

# 45 Configuring and Monitoring PNNI

The Private Network-to-Network Interface (PNNI) is the routing protocol for building world-wide ATM networks. It uses a hierarchical structure to efficiently distribute information about ATM network topology (i.e., ATM nodes and links), compute routes that meet requested bandwidth and Quality of Service (QoS) requirements, establish connections between ATM End Systems across the network, and scale from small local networks to large global networks.

The OmniSwitch with ATM switching functionality supports PNNI version 1.0. In addition, it supports the Interim Inter-Switch Signalling Protocol (IISP); see Chapter 46 for information on IISP. Both PNNI and IISP are ATM routing protocols that support interoperable multi-vendor ATM networks.

IISP is a static routing protocol in which addressing information about ATM devices is statically defined at each ATM switch interface, and signalling is used to forward connection messages hop-by-hop through the network. In contrast, PNNI is a dynamic routing protocol that intelligently determines routes through the network based on connection requirements, and uses source routing to forward call setup requests through the network.

This chapter is for users with a basic familiarity of the PNNI routing protocol. It provides a brief overview of PNNI, but it focuses on the implementation of PNNI within the OmniSwitch. More extensive information on PNNI can be found in the ATM Forum PNNI 1.0 specification.

## OmniSwitch PNNI Self-Configuration

The OmniSwitch implementation of PNNI is self-configuring while also providing commands that configure a diverse range of parameters. This chapter provides instructions on how to configure PNNI within the OmniSwitch at the global, node, and port level. Commands are provided that allow you to configure topology metrics, such as Administrative Weight, for each Class of Service supported on a port. In addition, various timers can be configured. Several other commands display extensive configuration information, status information, and statistics on neighboring nodes, adjacent links, and topology information packets.

## Load Balancing

The OmniSwitch supports load balancing, which is a technique for distributing traffic across the ATM network. In load balancing, traffic is distributed between two or more paths. As a result of load balancing, greater throughput is achieved across the network.

## Summarization and Reachability

PNNI allows address reachability information to be summarized, reducing the number of addresses that must be stored in topology databases. OmniSwitch PNNI nodes advertise a single address prefix to represent a group of ATM End System or node addresses that share a common prefix. See *Summarization and Reachability* on page 45-17 for more information.

---

### **PNNI Path Limitations**

The OmniSwitch implementation of PNNI supports at most 5 parallel paths between any two nodes. These paths can be physical CSM links or logical links (i.e., VP tunnels). Using factory-configured defaults, any one node, or OmniSwitch, should not have more than 30 neighbor nodes. If a node has more than 30 neighbor nodes, then performance may be affected unless you increase the hello or shortest path first timers. In addition, the maximum diameter of an OmniSwitch PNNI network should not exceed 50 nodes. If a larger diameter is required, then IISF links should be used.

# PNNI Configuration

The OmniSwitch implementation of PNNI is self-configuring. In releases prior to 4.1, the only user configuration required is to specify a CSM port as a PNNI port using the **map** command or through SNMP-based network management software.

In Release 4.1 and later, you can enable CSM auto-port configuration, which does not require user configuration. If CSM auto-port configuration has been enabled, then two CSM ports connected together will automatically become PNNI ports. You must disable CSM auto-port configuration to manually configure a CSM port. See Chapter 42, “Managing Cell Switching Modules (CSMs)” for more information on CSM auto-port configuration.

OmniSwitch PNNI contains carefully chosen defaults that enable you to bring up a PNNI node and begin communicating immediately. If you wish to fine tune the default configuration, several commands are available to configure PNNI parameters at the general, node, and port levels. Instructions for using these configuration commands begin on page 45-31. Other important PNNI variables, such as timer values and topology metrics, have defaults that are reasonable for most network environments.

## PNNI Port Type Configuration

The only step required by the user is to set up a CSM port as a PNNI port. Each CSM port by default supports the User-to-Network Interface (UNI). In order for a CSM port to support PNNI, it must be configured as a PNNI port. If the port is configured as an IISP port, or any port type other than PNNI, then PNNI will not run on that port. You make this change through the **map** command by changing the **Port Type** variable from UNI to PNNI. The **map** command is described in Chapter 42, “Managing CSM Modules (CSMs).”

In Release 4.1 and later, no reset is required for changing port type. In releases prior to 4.1, if you change the ATM port type from a PNNI port back to a UNI port (or vice versa) you *must* reset the switch if the port has already been used before.

## Running PNNI Software

Different sets of image files are required depending on whether you need to configure just a single-peer group network or if you also need to configure a multiple-peer group network. To configure a single-peer group network, you must load the following files into your switch:

**asm.img** (MPM and FCSM) *or* **asmc.img** (MPM-C)  
**cell.img**

To run the multiple-peer group version of PNNI, you must load the following files into your switch:

**asm\_mpg.img** (MPM and FCSM) *or* **asmc\_mpg.img** (MPM-C)  
**cell\_mpg.img**

### ◆ Note ◆

If you are using the multiple-peer group files listed above and you want to run LES/BUS software, you *must* use the **lsm\_mpg.img** image file instead of the **lsm.img** image file.

In addition, you must add the following line to your **mpm.cmd** file if you load the **asm\_mpg.img** (or **asmc\_mpg.img**), **cell\_mpg.img**, and **lsm\_mpg.img** files:

```
atm_load_mpg=1
```

You *must* put this line before the **cmInit** line and you must reboot the switch to implement the change. In the single-peer group version of PNNI, either change this line to **atm\_load\_mpg=0** or leave it out of the **mpm.cmd** file. See Chapter 11, “Managing Files,” for information on editing the **mpm.cmd** file.

### Loading the PNNI Module

PNNI functionality is built into the **cell.img** file in single peer group configurations and in the **cell\_mpg.img** file in multiple peer group networks. Once **cell.img** or **cell\_mpg.img** file is placed into the OmniSwitch’s flash memory file system, the switch will dynamically load the file once it senses the presence of a Cell Switching Module (CSM). During the initialization of the cell switching code, PNNI is enabled or disabled. Once the PNNI protocol is running in system memory, self-configuration begins and nodes can begin exchanging Hello packets, building topology databases, and establishing data connections (provided the associated port has been set to type PNNI). If the port type is set to any value other than PNNI, then PNNI will forward on that port based on user-defined static routes.

### Default ATM Address

To allow PNNI nodes to communicate without user intervention, each node is automatically assigned a default ATM address at system start-up. This address includes a prefix of **3903488001bc90000101** followed by Alcatel-specific OID information. Other important PNNI identifiers, such as the Node IDs and Peer Group IDs, can automatically be derived from this default ATM address. You can optionally change this ATM address later through the **pncfg** command (described on page 45-39).

#### ♦ Important Note ♦

The default ATM address for the node is locally defined and should not be used when attaching to a public network. To obtain a unique NSAP address for your enterprise consult RFC 1629 or ANSI.

In PNNI hierarchy levels 103 (hexadecimal 67) and lower in a multiple-peer group network, the Node ID and ATM address can be the same. In levels 104 (hexadecimal 68) and higher, the Node ID and ATM address *must* be different.

### Static Routes/IISP

Static routes must be user configured. PNNI is a dynamic routing protocol that can determine path selection on its own, so static routes are optional. If static routes or the Interim Inter-Switch Signalling Protocol (IISP) are necessary, then use the Route Management menu (described in Chapter 46, “Managing IISP and PNNI Routes”) for configuration.

## Elements of a PNNI Network


PNNI relies on a hierarchical model to distribute information throughout the network about all ATM nodes and links. Each node is part of one or more *peer groups* that serve as its local routing domain. All nodes in a single peer group continuously update each other on their current status.

Each node shares a common map of the peer group, including information on how nodes are connected, which End Systems (or the address summaries representing those devices) are attached to which nodes, and the topology attributes supported at each node. This information is kept in each node's *Topology Database*, which contains individual entries referred to as *PNNI Topology State Elements (PTSEs)*. Whenever there is a change within the peer group, all nodes are updated through regular exchanges of topology information. Topology information is exchanged between nodes via *PNNI Topology State Packets (PTSPs)*.

### ◆ More Information on Topology Exchanges? ◆

A more detailed discussion of PNNI topology information exchanges is provided in *PNNI Network Initialization* on page 45-18. The packet types used in the PNNI protocol are described in *PNNI Packet Types* on page 45-14.

Within a single peer group, nodes are connected by *horizontal links*, which are physical links between the nodes. A *Designated Transit List (DTL)* is a source route between nodes through a peer group. A DTL provides a complete source route through the peer group and will be used to route call requests through the peer group.



### A Single PNNI Peer Group

PNNI has the capability to support multiple peer groups. The OmniSwitch supports multiple peer group configurations in release 4.1 and later.

### Multiple Peer Group Networks

In Release 4.1 and later, multiple peer group configurations, where information on each peer group is summarized by *Peer Group Leader (PGL)* nodes and aggregated into virtual ATM nodes known as *Logical Group Nodes (LGNs)*, are supported. These Logical Groups Nodes distribute summary information on their peer groups to other LGNs. Depending on the size of the network, there may be several levels of low-level nodes and Logical Group Nodes.

