. If all summaries have been sent and acknowledged, the master will send an empty database summary packet with the init and more bits set to zero, and will consider the transactions complete upon receipt of the slaves acknowledgement.

4. **Slave sends final acknowledgement** – Like every other packet it receives, the slave must acknowledge the empty packet by transmitting the matching packet for that sequence. In this case the packet would be consist of the master's empty packet sequence number with the init and more bits set to zero.

The PTSE Exchange Process – With the information received from the database synchronization process, each node looks for corresponding Node ID/PTSE ID entries in its database topology tables. If it cannot find the corresponding entry, a PTSE request is issued. There is no set process other than the each PTSPs fulfilling the PTSE request must only contain PTSEs belonging to one originator. After receiving the needed PTSEs the requestor must acknowledge their receipt by sending a PTSE Acknowledgement. The acknowledgement may contain multiple acknowledgements from multiple nodes or originators, as long as the actual PTSE ACK instances are bundled according to their Originating Node ID. The exchange and acknowledgement of all PTSEs complete the PNNI initialization process

PNNI Signaling Call Complete – In addition to the routing process we will be looking at when a PNNI signaling event "Call Complete". For this event to take place PNNI routing will have to be initialized on all points between the call originator and the call destination, the most apparent sign at PNNI routing is active.

Experiment Procedures for all PNNI Tests

With the equipment indicated, we proceeded to follow the following steps for collecting and interpreting data in our PNNI tests. For each test, 10 captures were performed.

- **Set up HP Internet Advisor** – Configure unit to run in active monitor mode with filters set for PNNI signaling (VC label 0/5) and PNNI routing (VC label 0/18).

- **Set up AdTech SX/14 Delay Simulator –** Where applicable, set unit to 250ms East to West delay as well as the 250ms for the West to East delay to simulate communication between two groundstations connected via satellite.

Configure switches to reside in the same peer group – PNNIv1.0 specifications state, nodes residing in different peer groups will not exchange PNNI routing information directly, but will uplink their information to higher levels of their hierarchy which, if a common peer group is found, will result in the ability to route between the dissimilar groups. Because we wanted to understand the initialization process at its simplest level, we configured our test nodes with an identical peer group identifier.

- **Configure a PNNI-based SVC from ATMSW1 to ATMSW2** – While routing can begin to take place during initialization or changes in large PNNI routing domains, it is unclear when the earliest opportunity for routing to occur on our smaller scale. Although we speculated that routing can occur immediately after the exchange of Horizontal Link Information Group (HLIG) PTSEs (The PTSE type which describes a common link between two peer nodes in terms of QOS parameters), we opted to use when an actual call completion takes place on a pre-configured PNNI-based SVC as the earliest possible time for network use.

- **Reboot ATMSW1** – The PNNI v1.0 specification clearly states that each interface on a switch is a PNNI node. It also states that a link outage will result in a reset of the routing and signaling protocols in the PNNI nodes for both interfaces connected via that link. This allowed us to merely reboot ATMSW1. While we could have opted to disconnect the cables from one of the switches, the reboot sequence ensured that enough time had expired to expire any PTSEs in ATMSW2's topology database.

- **Start HP Internet Advisor** – While ATMSW1 is rebooting this will allow enough time for the device to clears its buffers and begin recording data, ensuring that the earliest events in the initialization process are not missed. The captures lasted for two minutes after ATMSW1 displayed its login prompt, to insure a complete capture of PNNI initializations.

- **Stop HP Internet Advisor**

- **Export data and find occurrences of each process** – Data was exported from the Internet Advisor into a database capable of indexing the events and taking the IA absolute times and calculating delta times between completion of the aforementioned PNNI process.

Equipment utilized for Control PNNI tests (no delay)

While multiple equipment components and configurations were tried, the following items were used in order to analyze a PNNI initialization between two ATM switches:

- Two Cisco LS1010 with OC-3 interfaces running IOS 12.0 patched to 12.0.4
- One HP Internet Advisor WAN (J2300C) with an OC-3 interface running ATM advisor 11.0 patched to level 11.1

Units were set up with their configuration and network connections constant between test runs. The diagram in figure 1 illustrates the initial setup with no delay involved.