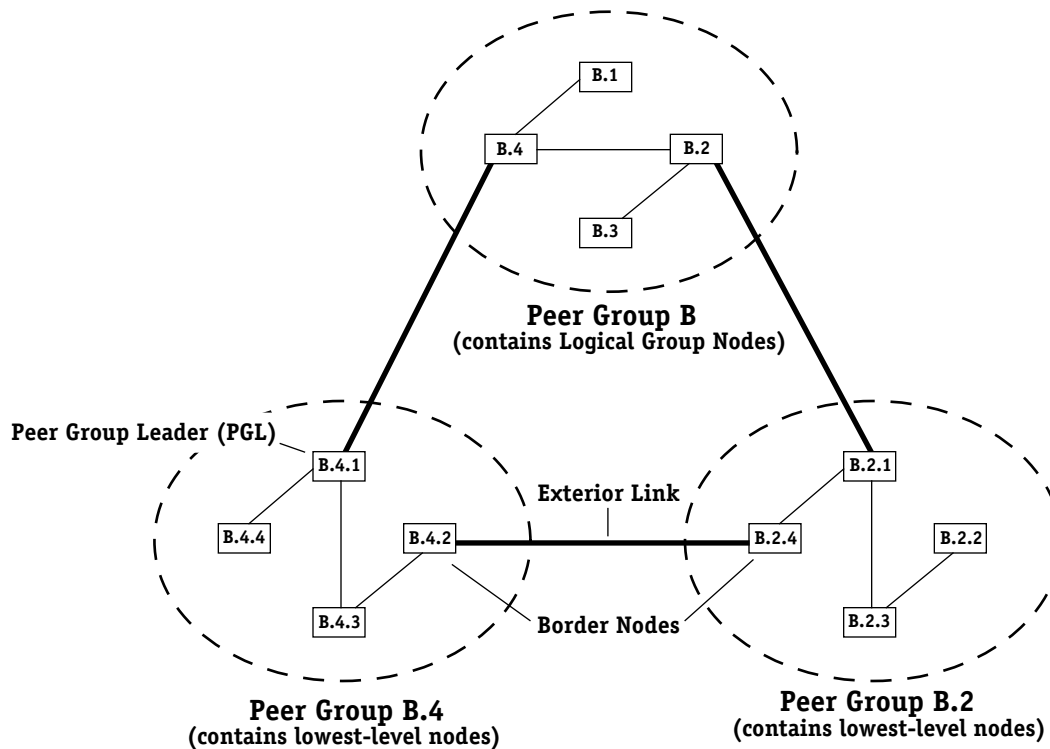
#### ◆ Note ◆

Multiple-peer group operation is optional feature and it is not included as part of the standard **cell.img** and image files. Instead, you must load the image files designed for this feature. See *Running PNNI Software* on page 45-3 for more information on multiple-peer group PNNI files.

The level of a peer group or node indicates the position of that peer group or node within the PNNI hierarchy. PNNI levels range from 0 to 104. A lower PNNI level number indicates a node that is higher up in the PNNI hierarchy (greater summarization of PNNI network topology information). A higher PNNI level number indicates a node that is lower in the hierarchy. Nodes that are located at the lowest-level of a particular chain in the PNNI hierarchy are referred to as *lowest-level nodes*.

In a single peer group network, just one peer group is supported and there is just one level in the PNNI hierarchy. The default number for this level is 80 decimal (50 hexadecimal). You must configure a node to operate at a specific level to implement multi-peer group operation.

Another type of node in a multiple peer group configuration is a *border node*. Border nodes lie on the edges of their respective peer groups and form physical links with border nodes in other peer groups. However, border nodes do not summarize information on their peer groups; summarization is left to the Peer Group Leader. Instead border nodes inform other nodes in their peer group about their access to another peer group. These other nodes may use this link to set up connections to the neighboring peer group (or *through* the neighboring peer group). The border node's link is referred to as an *exterior link*.



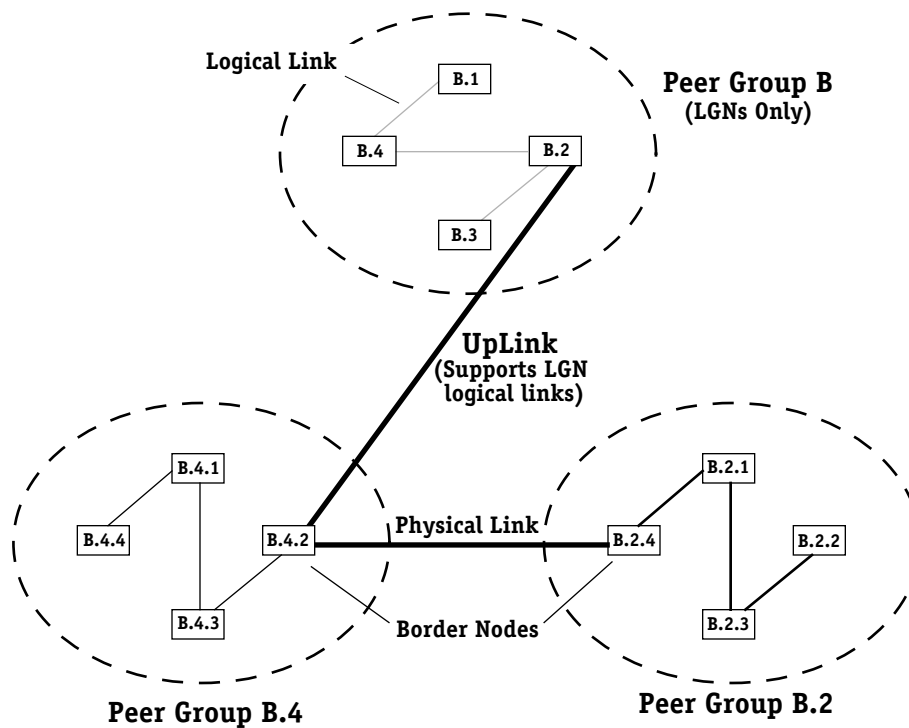
### A PNNI Network With Multiple Peer Groups

Logical Group Nodes (LGNs) are virtual PNNI entities that aggregate information on an entire peer group. They are not physical nodes the way lowest-level PNNI nodes are physical nodes. In addition, LGNs do not have actual physical links with each other or with their child nodes. Instead, LGNs have logical links with each other that rely on physical links between border nodes in neighboring peer groups. These logical links between LGNs representing one peer group and border nodes in another peer group are called *uplinks*. The border nodes supporting these uplinks are called *upnodes*.

There must be a physical means for transporting topology information between LGNs. These LGNs depend on border nodes in other peer groups to provide information on the neighboring peer group. The example below shows a border node in peer group B.4 (node B.4.2) with a logical uplink to the LGN for peer group B.2. The LGN in peer group B.2 relies on this link to obtain topology information on peer group B.4.

#### ◆ Note ◆

If you have a network with more than 30 PNNI nodes, you should upgrade the SIMM memory in your MPM II or MPM 1G to at least 32 MB.



**Uplink in a PNNI Network**


### Peer Group Leader (PGL) Election Algorithm

The nodes in a logical peer group use an election process to determine which node will be the peer group leader (PGL). A node participates by advertising its leadership priority to all the nodes in the peer group through PTSEs. Each node will elect the node with the highest non-zero PGL priority as PGL.

The peer group must have at least two nodes with non-zero PGL priorities for an election to take place. All nodes will participate in the election process unless their leadership priority and preferred PGL is set to zero. If the nodes cannot reach a unanimous decision, then a 2/3 vote will suffice. In the case of a tie, the node ID will be used as a tie breaker. Once a node is elected PGL, its leadership priority is incremented to ensure stability and prevent future election deadlocking.

As shown in the figure on the following page, all nodes, except for Node B (which does not participate in the election process), advertise its leadership priority through PTSEs. Node C, which has the highest leadership priority, will be elected PGL unanimously since Nodes A and D will elect Node C PGL and Node C will elect itself PGL.





### How Nodes Elect a Peer Group Leader (PGL)

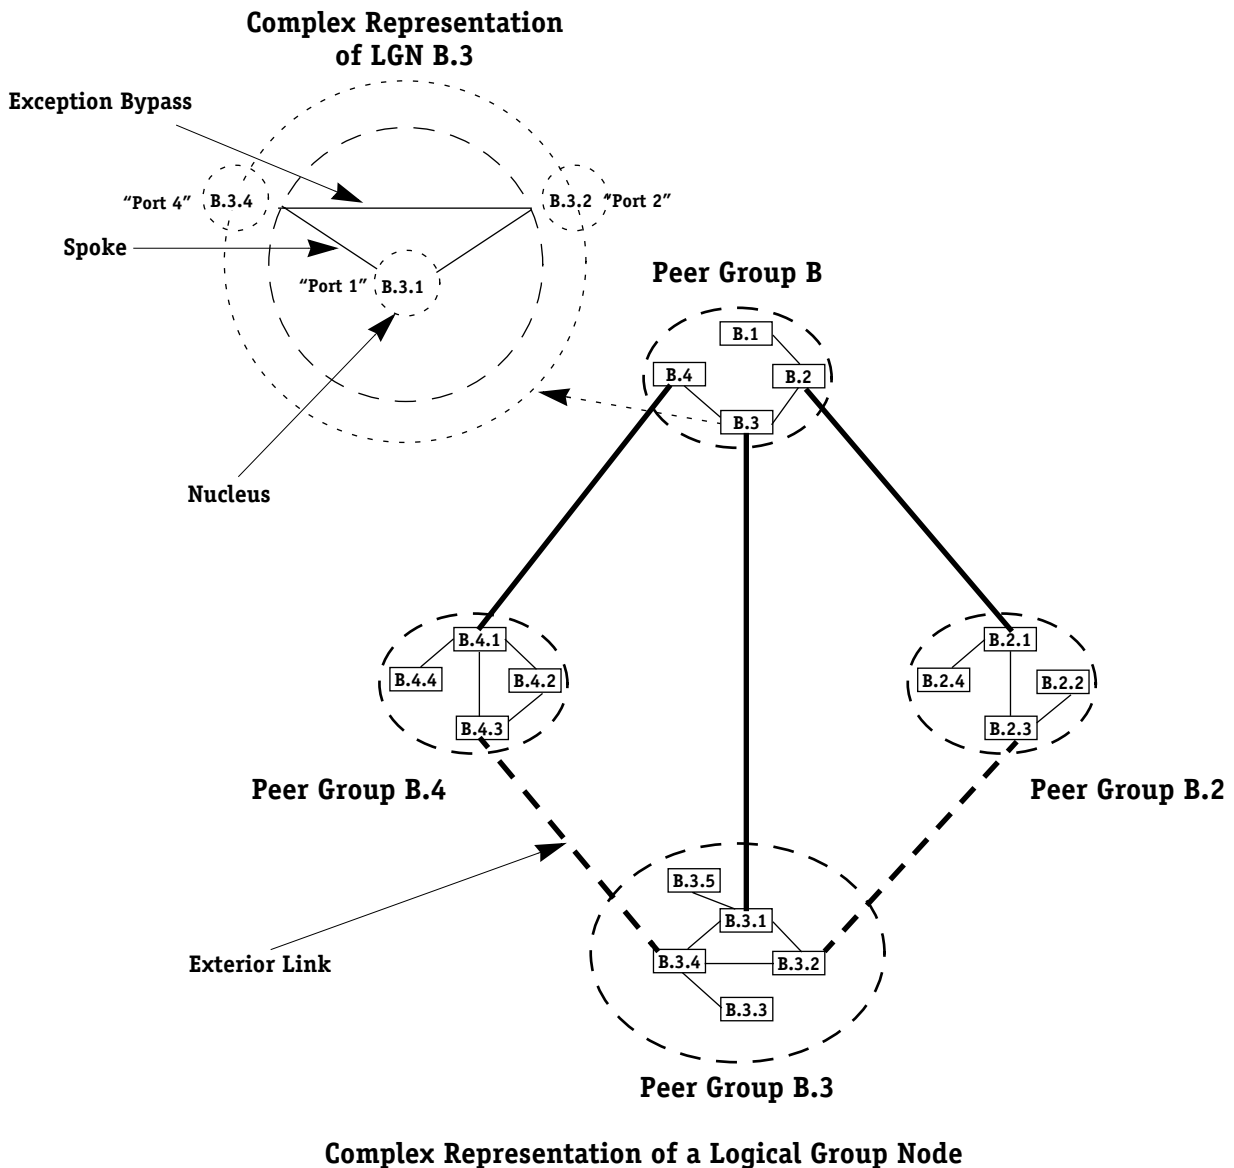
The PGL election algorithm is continuously running process. If a node that is PGL fails or loses its connectivity, then the node with the next-highest leadership priority will become PGL. If a node with a higher priority leadership priority is added to the peer group, it will become the new PGL.

You can configure the PGL priority of a node with the **pncfg** command, which is described in *Configuring Node-Specific Parameters* on page 45-39. To view the current leadership priority of a node, use the **pninfo** command, which is described in *Viewing Node-Specific Information* on page 45-52.

### Complex Representation

Topology aggregation is the process used in PNNI to summarize the topology information of a child peer group to its parent group. This summarization can greatly reduce the volume of information advertised in the parent peer group. Thus, effective topology aggregation is a necessity for PNNI to scale to large multiple-peer group networks. In PNNI, this concept of topology aggregation is known as *complex node representation*.

The figure on the following page shows the complex representation of logical group node (LGN) B.3 using a symmetric star topology. The Peer Group Leader (PGL) node, B.3.1, at the center of the star is referred to as the nucleus. The logical connections between the nucleus and another node in the peer group as known as a spoke.



In most cases, peer groups are not symmetrical. Therefore, exceptions can be used to represent the connectivity of nodes that is significantly different from the default. For example, a node in a peer group could be in a different location from other nodes in a peer group. In complex representation, the local nodes would have spokes to the nucleus with default attributes whereas the remote node would have a spoke with exception attributes to the nucleus. A bypass exception represents connectivity between two nodes that bypasses the nucleus because it is more efficient.

You can configure complex representation with the `pncfg` command, which is described in *Configuring Node-Specific Parameters* on page 45-39. To determine if a node is using complex representation, use the `pninfo` command, which is described in *Viewing Node-Specific Information* on page 45-52.

# PNNI Identifiers

Node ATM addresses are important in the PNNI identification scheme. But PNNI requires additional methods of identification to describe its hierarchical structure and its peer groups. The following descriptions highlight some key PNNI identifiers that you are likely to encounter while configuring and monitoring your PNNI network.

## Level Identifier

The level identifier is the level within the PNNI hierarchy where a node exists. This value may range from 0 to 104, with higher values indicating nodes lower in the PNNI hierarchy. This level is used to determine the default Node ID and the default Peer Group ID for a node. (The default node level is 80 decimal.) In a single-peer group configuration, all nodes will be at the same level.

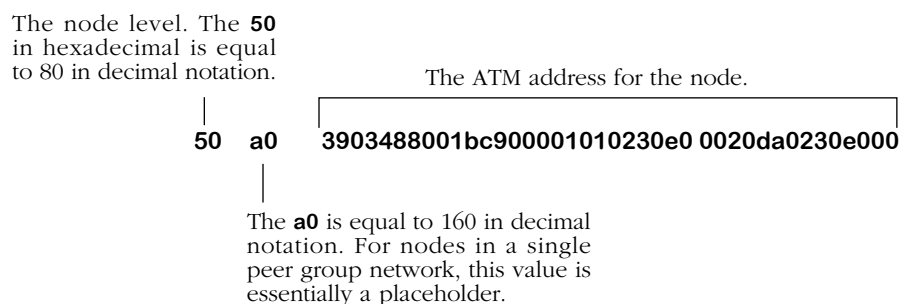
## Node ID

The Node ID is a 22-octet identifier for an OmniSwitch node within the PNNI network. The first octet consists of the level of the node within the PNNI hierarchy. By default, OmniSwitch nodes reside at level 80 decimal.

### ◆ Important Note ◆

In releases previous to 3.2, the default node level was 96 decimal. In release 3.2 and later the default level is 80 decimal. Therefore, if you have an existing OmniSwitch network with PNNI nodes residing at the old default level, then you will have to reconfigure the Node IDs for either the new nodes or the old nodes if you want all nodes to be in the same peer group.

In a single-peer network, the next octet equals 160, and the remaining 20 bytes consist of the ATM address of the node itself. The following is an example of a single-peer network Node ID expressed in hexadecimal notation:

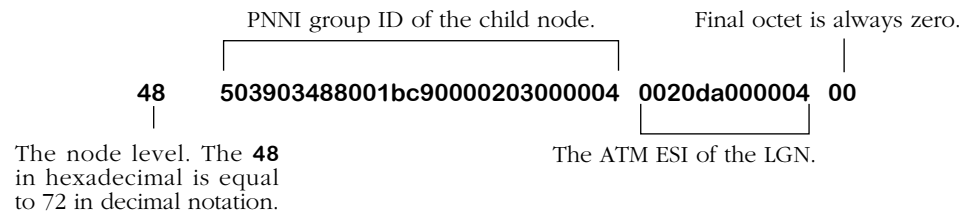


### Typical Node ID in Single Peer Group Network

Logical Group Nodes (LGNs) use a slightly different Node ID structure. The first octet is the node level, the next 14 octets are equal to PNNI node peer group ID of the child node whose election as the Peer Group Leader (PGL) produced this LGN, the next 6 octets are the ATM End System Identifier (ESI) of the LGN, and the final octet is a zero.

In PNNI hierarchy levels 103 (hexadecimal 67) and lower in a multiple-peer group network, the Node ID and ATM address can be the same. In levels 104 (hexadecimal 68) and higher, the Node ID and ATM address *must* be different.

The following is an example of an LGN in a multiple-peer network expressed in hexadecimal notation:



### Logical Group Node (LGN) ID in a Multi-Peer Group Network

#### Note on Default ATM Node Addresses

A default ATM address will be assigned to each PNNI node without user configuration. This default ATM address is equal to

**3903488001bc90000101 xxyyzz 0020da xxyyzz 00**

where **xxyyzz** is the Alcatel-specific OID for the MPM in this chassis.

#### Port ID

Port IDs identify each CSM port in an OmniSwitch, but also identify logical links between nodes. A link between two nodes is identified by the Node IDs of the two nodes as well as the Port IDs of the CSM ports on each end of the link.

PNNI uses Port IDs that are internal to the local switch. ATM switches use different methods for calculating Port IDs, but the values are relevant only to the local switch. In some screen displays, Port IDs may also use the standard OmniSwitch Slot/Port notation to identify ports on each end of a link. For example, the notation **5/2** would indicate port 2 on the CSM module located in slot 5 of the OmniSwitch. The **ppinfo** command displays the local mapping of the slot/port to the advertised PNNI port identification.