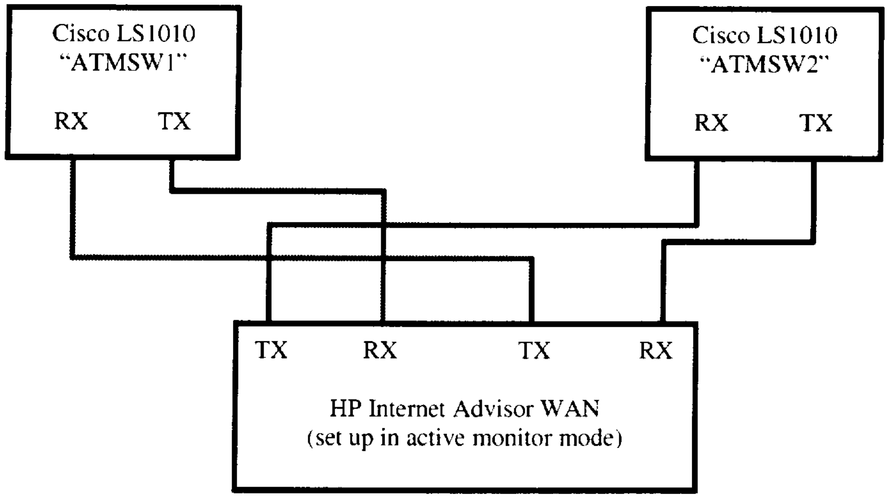

Figure 1.—Equipment setup for capturing ATM traffic (control setup, no delay).

Experiment Results for Control PNNI tests (no delay)

The following results were observed in conducting out tests:

Absolute Times for Controlled Benchmark Tests No Delay

Test	2-Way Inside Achieved	Database Sync'd	PTSEs Exchanged	PNNI SVC Creation
1	1.270376	2.271804	3.3514383	5.4383557
2	1.3247219	2.3264068	3.4045185	5.4422744
3	.9997622	1.9981576	3.0780246	5.1327254
4	1.0014377	2.001794	3.0796575	5.1134101
5	1.0913872	2.092738	3.1738224	5.236591
6	1.2672618	2.2680489	3.3472539	5.4380211
7	1.4002114	2.4000758	3.4794587	5.5344226
8	.9991141	2.000663	3.0813059	5.1367305
9	1.0019297	2.0023069	3.082378	5.1369432
10	1.4108342	2.4136035	3.4926239	5.5393456
AVG	1.1767036	2.1775598	3.2570481	5.3148819

Delta Times for Controlled Benchmark Tests No Delay

Test	2-Way Inside Achieved	Database Sync'd	PTSEs Exchanged	PNNI SVC Creation
1	1.270376	1.0014444	1.0796179	2.0868173
2	1.3247219	1.0016848	1.0781117	2.0377559
3	.9997622	.9983954	1.079867	2.0547008
4	1.0014377	1.0003563	1.0778634	2.0337526
5	1.0913872	1.0013508	1.0810844	2.0627686
6	1.2672618	1.0007871	1.079205	2.0907672
7	1.4002114	.9998644	1.0793829	2.0549639
8	.9991141	1.0015489	1.0806429	2.0554246
9	1.0019297	1.0003772	1.0800711	2.0545652
10	1.4108342	1.0027692	1.0790204	2.0467217

Our absolute times show that all calls were able to take place within 5.5 seconds on our control network with no delay. The delta times show that the control setup provided repeatable results with only very minute variations, which can most likely be explained by the timing of the processes on the ATM switches.

Equipment utilized for Control PNNI tests (250ms delay)

While multiple equipment components and configurations were tried, the following items were used in order to analyze a PNNI initialization between two ATM switches:

- Two Cisco LS1010 with OC-3 interfaces running IOS 12.0 patched to 12.0.4
- One HP Internet Advisor WAN (J2300C) with an OC-3 interface running ATM advisor 11.0 patched to level 11.1
- One AdTech SX/14 Delay simulator

Units were set up with their configuration and network connections constant between test runs. The diagram in figure 2 illustrates the initial setup with the delay simulator in line.