#### Peer Group ID

The Peer Group ID identifies a peer group within the PNNI hierarchy. The first octet is the level within the PNNI hierarchy where nodes in this peer group are located. The next 13 octets are the prefix for the ATM End System Address of the node. The prefix

**3903488001bc90000101000000**

is the default peer group ID for each node. As an example, the Peer Group ID for a node in a peer group located at level 80 would be

**50 3903488001bc90000101000000**

where **50** hexadecimal is equal to 80 in decimal notation.

## Summary Addresses

PNNI use a summary address to represent end systems that are attached to a node. Nodes use these summary addresses during exchanges with other nodes in the same peer group. ILMI uses the summary address when satisfying registration requests to attached ATM devices.

The address summary is configurable through the **pncfg** command, which is described in *Configuring Node-Specific Parameters* on page 45-39.

## Single-Peer Networks

In a single-peer group network, the default summary address used by a node is the 13-byte prefix:

**3903488001bc90000101 xxyyzz**

where **xxyyzz** is the 3-byte Alcatel-specific address in the MPM MAC address for that node.

## Multiple-Peer Group Networks

In a multiple-peer group network, there are two summary address: the default summary address and the configured summary address. The configured summary address is taken as the ILMI prefix. The default summary address is derived when the higher node is configured.

# PNNI Packet Types

There are five types of packets used by the PNNI routing protocol to maintain communication between nodes and keep node databases in synchronization. These five packet types are described below.

## Hello packets

Hello packets are exchanged between a node and all of its neighboring nodes. These packets contain Node IDs, Peer Group IDs, Port IDs, and information on Hello protocol timers. The main purpose of Hello packets is to establish the state of the links connecting nodes. In a single peer group model, each link must attain a state of 2-WAY-INSIDE before database synchronization begins. See *Viewing Link Information* on page 45-62 for more information on Hello states.

## Database Summary packets

After two nodes learn about each other through Hello packet exchanges, they inform each other about the entire contents of their topology databases through the exchange of Database Summary packets. Database Summary packets include information on all the PTSEs in a node's topology database.

## PNNI Topology State Element (PTSE) Request packets

When a node receives a Database Summary packet from a neighboring node, it finds all the new PTSEs included in that packet. PNNI nodes need to keep up-to-date topology information, so any new PTSEs from the neighboring node must be obtained. The node sends a request for these new PTSEs in the form of PTSE Request packets. When the neighboring node receives these Request packets it forwards the PTSEs requested by the node.

## PNNI Topology State Packets (PTSPs)

PTSPs bundle topology information from one node for transport to another node. Topology information about a peer group is propagated through a peer group via PTSPs. PNNI Topology State Elements (PTSEs) are individual pieces of information in a node's topology database. A PTSP bundles one or more of these PTSEs and sends them to a node that has requested the particular information included in the PTSE.

## PNNI Topology State Element (PTSE) Acknowledgment packets

When a node receives a particular PTSE from another node it responds with a PTSE Acknowledgment packet. These packets inform the node that sent the PTSEs (bundled inside PTSPs) that the information was transported successfully.

## Metrics and Attributes

PNNI frequently advertises information on the state of links and nodes by means of topology metrics and attributes. A topology *metric* is a parameter that is combined along a path to determine whether this path meets the requirements of a given call request. For example, the administrative weight of a path is the sum of the weights of all links and nodes along the path.

A topology *attribute* is a parameter that is considered individually at each node when determining whether a path meets requirements. For example, if any node along a path violates the Cell Loss Ratio (CLR) for a given call, then that entire path must be rejected.

The following metrics and attributes are used by PNNI for all ATM traffic types. You can configure values for all topology metrics (i.e., Administrative Weight, Cell Transfer Delay, and Cell Delay Variation) on a port-by-port basis via the **ppcfg** command (described in *Configuring Port Parameters* on page 45-46). Topology attributes are discovered dynamically by PNNI from the ATM switch fabric and do not require user configuration.

### Metrics

|                       |  |
|-----------------------|--|
| Administrative Weight | Indicates the preference of a given link relative to other links. Lower administrative weight values have higher priority on the link than higher administrative weight values.          |
| Cell Transfer Delay   | The average time, in microseconds, it takes for cells to transmit from any incoming port to an outgoing port in the switch for a particular Class of Service. The default for CTD is 10. |
| Cell Delay Variation  | Also referred to as “jitter,” this metric is the change that occurs in cell spacing from the time cells leave one node and arrive at another node.                                       |

### Attributes

|                     |  |
|---------------------|--|
| Maximum Cell Rate   | The maximum bandwidth usable by a connection.  |
| Available Cell Rate | The amount of bandwidth available on this link or node. This value is dynamic and changes depending on usage of the link or node.  |
| Cell Loss Ratio     | The ratio of the number of lost cells to the total number of cells transmitted across a link or node. There is a Cell Loss Ratio for CLP=0 traffic ( $CLR_0$ ) and for CLP=0+1 traffic ( $CLR_{0+1}$ ).  |
| Cell Rate Margin    | The difference, in cells per second, between the total bandwidth allocation and the sustainable cell rate allocation. This attribute indicates the “safety margin” of available bandwidth above the amount of bandwidth allocated for the sustainable cell rate. Not supported in the current release. |
| Variance Factor     | A relative measure of the Cell Rate Margin normalized by the variance of the aggregate rate. Not supported in the current release.   |

The table on the next page describes how each ATM Class of Service uses these metrics and attributes. An Administrative Weight is always required to enable a particular Class of Service on a CSM port. However, the use of other metrics and attributes varies by Class of Service.

**Topology Metrics and Attributes and ATM Service Classes**

| Metric/Attribute      | User Configurable? | Service Classes |          |          |          |          |
|-----------------------|--------------------|-----------------|----------|----------|----------|----------|
|                       |                    | CBR             | rt-VBR   | nrt-VBR  | ABR      | UBR      |
| Administrative Weight | Yes                | Required        | Required | Required | Required | Required |
| Cell Transfer Delay   | Yes                | Required        | Required | Required | N/A      | N/A      |
| Cell Delay Variation  | Yes                | Required        | Required | N/A      | N/A      | N/A      |
| Maximum Cell Rate     | No                 | Optional        | Optional | Optional | Required | Required |
| Available Cell Rate   | No                 | Required        | Required | Required | (1)      | N/A      |
| Cell Loss Ratio       | No                 | Required        | Required | Required | N/A      | N/A      |
| Cell Rate Margin      | No                 | N/A             | Optional | Optional | N/A      | N/A      |
| Variance Factor       | No                 | N/A             | Optional | Optional | N/A      | N/A      |

**Table Note:**

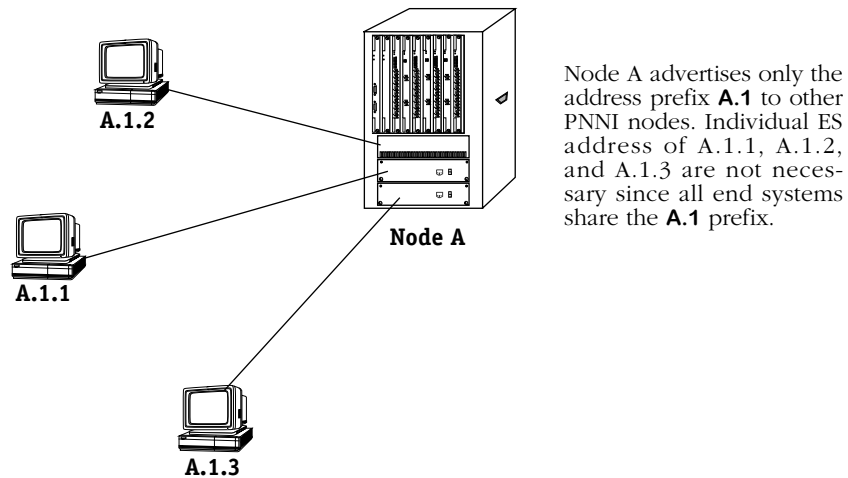
(1) For ABR traffic, the Available Cell Rate is the same as the Minimum Cell Rate.



## Summarization and Reachability

PNNI allows address reachability information to be summarized. PNNI nodes can advertise a single address prefix to represent a group of ATM End System or node addresses that share a common prefix. This summarization reduces the amount of information each node needs to store in its topology database. PNNI matches  $n$  addresses with the summary address. If they match, then the summary address is advertised.

The following diagram illustrates how PNNI summarizes reachability information. In the diagram, Node A summarizes addresses for the three attached End Systems into a single address prefix of **A.1**. This address prefix is the shortest address common to all three systems.



**How PNNI Summarizes Reachable Addresses**


## PNNI Network Initialization

Before OmniSwitch nodes can set up virtual connections between ATM end stations, they need to initialize properly. An OmniSwitch node needs to know its own private ATM address, its own Node ID, the ATM address prefixes of attached ATM End Systems, and information about nodes representing reachable ATM End Systems. Much of this information is self-configured by PNNI, but can also be user-configured through User Interface software or SNMP-based network management software.

Once an OmniSwitch node starts up, it must go through steps to build an accurate topology database. Building this database is essential for the proper functioning of the PNNI routing protocol. OmniSwitches use this topology information to select routes through the network before establishing virtual connections between end systems. In addition, the PNNI topology database is updated periodically to ensure that nodes have an accurate view of the network. The following steps outline what happens at network initialization for a single OmniSwitch where PNNI is enabled.

### Step 1. Discover Attached ATM End Stations

The first step in OmniSwitch PNNI initialization is to discover the ATM end systems that are attached to the node. These end systems include directly attached stations and LAN switches with ATM uplinks to the OmniSwitch. The OmniSwitch can discover ATM end systems through Integrated Local Management Interface (ILMI) address registration. ILMI provides information on the addresses of attached ATM end systems. Address information can also be configured manually for End Systems that do not support ILMI; such configuration takes place at the End System.




**OmniSwitch ATM Switch Discovering Attached ATM Devices**

## Step 2. Discover Neighbor Nodes

In order to set up Routing Control Channel (RCC) virtual connections, each OmniSwitch needs to know about attached PNNI nodes in its peer group. It will be through these neighboring nodes that the OmniSwitch obtains information about the rest of the network. Each OmniSwitch node discovers the identity of its neighboring nodes, or adjacent nodes, and becomes a member of a peer group via the Hello protocol. Nodes in a single peer group exchange Hello messages on each physical link. A Hello message contains the ATM address of the sending node.

The diagram below illustrates the node discovery process. Through Hello message exchanges, Node A.1.1 discovers its neighboring nodes A.1.2, A.1.3, and A.1.4.




**OmniSwitch Discovering Neighboring Nodes in a Single Peer Group**

### Step 3. Send Topology Information for Updating Other Nodes

Once nodes on both sides of a physical link know each other's identity through the exchange of Hello messages, they can begin exchanging topology information. The ATM network will change at times as nodes and devices are added and virtual connections are established and torn down. Because of these changes, the various OmniSwitch nodes need a way to advertise the current information in their topology databases. Entries within a node's topology database are referred to as PNNI Topology State Elements (PTSEs).

Topology information is exchanged via PNNI Topology State Packets (PTSPs). PTSPs contain information on nodes, links, and reachable addresses in the network. Periodically, an OmniSwitch node will send PTSPs that update other nodes of its current state (virtual connections, Quality of Service supported). The topology update messages—PTSPs—are flooded by a reliable hop-by-hop mechanism that ensures that all nodes in the network have updated topology information. In addition, each OmniSwitch can receive PTSPs from other nodes to help build its topology database.

OmniSwitch nodes use the information in their topology database to compute path selection for attached ATM devices that request virtual connections.



OmniSwitches Sending and Receiving PTSPs

### Step 4. Compute the Topology of the Peer Group

Periodically an OmniSwitch will run a Shortest Path First (SPF) algorithm. The interval at which this algorithm is executed is configurable through the **pgcfg** command. If the OmniSwitch has not received any PTSEs from other nodes, then the topology remains intact and the SPF will be deferred to the next period.

If new PTSEs have been received, then the topology database is re-computed. This database is used by call control within the OmniSwitch to forward calls.

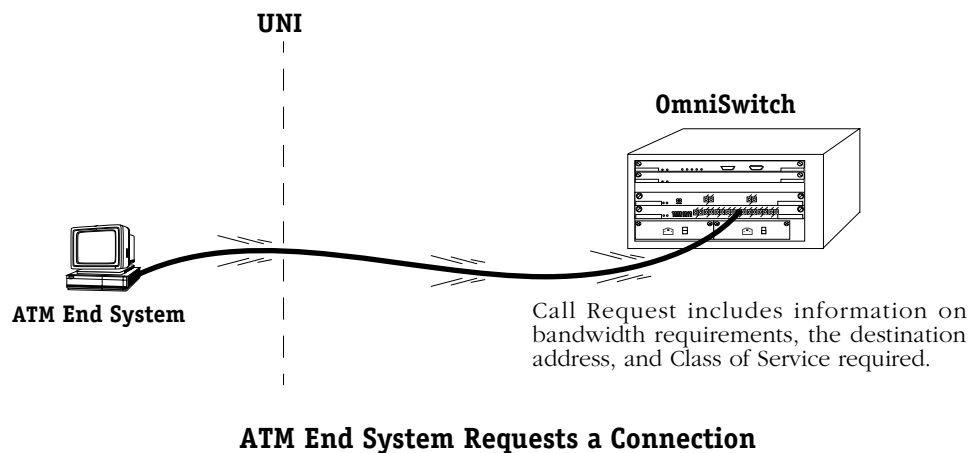
## Establishing a Connection

It is easier to understand the features of PNNI in the OmniSwitch by following a simple example. In this example, the steps required for establishing a point-to-point connection using PNNI are illustrated. This example assumes that the connection is a Switched Virtual Circuit (SVC) and that the OmniSwitch nodes involved are using the PNNI routing protocol. It also assumes that the ATM network has initialized properly as explained in *PNNI Network Initialization* on page 45-18.

### Step 1. Receive a Call Request

The first step is initiated by the ATM End System (ES) that wants to communicate over the ATM network with another ATM ES. An ATM ES could be a workstation, server, LAN switch, or router. The source ATM ES requests a connection with another ATM ES, the destination. This request will be received by the OmniSwitch node to which the source ATM ES is directly connected. The OmniSwitch may receive this request in the form of a signalling protocol message (as in the case of an SVC), or as a request from the Network Management Software on behalf of the station (as in the case of a soft PVC).

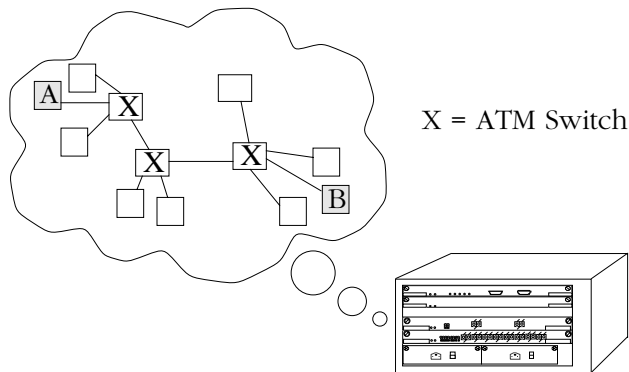
In the case of an SVC, the signalling message will contain information on the requested cell rate, the address of the destination station, a virtual connection identifier (VPI/VCI value) and a requested Class of Service.



### Step 2. Locate Called Parties

In the first step, the OmniSwitch PNNI node learned the address of the destination End System. An OmniSwitch using the PNNI protocol continuously updates its internal map of the ATM network. This map—referred to as a “Topology Database”— includes reachability, link status, and node status information. This map is updated through topology update messages called PNNI Topology State Packets (PTSPs), which are received from other PNNI nodes.

In this second step, the OmniSwitch searches its topology database for the location of the destination device within the ATM network. Within this database, the OmniSwitch finds the location of the end device and the ATM switch to which the destination is attached. This information is also referred to as “reachability” information. By looking through its database, the OmniSwitch decides whether or not this device is reachable. The illustration below shows the OmniSwitch node locating the source (A) and the destination (B) in its topology database. If the end device is not reachable, then connection establishment may be terminated at this point.



**OmniSwitch Locates the Destination in Topology Database**

The topology database contains information on End System and node reachability. This database specifies the nodes where an End System may be reached and the set of paths that lead to that node.

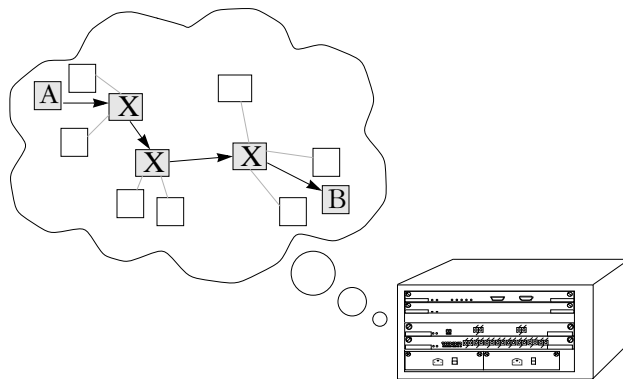
PNNI summarizes reachability information in a node's topology database. Summarization reduces the amount of information needed to be stored in the database. A switch does not require a database of all end stations, just enough information to uniquely identify the device's link to a PNNI node. Therefore, each OmniSwitch node can advertise a single prefix (known as a summary address) instead of the entire ATM address of an End System.

### Step 3. Path Selection

Once the OmniSwitch knows where the destination system is located (i.e., the ATM address of the node that represents the destination), it computes the path through the ATM network that will meet the bandwidth and Class of Service requirements of the original request. The OmniSwitch uses its topology database to determine which paths are possible based on the dynamic state of the network.

The OmniSwitch knows the node state, link state, and reachability information for the ATM network. From this information it can determine which path is best. It may find a switch that is located within the path but does not support enough bandwidth or a particular Class of Service to sustain the connection requirements of the source device. In such a case, that switch will be eliminated from consideration when constructing the path.

The selected path is identified as a sequence of Node IDs and, optionally, Port IDs along a path from End System A (source) to End System B (destination). This information, referred to as a Designated Transit List (DTL), is added to the original Setup message. The illustration below shows an OmniSwitch node selecting a path through the various nodes in the ATM network to the destination (B).