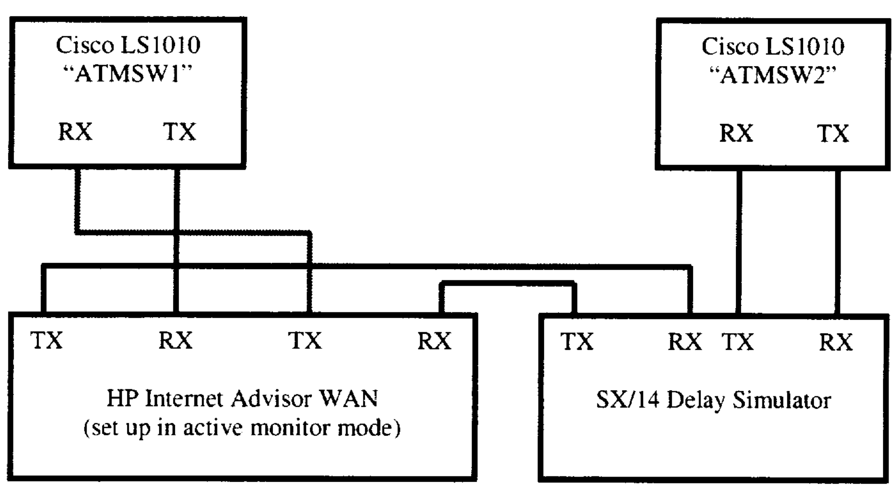

Figure 2.—Equipment setup for capturing ATM traffic (control setup, 250ms delay).

Experiment Results for Control PNNI tests (250ms delay)

The following results were observed in conducting out tests:

Absolute Times of Controlled Benchmark Test with 250ms Delay

Test	2-Way Inside Achieved	Database Sync'd	PTSEs Exchanged	PNNI SVC Creation
1	1.629395	3.1941977	4.7730553	6.8331295
2	1.4980075	3.0014934	4.597129	7.136171
3	1.6941739	3.1941331	4.7742565	7.373022
4	1.6594261	3.1621575	4.7398	7.3221899
5	1.5678668	3.0700849	4.6482123	7.2106451
6	1.4981626	2.99876	4.5775795	7.1400966
7	1.499303	3.0024638	4.582431	7.1451925
8	1.541992	3.0420855	4.648709	7.179246
9	1.5015733	3.0021478	4.581903	7.11532
10	1.501553	3.0023777	4.5798259	7.17787
AVG	1.55914532	3.0669901	4.6502901	7.1632882

Delta Times of Controlled Benchmark Test with 250ms Delay

Test	2-Way Inside Achieved	Database Sync'd	PTSEs Exchanged	PNNI SVC Creation
1	1.629395	1.5012582	1.5788576	2.0600742
2	1.4980075	1.5034859	1.5956355	2.539042
3	1.6941739	1.4999592	1.5801234	2.5987655
4	1.6594261	1.5027314	1.5776425	2.58239
5	1.5678668	1.5022181	1.5781274	2.5624328
6	1.4981626	1.5017134	1.5777035	2.562517
7	1.499303	1.5031607	1.5799672	2.5627614
8	1.541992	1.5000935	1.6066235	2.5305369
9	1.5015733	1.5005745	1.5797552	2.533417
10	1.501553	1.5008247	1.5774481	2.598044

All calls being placed over this delayed network are taking place within 7.5 seconds. With the 250ms delay one would expect the call setup to be approximately 5.25 seconds longer, based upon the extra time for each packet to reach the other end of the link. This is not the case because some of these processes take place concurrently on both ends. The HELLO, and PTSE exchange processes have transactions that are independent of their neighbor. Once again, we see a fair amount of similarity in the delta times between events.

Equipment utilized for Live Network PNNI tests (no delay)

While multiple equipment components and configurations were tried, the following items were used in order to analyze a PNNI initialization between two ATM switches:

- One Cisco LS1010 with OC-3 interface, running IOS 12.0 patched to 12.0.4 to be used ATMSW1

- One Fore ASX1000 running ForeThought v5.2.0 to be used as ATMSW2
- One HP Internet Advisor WAN (J2300C) with an OC-3 interface running ATM advisor 11.0 patched to level 11.1

The units were set up with their configuration and network connections constant for ATMSW1 but ATMSW2 was also hooked into the NREN network, which had signaling enabled at the time. ATMSW2 was configured as a Gateway domain to utilize both the ATMF and ForeThought PNNI protocols.

The diagram in figure 3 illustrates how ATMSW1 was hooked into the ASX1000.