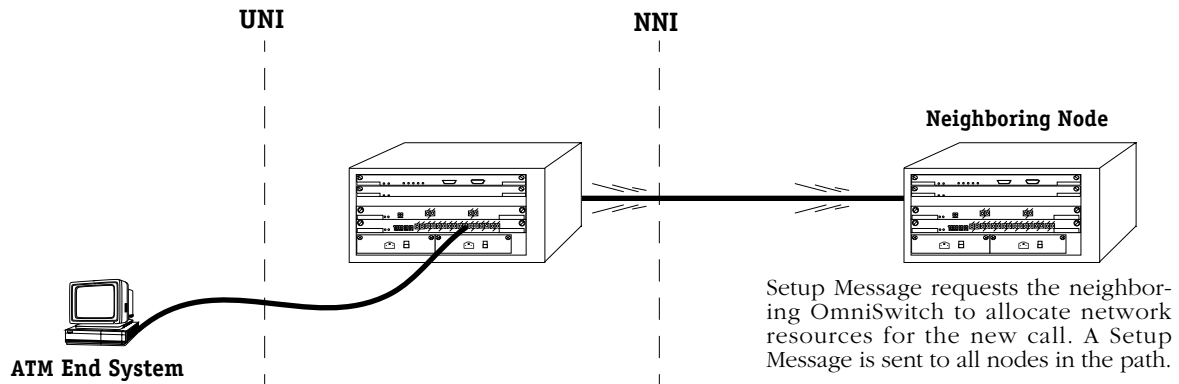


**OmniSwitch Selects a Path to the Destination**

### Step 4. Send Setup Message

The OmniSwitch node now sends a Setup Message along the path selected in the previous step to the first node in the path. The neighboring OmniSwitch (or other ATM switch) will then process the setup request. If this node accepts the request, it forwards the setup message to the next node in the path.

This step will be repeated at each node on the path. However, a node only passes on the request if it accepts the request. The setup message is source routed along the path until all nodes receive the request. Each node in the path is concerned only with the connection to the neighboring node upstream and downstream. The illustration below shows the first OmniSwitch sending the Setup Message to the next node in the path.



**OmniSwitch Sends Setup Message to Neighboring Node**



## Step 5. Process Setup Message

At each hop along the path selected, each ATM switch processes the setup message to determine if it can meet the connection requirements of the request. To process the setup message, a switch uses a specialized part of its functionality referred to as Call Admission Control (CAC). CAC is a set of functionality concerned with determining whether a particular connection request can be met by the available resources in the switch. During this process, CAC looks at its available bandwidth and the Classes of Service it supports. If it determines it can set up the connection, then it assigns a connection identifier (VPI or VPI/VCI) and allocates bandwidth for the connection.


### ◆ Note ◆

The OmniSwitch uses the generic CAC algorithm defined in the ATM Forum PNNI 1.0 specification.

This step is repeated at each switch that receives the setup message. If CAC determines the connection cannot be made, the call will be “cranked back” up to the entry border node in that peer group and *not* to the switch which had originated the setup. It will be cranked back to the switch that originated the setup message only if the call crankback is in the same peer group as that of the originator. An alternate route will be selected for the call.

*Crankback* is the mechanism for releasing an in-progress connection setup due to a failure, such as a link failure or bandwidth allocation failure. *Alternate routing* is a mechanism that allows the call to be re-established on an alternate path when a setup fails.

Even if other switches earlier along the path determined the connection could be set up, the call still cranksback to the originating switch. The Path Selection computations performed in Step 3 are then repeated with knowledge of the node that could not set up the request.




**OmniSwitch Using CAC: Can the Connection Be Accepted?**

### Step 6. Send Call Proceeding Message

After accepting a connection request, each OmniSwitch forwards the setup message to the next hop in the path. This next switch will use its CAC to determine if it can set up the connection as described in Step 5.


After accepting the connection request, a second message is sent backward along the path. This second message, referred to as a Call Proceeding message, goes back to the node that forwarded the call setup message. At each node along the path, a node knows whether the next node was able to set up the connection. These Call Proceeding messages continue along the ATM switches until the final node in the path forwards the final call setup message to the destination ATM end station.



#### Messages Forwarded to Neighboring Nodes

### Step 7. Send Connect Message

After each node in the path accepts the setup message, the destination ES initiates a Connect message back to the source ES. The node to which the destination ES is attached receives this message and forwards it to the next node in the path leading to the source ES. Upon receipt of the Connect Message, a node establishes a leaky bucket algorithm according to the Class of Service and Traffic Descriptors included in the original call setup message. Leaky buckets are described in Chapter 42, "Managing Cell Switching Modules (CSMs)."



#### Connect Message Sent Along Path

This process continues until each node in the path receives the Connect message and sets up a leaky bucket. Finally the initial source ES receives the Connect message and prepares to send data. The connection between the source and the destination is now established end-to-end. All nodes have reserved bandwidth, set up leaky buckets, and assigned virtual circuits to support the connection.

### Step 8. Data Flow

After all the virtual connections have been set up on each node along the path, the source ES begins sending data over the connection.

# The PNNI Menu and Submenus

The PNNI menu contains six submenus that provide command options for configuring PNNI attributes and displaying PNNI configuration parameters and statistics. The main PNNI menu displays as shown below:

| Command   | ATM PNNI Menu                           |
|-----------|---|
| Pconfig   | Enter the PNNI configuration submenu    |
| Proute    | Enter the PNNI route management submenu |
| Pinfo     | Enter the PNNI information submenu      |
| Pstats    | Enter the PNNI statistics submenu       |
| Padmin    | Enter the PNNI administration submenu   |
| Main      | File                                    |
| Interface | Security                                |
|           | Summary                                 |
|           | System                                  |
|           | VLAN                                    |
|           | Services                                |
|           | Networking                              |
|           | Help                                    |

This chapter describes the **Pconfig**, **Pinfo**, **Pstats**, **Ptest**, and **Padmin** submenus. The **Proute** submenu is described in Chapter 46, “Managing IISP and PNNI Routes.” The following sections provide a brief description of each PNNI submenu.

## Pconfig

The **Pconfig** submenu contains options for configuring general PNNI parameters, node-level parameters, and port-level parameters. The submenu displays as shown below.

| Command        | ATM PNNI Configuration Menu                          |
|----------------|--|
| pgcfg          | Configure PNNI general parameters                    |
| pncfg          | Configure PNNI node-specific operation parameters    |
| ppcfg          | Configure PNNI port (interface) operating parameters |
| pndel          | Remove a PNNI logical group node                     |
| Related Menus: |  |
| Pconfig        | Proute   |
|                | Pinfo  |
|                | Pstats   |
|                | Padmin   |

Descriptions of **Pconfig** submenu commands begin on page 45-31.

### Proute

The **Proute** submenu contains options for configuring and monitoring PNNI static routes. The submenu displays as shown below:

| Command        | ATM PNNI Route Management Menu                             |
|----------------|--|
| proutea        | View the Table of routes from nodes to reachable addresses |
| prouten        | View the Table of routes to other nodes                    |
| prpadd         | Add a PNNI static route property (type, metrics, TNS)      |
| prpdcl         | Delete a PNNI static route property                        |
| pradd          | Add PNNI static route address(es) to a route property      |
| prdel          | Delete PNNI static route address(es) from a route property |
| prp            | View PNNI configured route properties                      |
| prrt           | View PNNI configured route prefixes                        |
| Related Menus: |  |
| Pconfig        | Proute    Pinfo    Pstats    Padmin                        |

All **Proute** submenu options are described in Chapter 46, “Managing IISP and PNNI Routes.”

### Pinfo

The **Pinfo** submenu contains options for displaying PNNI general, node, and port attributes. In addition, various connections statistics are available through these commands. The submenu displays as shown below.

| Command        | ATM PNNI Information Menu                                   |
|----------------|---|
| pginfo         | View general PNNI information                               |
| pninfo         | View node-specific PNNI information                         |
| ptinfo         | View PNNI timer information                                 |
| pnbrs          | View PNNI neighbor information                              |
| ppinfo         | View PNNI port information                                  |
| plink          | View PNNI link information                                  |
| pptse          | Dump PTSE database (output may be lengthy)                  |
| pdtl           | View the PNNI DTL Table                                     |
| padj           | View end-point adjacencies reported by call control to PNNI |
| psmap          | View PNNI scope mapping information                         |
| pmap           | View the Pnni Map Table                                     |
| pnmap          | View the Pnni Nodal Map Table                               |
| pcalls         | View the status of ongoing PNNI calls                       |
| Related Menus: |   |
| Pconfig        | Proute    Pinfo    Pstats    Padmin                         |

Descriptions of **Pinfo** submenu commands begin on page 45-49. Many PNNI informational commands provide a “-s” flag option that allows you to summarize displays (see *Summary Form of PNNI Commands* on page 45-30).

◆ **Note** ◆

The **psmap** command will not display unless you are running the multiple-peer group version of PNNI.

## Pstats

The **Pstats** submenu contains options for displaying port, error, and PTSE statistics. The submenu displays as shown below.

| Command        | ATM PNNI Statistics Menu            |
|----------------|-------------------------------------|
| pgstats        | View PNNI port basic statistics     |
| pestats        | View PNNI port error statistics     |
| ppstats        | View PNNI port ptse statistics      |
| Related Menus: |                                     |
| Pconfig        | Proute    Pinfo    Pstats    Padmin |

Descriptions of **Pstats** submenu commands begin on page 45-81.

## Padmin

The **Padmin** submenu contains options for bringing the PNNI protocol entity up and down and for resetting PNNI statistical counters. The submenu displays as shown below.

| Command        | ATM PNNI Administration Menu                        |
|----------------|---|
| phalt          | Halts all PNNI operations - clears all databases    |
| preset         | Reset the PNNI interface counters                   |
| prestart       | Restart the PNNI entity (from a halted state)       |
| pvcfg          | View PNNI information in the configuration file     |
| prmcfg         | Remove PNNI information from the configuration file |
| prtst          | Route test to verify reachability to an ATM address |
| Related Menus: |   |
| Pconfig        | Proute    Pinfo    Pstats    Padmin                 |

Descriptions of **Padmin** submenu commands begin on page 45-86.

### Summary Form of PNNI Commands

You can use the **-s** option on many PNNI display commands to display a summary form instead of a lengthy list of parameters. For example, to display the summary form of the **pnbrs** command, enter

```
pnbrs -s
```

The following PNNI commands support the summary option:

**pnbrs**

**ppinfo**

**plink**

**pptse**

**pmap**

**pnmap**

**pdtl**

**proutea** (described in Chapter 46, “Managing IISP and PNNI Routes”)

**prouten** (described in Chapter 46, “Managing IISP and PNNI Routes”)

### Displaying PNNI Command Help (Multi-Peer Group PNNI Only)

In the multiple-peer group version of PNNI, you can display the syntax for many PNNI commands by entering the command followed by a space, hyphen (-), and the word **help**. For example, to display the syntax for the **pninfo** command, enter

```
pninfo -help
```

The following screen will be displayed.

```
Usage: pninfo [-s | <node level>]
where -s is for summary mode and
<node level> is between 1 and 104
```

#### ◆ Note ◆

You can abbreviate **-help** as **-h**.

Text displayed within square brackets is optional. A vertical line (|) separates mutually exclusive parameters. And a user-supplied value can be indicated “less-than” sign (<) and a “greater-than” sign (>) or by two more words written without spaces (e.g., **portld**).

You can display syntax help on all PNNI commands *except* for the following:

```
pgcfg
ppcfg
prpadd
prdel
prp
prt
padj
phalt
prestart
pvcfg
prmcfg
prtst
```

## Configuring General PNNI Parameters

The **pgcfg** command allows you to configure general PNNI parameters for the local OmniSwitch. It allows you to globally enable PNNI on all CSM ports in the switch, set timers, operational limits, and the topology metrics to use for path selection computations. The values you set here apply to all CSM ports in the switch. Defaults are supplied for all parameters. You can begin configuring these parameters by entering

**pgcfg**

at a system prompt. A screen similar to the following displays.

### Modifying PNNI General Operating Parameters (Will be defaulted to bracketed values if unspecified)

- 1) PNNI enabled on all interfaces configured as  
    type=PNNI in the "map" and cvpt cmds (t,f) [t]:Unspecified
- 2) Shortest path first calc timer (seconds) [ 20]:Unspecified
- 3) Period at which the process to age PTSEs  
    in the database is activated (seconds) [ 10]:Unspecified
- 4) AvCR proportional multiplier (1-99%) [50]:Unspecified
- 5) AvCR minimum threshold (1-99%) [ 3]:Unspecified
- 6) CTD proportional multiplier (1-99%) [50]:Unspecified
- 7) CDV proportional multiplier (1-99%) [25]:Unspecified
- 8) Configure Operation Limits sub-menu (for performance tuning)
- 9) Configure Routing Table Operation sub-menu
- 10) Configure Multi-Peer Group Operation sub-menu

To configure a parameter, type "item = value" (as in 1=t)

To quit out of configuration, type "quit"

To save the configured info, type "save"

-> 1=f

Note that since you've not enabled PNNI on any ports, you must explicitly enable the ports on which PNNI is to operate by using the **ppcfg** command.

-> save

Do you want these parameters to take effect immediately? (y)

To alter a parameter, enter the line number for the parameter, followed by an equal sign (=), and then the new value. For example, to change the **PTSE Aging Timer** (line 3) from 10 seconds to 15 seconds, you would enter:

**3=15**

When you have completed configuring parameters, enter **save**. Your new values will be saved and you will exit this menu. If you want to exit this menu without saving changes, simply enter **quit**.

### ◆ Note ◆

Option 10, **Configure Multi-Peer Group Operation sub-menu**, will not display unless you are running the multiple-peer group version of PNNI.

Options on line numbers 8, 9, and 10 enter submenus with additional configuration parameters. Simply enter the submenu's line number (8, 9, or 10) and press **<Enter>** to go to that submenu. Descriptions for the parameters under option 8 can be found in *Configuring PNNI Operation Limits* on page 45-34. Descriptions for the parameters under option 9 can be found in *Selecting Metrics Used in Path Computations* on page 45-36. Descriptions for the parameters under option 10 can be found in *Configuring Multi-Peer Group Operation* on page 45-37.

### 1) PNNI enabled on all interfaces configured as PNNI

Indicate whether you want to enable the PNNI routing protocol for all CSM ports in this OmniSwitch chassis that have been configured as PNNI type ports through the **map** command (see Chapter 42 for information on the **map** command). You can also enable PNNI on a port-by-port basis through the **ppcfg** command. The setting configured through **ppcfg** takes precedence over the global setting you configure here. For example, if you enable PNNI here (by entering a **t**), but you disable PNNI on a specific port (through **ppcfg**), then PNNI will not run on that port. If you disable PNNI here (by entering an **f**), but enable PNNI on a specific port (through **ppcfg**), then PNNI will run on that port.

### 2) Shortest path first calc timer (seconds)

The time interval, in seconds, between recomputations of routes in the routing tables based on the current contents of the topology database. This timer may range from 1 to 255 seconds. Note that shorter periods increase the sensitivity and reaction time to topology changes, but also consume more switch CPU time in networks where topology changes occur frequently. The default value of 20 is sufficient for most networks of a modest size.

### 3) Period at which the process to age PTSEs is activated (seconds):

The time, in seconds, before the process to age PTSEs starts. When the aging process begins, PTSEs are aged out of the topology database. If a PTSE's lifetime goes to zero (0), then it will be removed. However, a PTSE may be refreshed by the originating node before this entry ages out. This timer may range from 1 to 255 seconds.

### 4) AvCR proportional multiplier (1..99%)

The Available Cell Rate (AVCR) Proportional Multiplier expressed as a percentage. Valid values are integers from 1 to 99. This value is used in determining what defines a significant change in the Available Cell Rate, which is a measure of the bandwidth available for each service class. The switch will measure the AVCR at the current time and at a previous time. The percent indicated here is multiplied by the previous AVCR. If the difference between the current AVCR and the previous AVCR is greater than the product of the Proportional Multiplier (i.e., this value) and the previous AVCR, then the change in AVCR will be considered significant by PNNI.

### 5) AvCR minimum threshold (1..99%)

The Available Cell Rate (AVCR) Minimum Threshold expressed as a percentage. Valid values are integers from 1 to 99. This value is used in computing the lowest level of significant change in the Available Cell Rate, which is the bandwidth available for each service type. The value you indicate here is multiplied by the Maximum Cell Rate to yield the minimum difference (allowed in computations) between the current AVCR and a previously measured AVCR. If the previously measured AVCR multiplied by the AVCR Proportional Multiplier (indicated on line 4 in this menu) is less than the Maximum AVCR multiplied by the value indicated here, then this value will be used in computations to determine upper and lower limits of significance in AVCR change.



**6) CTD proportional multiplier (1..99%)**

The Cell Transfer Delay (CTD) Proportional Multiplier expressed as a percentage. Valid values are integers from 1 to 99. This value is used in determining what defines a significant change in the Cell Transfer Delay, which is the time it takes cells to transmit across a link within a single peer group. The switch will measure the CTD at the current time and at a previous time. The percent indicated here is multiplied by the previous CTD. If the difference between the current CTD and the previous CTD is greater than the product of the Proportional Multiplier (i.e., this value) and the previous CTD, then the change in CTD will be considered significant by PNNI.

**7) CDV proportional multiplier (1..99%)**

The Cell Delay Variation (CDV) Proportional Multiplier expressed as a percentage. Valid values are integers from 1 to 99. This value is used in determining what defines a significant change in the Cell Delay Variation, which is a measure of “jitter” or the change in cell spacing over a given link. The switch will measure the CDV at the current time and at a previous time. The percent indicated here is multiplied by the previous CDV. If the difference between the current CDV and the previous CDV is greater than the product of the Proportional Multiplier (i.e., this value) and the previous CDV, then the change in CDV will be considered significant by PNNI.

### Configuring PNNI Operation Limits

Option 8 on the main **pgcfg** configuration menu opens a submenu for fine-tuning performance and resource utilization in the local switch. The options on this submenu set limits for various PNNI attributes, such as nodes in the network, Designated Transit Lists (DTLs), and static routes. PNNI uses these values to pre-allocate storage for the related elements. When limits are set too high, memory can be wasted. If limits are set too low, then the database can become overloaded.

#### ◆ Note ◆

If you are unsure of a limit, use the default supplied. If the default clearly does not fit your network configuration, then configure it appropriately.