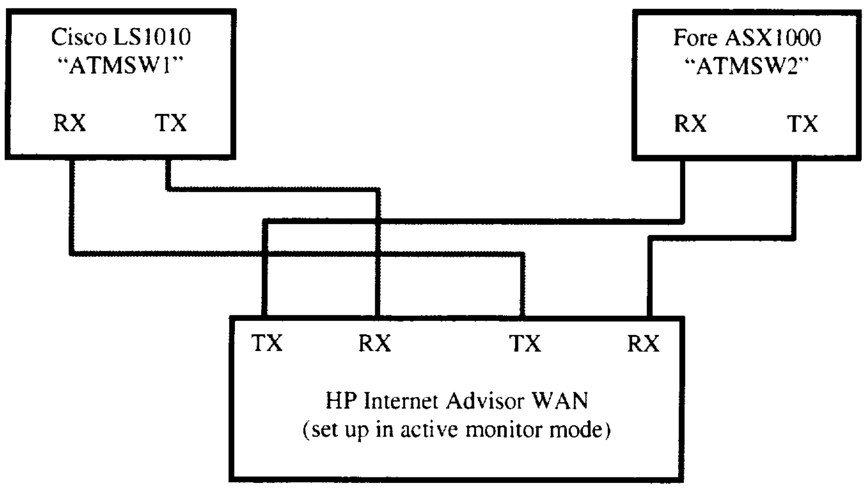

Figure 3.—Equipment setup for capturing ATM traffic (live network setup, no delay).

Experiment Results for Live Network PNNI tests (no delay)

The following results were observed in conducting out tests:

Absolute Times of Live Benchmark Tests No Delay

Test	2-Way Inside Achieved	Database Sync'd	PTSEs Exchanged	PNNI SVC Creation
1	1.3431769	1.8450259	4.3417151	6.4084682
2	1.0855029	1.5868395	4.0847737	6.0972099
3	2.9726585	3.4735137	5.9717264	8.0253593
4	3.4650046	3.9569697	6.4544312	8.5308122
5	20.0013074	20.5030062	23.0003360	25.10085746
6	1.972823	2.9709369	4.4726357	6.5202186
7	1.1394578	1.6414967	3.6406292	5.7227762
8	1.0804916	1.5822726	4.0796885	6.1365336
9	1.2022865	1.7034022	4.2013033	6.25427893
10	1.1802478	1.6819977	4.1793671	6.223996
AVG	1.5441649	2.0442454	4.1426270	6.6114611

Delta Times of Live Network Tests No Delay

Test	2-Way Inside Achieved	Database Sync'd	PTSEs Exchanged	PNNI SVC Creation
1	1.3431769	.501849	2.4966892	2.0667531
2	1.0855029	.5013366	2.4979342	2.0124362
3	2.9726585	.5008552	2.4982127	2.0536329
4	3.4650046	.4919651	2.4974615	2.076381
5	20.0013074	.5016988	2.4973298	2.1005214
6	1.972823	.9981139	1.5016988	2.0475829
7	1.1394578	.5020389	1.9991325	2.082147
8	1.0804916	.501781	2.4974159	2.0568451
9	1.2022865	.5011157	2.4979011	2.0529756
10	1.1802478	.5017499	2.4973694	2.0446289

Factoring out the abnormal values of test #5, the times for call completion still varies greatly from 5.7 to 8.5 seconds. These differences may be due to the switch load, or interoperability issues between Cisco and Fore's implementation of the ATMF PNNI protocol. In fact, it was noted during the tests that the two vendors use different PTSE sequence number schemes, although both schemes were within ATMF specifications.

When one looks at the delta times, one can note more similarities in event times between tests. Note the delta times for the PTSEs Exchanged event. In most cases, the event takes 1.4 seconds longer than the same event in the control test with no delay. Once again it is unknown if this lag due to switch load, or FTPNNI PTSE Table size, or if the lag is due to differences in vendor implementation of the protocol. Another unexplainable phenomenon is the shorter times for the Database Sync'd event.

Equipment utilized for Live Network PNNI tests (250ms delay)

While multiple equipment components and configurations were tried, the following items were used in order to analyze a PNNI initialization between two ATM switches:

- One Cisco LS1010 with OC-3 interface, running IOS 12.0 patched to 12.0.4 to be used ATMSW1
- One Fore ASX1000 with OC-3 interface running ForeThought v5.2.0 to be used as ATMSW2
- One HP Internet Advisor WAN (J2300C) with OC-3 interface running ATM advisor 11.0 patched to level 11.1
- One AdTech SX/14 Delay simulator

The units were set up with their configuration and network connections constant for ATMSW1 but ATMSW2 was also hooked into the NREN network which had signaling enabled at the time. ATMSW2 was configured as a Gateway domain to utilize both the ATMF and ForeThought PNNI protocols.

The diagram in figure 4 illustrates how ATMSW1 was hooked into the ASX1000.

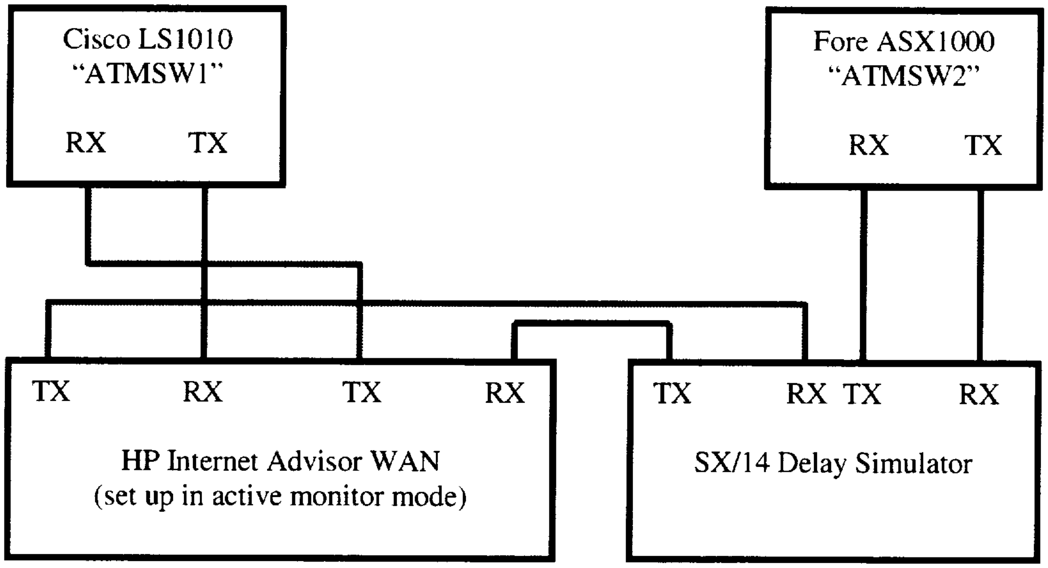

Figure 4.—Equipment setup for capturing ATM traffic (live network setup, 250ms delay).

Experiment Results for Live Network PNNI tests (250ms delay)

The following results were observed in conducting out tests:

Absolute Times of Live Network Tests 250ms Delay

Test	2-Way Inside Achieved	Database Sync'd	PTSEs Exchanged	PNNI SVC Creation
1	1.5935471	2.7008993	5.6010933	8.1680453
2	2.622147	3.7591883	6.8554613	9.6176632
3	2.4995214	3.5968604	6.697714	9.255183
4	3.9600012	5.0359014	7.9355423	10.5293633
5	1.783214	2.960063	5.9220703	8.3397623
6	1.5984582	2.7369593	5.6772457	8.4463962
7	1.3994821	2.5494878	5.5523199	8.0481799
8	1.77343	2.8732372	5.7685372	8.514646
9	1.6365397	2.7551319	5.800114	8.700165
10	1.7207131	2.7979431	5.6976182	8.2381002
AVG	2.0587053	3.1765671	6.1507716	9.607094

Delta Times of Live Network Tests 250ms Delay

Test	2-Way Inside Achieved	Database Sync'd	PTSEs Exchanged	PNNI SVC Creation
1	1.5935471	1.1073522	2.9001940	2.566952
2	2.622147	1.1370413	3.096273	2.7622019
3	2.4995214	1.097339	3.1008536	2.557469
4	3.9600012	1.0759002	2.8996409	2.593821
5	1.783214	1.176849	2.9620073	2.417692
6	1.5984582	1.1385011	2.9402864	2.7691505
7	1.3994821	1.1500057	3.0028321	2.49586
8	1.77343	1.0998072	2.8953	2.7461088
9	1.6365397	1.1185922	3.0449821	2.900051
10	1.7207131	1.07723	2.8996751	2.540482

The results indicate that PNNI initialization to the point where calls can be routed takes an average of 9.6 seconds. The average call setup time between our live network tests differ by 3 seconds. That time is somewhat longer than expected considering the average difference for the same event between our two control tests differ by only 1.8 seconds. The delta times for the events also indicate that the protocol initialization results are not repeatable. Once again it is unknown if some or all of those differences are due to switch load for the NREN connected switch.

Average Absolute Times for Events for All Tests

Test	2-Way Inside Achieved	Database Sync'd	PTSEs Exchanged	PNNI SVC Creation
Control No Delay	1.1767036	2.1775598	3.2570481	5.3148819
Control 250ms Delay	1.55914532	3.0669901	4.6502901	7.1632882
Live No Delay *	1.5441649	2.0442454	4.1426270	6.6114611
Live 250ms Delay	2.0587053	3.1765671	6.1507716	9.607094

Conclusion

We have achieved the primary goal of documenting and measuring the PNNI initialization process.

In general, our results were as expected, the delay for our test compounds in the lock step transactions used in the protocol. That is, delaying a packet that is required by ATMSW1 before it can move forward in the initialization process will obviously delay ATMSW1's initialization. In turn ATMSW1's responses, awaited by ATMSW2, will also be delayed postponing the initialization process even further. Some instances such as the initial hellos, the PTSE Requests, and the keepalive hellos are sent as a result of internal processes are not directly impacted by delays in packet reception. It is speculated that the compounding of these delays could pose problems for hybrid networks consisting of multiple peer group nodes all connected via satellite links. The utilization of PNNI over multiple long delay paths might require changes to the specification to make the protocol more delay tolerant.

With regards to simpler hybrid networks, the protocols quick initialization times show promise for Low Earth Orbit Satellite (LEO) applications in which an ATM switch is connected to a groundstation system tracking two LEOs and would have switch between a LEO moving out of range and its successor that is moving into range. More importantly, the ability of the protocol to make those decisions based upon factors influenced by satellite communications such as Bit Error Rate (BER) and Cell Transfer Delay (CTD) would allow the change in routing from one LEO to the other to take place in a timely fashion.

While we achieved our initial objectives, our tests are far from conclusive on this protocol. Further tests could evaluate the protocols initialization and recovery characteristics within different levels of a large PNNI hierarchy, or testing how the protocol functions in a LEO situation, which is characterized by changing Cell Transfer Delays and Bit Error Rates. In addition, more interoperability tests should be performed. Although the Cisco equipment used in out tests was fully PNNIv1.0 compliant, other vendors have only recently introduced versions of the PNNI protocol, which claim to be fully ATMF compliant. Out tests alone indicated several minor vendor problems that had to be addressed because the protocol was not adhering to the written specifications. Even our analysis equipment had to be patched to address problems in the initialization decodes which did not appear to be according to specification.

References

The ATM Forum Technical Committee: Private Network-Network Interface Specification Version 1.0.

Metz, Christopher Y. and Sackett, George C.: ATM and Multi-protocol Networking.

Ibe, Oliver C.: Essential of ATM Networks and Services.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
December 1999

3. REPORT TYPE AND DATES COVERED
Final Contractor Report

4. TITLE AND SUBTITLE
PNNI Performance Validation Test Report

5. FUNDING NUMBERS
WU–632–50–5A–00
NAS3–99175

6. AUTHOR(S)
Robert P. Dimond

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
GTE Technology Organization
21000 Brookpark Road
Cleveland, Ohio 44135

8. PERFORMING ORGANIZATION REPORT NUMBER
E–11974

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
John H. Glenn Research Center at Lewis Field
Cleveland, Ohio 44135–3191

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
NASA CR—1999-209412

11. SUPPLEMENTARY NOTES
Project Manager, William D. Ivancic, Communications Technology Division, NASA Glenn Research Center, organization code 5610, (216) 433–3494.

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified - Unlimited
Subject Category: 17 Distribution: Nonstandard

This publication is available from the NASA Center for AeroSpace Information, (301) 621–0390.

12b. DISTRIBUTION CODE

13. ABSTRACT ***(Maximum 200 words)***

Two PNNI neighboring peers were monitored with a protocol analyzer to understand and document how PNNI works with regards to initialization and recovery processes. With the processes documented, pertinent events were found and measured to determine the protocols behavior in several environments, which consisted of congestion and/or delay. Subsequent testing of the protocol in these environments was conducted to determine the protocol's suitability for use in satellite-terrestrial network architectures.

14. SUBJECT TERMS
Satellite; ATM; Network; PNNI; Protocol

15. NUMBER OF PAGES
20

16. PRICE CODE
A03

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102