The submenu displays as follows:

#### Modifying PNNI General Operating Limit Parameters (Will be defaulted to bracketed values if unspecified)

|   |                    |
|---|--------------------|
| 10) Approximate number of nodes in network      | [100]:Unspecified  |
| 11) Approximate Max neighbors for this node     | [30]:Unspecified   |
| 12) Approximate Max PTSEs in net                | [1000]:Unspecified |
| 13) Approximate Max Information Groups in PTSEs | [1500]:Unspecified |
| 14) Max reachable addresses in this network     | [400]:Unspecified  |
| 15) Max transit networks in this network        | [50]:Unspecified   |
| 16) Max Designated Transit Lists (DTLs)         | [200]:Unspecified  |
| 17) Max DTL entries                             | [200]:Unspecified  |
| 18) Max outstanding route requests              | [500]:Unspecified  |
| 19) Max configured static routes                | [300]:Unspecified  |
| 20) Max configured static route groups          | [200]:Unspecified  |
| 21) Max point to multipoint endpoints           | [500]:Unspecified  |
| 22) Max number of Paths                         | [1000]:Unspecified |
| 23) Max number of retransmissions               | [250]:Unspecified  |
| 24) Max number of entries in any routing tbl    | [2000]:Unspecified |

To configure a parameter, type "item = value" (as in 11=100)

To quit out of configuration, type "quit"

To save the configured info, type "save"

To return to General Operating Parameters, type "return"

Operational limit parameters are described below. Unless otherwise specified, the values for these limits may range from 1 to 1,000.

#### 10) Max Nodes in Network

The maximum number of nodes that are allowed in this PNNI network.

#### 11) Max neighbors

The maximum number of neighboring node connections for each switch. A good rule of thumb is to never attach more than 30 neighbors without making appropriate adjustments to the hello timers.

#### 12) Max PTSEs

The maximum number of PTSEs that can be held in the topology database of a switch. Up to 65,535 PTSEs may exist in a switch.

**13) Max Information Groups in PTSEs**

The maximum number of Information Groups this switch will store in its topology state database. The switch will store up to 65,535 Information Groups.

**14) Max reachable addresses in this network**

The maximum number of ATM End System addresses that will be available to nodes in this PNNI network. This value may range from 1 to 65,535.

**15) Max transit networks in this network**

A transit network is a route used to tunnel call requests from an ATM End System in one peer group to an ATM End System in another peer group. A transit network differs from a Designated Transit List (DTL) in that it provides a route to links outside the peer group while a DTL provides a route to nodes within the same peer group. This value may range from 1 to 255.

**16) Max Designated Transit Lists (DTLs)**

The maximum number of Designated Transit Lists (DTLs) that may be set up through the PNNI network. A DTL is a complete source route through a peer group. This value may range from 1 to 65,535.

**17) Max DTL entries**

The maximum number of entries (hops) that can be used to make up a single DTL. An entire DTL includes all hops that comprise a source route through a peer group. In a single peer group, this value is the maximum path length plus 1. This value may range from 1 to 500.

**18) Max outstanding route requests**

The maximum number of call setup messages that can be outstanding at one time. Call setup messages are initiated by ATM End Systems. When a call setup is in progress, a call descriptor is held by PNNI until the connection is set up or the setup cranksback. This value may range from 1 to 65,535.

**19) Max configured static routes**

The maximum number of static routes that can be configured in the PNNI network. A static route is a vector through the network that routes to a particular class of NSAP addresses. This value may range from 1 to 10,200.

**20) Max configured static route groups**

The maximum number of static route groups allowed. A group of static routes is ordered by route properties. This value limits the amount of space reserved for route properties. This value may range from 1 to 255.

**21) Max point to multipoint endpoints**

The maximum number of multicast virtual circuits that may be configured within the PNNI network. This number is limited by the type and number of CSM modules installed. CSM-155 modules each support 8,000 multicasts, and CSM-622 modules support 16,000 multicasts.

**22) Max number of Paths**

The maximum number of paths that may be stored in this switch. A path is a physical connection between two switches.

### 23) Max number of retransmissions

PTSE requests and PTSEs will be retransmitted this number of times before PNNI declares that the neighboring node is down. This value may range from 1 to 255.

### 24) Max number of entries in any routing tbl

The maximum number of routing table entries. This value may range from 1 to 65,535.

## Selecting Metrics Used in Path Computations

Option 9 on the main **pgcfg** configuration menu opens a submenu for configuring the Routing Table. These submenu options allow you to choose the topology metric on which each Class of Service will be sorted. Each Class of Service can be sorted on Administrative Weight, Cell Transfer Delay (CTD), or Cell Delay Variation (CDV). The submenu displays as follows:

### Modifying PNNI Routing Table Configuration

This sub-menu allows a user to configure a PNNI routing table. When the next route calculation executes, a routing table will be generated for the configured quality of service on the metric specified. Note that there may be only one metric per QOS specified. (The default metric in brackets will be used if the value is unspecified.)

Valid metric values are:

Admin Weight:     **AW**  
Cell Transit Delay:   **CTD**  
Cell Delay Variation: **CDV**

|  |                  |
|--|------------------|
| 25) QOS Class CBR sorted on Metric     | [AW]:Unspecified |
| 26) QOS Class rt-VBR sorted on Metric  | [AW]:Unspecified |
| 27) QOS Class nrt-VBR sorted on Metric | [AW]:Unspecified |
| 28) QOS Class ABR sorted on Metric     | [AW]:Unspecified |
| 29) QOS Class UBR sorted on Metric     | [AW]:Unspecified |

To configure a parameter, type "item = value" (as in 26=AW)

To quit out of configuration, type "quit"

To save the configured info, type "save"

To return to General Operating Parameters, type "return"

### 25) QOS Class CBR sorted on Metric

The metric on which Constant Bit Rate (CBR) traffic will be sorted during path selection computations. CBR traffic can be sorted by Administrative Weight (AW), Cell Transfer Delay (CTD), or Cell Delay Variation (CDV). The default metric is Administrative Weight.

### 26) QOS Class VBR-RT sorted on Metric

The metric on which real-time Variable Bit Rate (VBR-RT) traffic will be sorted during path selection computations. VBR-RT traffic can be sorted by Administrative Weight (AW), Cell Transfer Delay (CTD), or Cell Delay Variation (CDV). The default is Administrative Weight.

### 27) QOS Class VBR-NRT sorted on Metric

The metric on which non-real-time Variable Bit Rate (VBR-NRT) traffic will be sorted during path selection computations. VBR-NRT traffic can be sorted by Administrative Weight (AW), Cell Transfer Delay (CTD), or Cell Delay Variation (CDV). The default is Admin Weight.

**28) QOS Class ABR sorted on Metric**

The metric on which Available Bit Rate (ABR) traffic will be sorted during path selection computations. ABR traffic can be sorted by Administrative Weight (AW), Cell Transfer Delay (CTD), or Cell Delay Variation (CDV). The default metric is Administrative Weight.

**29) QOS Class UBR sorted on Metric**

The metric on which Unspecified Bit Rate (UBR) traffic will be sorted during path selection computations. UBR traffic can be sorted by Administrative Weight (AW), Cell Transfer Delay (CTD), or Cell Delay Variation (CDV). The default metric is Administrative Weight.

**Configuring Multi-Peer Group Operation**

Option 10 on the main **pgcfg** configuration menu opens a submenu for configuring Multi-Peer Group PNNI operation. These submenu options allow you configure a node for operation as a Logical Group Node (LGN). The submenu displays as follows:

**Modifying PNNI Multi-Peer-Group Operating Parameters**  
(Will be defaulted to bracketed values if unspecified)

31) Interval for initiating SVC-RCC (secs) [ 1]:Unspecified

The following are for all LGN levels operating within this switch:

32) Max RCC expected to be established [200]:Unspecified

33) Max LGN Hor Links expected to be advertised [100]:Unspecified

34) Max DTL Hops [ 50]:Unspecified

To configure a parameter, type "item = value" (as in 31=10)

To quit out of configuration, type "quit"

To save the configured info, type "save"

To return to General Operating Parameters, type "return"

:

Multi-peer group (MPG) operation parameters are described below. Options 32, 33, and 34 are used to set Logical Group Node (LGN) parameters for the local OmniSwitch only.

**31) Interval for initiating SVC-RCC (secs)**

Enter the maximum amount of time (in seconds) to establish a Switched Virtual Circuit (SVC) Routing Channel Connection (RCC). This value can be from 1 to 1800 seconds (the default is 1 second). SVC-RCCs are used to route PNNI management information, such as PNNI Topology State Packets (PTSPs) and Database Summary packets.

**32) Max RCC expected to be established**

Enter the maximum amount of Routing Channel Connections (RCCs) for this OmniSwitch. This value can be from 1 to 1000 (the default is 200). Switched Virtual Circuit (SVC)-RCCs are used to route PNNI management information, such as PNNI Topology State Packets (PTSPs) and Database Summary packets.

**33) Max LGN Hor Links expected to be advertised**

Enter the maximum number of horizontal links to Logical Group Nodes (LGNs) to be advertised by this node. This value can be from 1 to 1000 (the default is 100).

### **34) Max DTL Hops**

Enter the maximum number of Designated Transit List (DTL) hops in this node. This value can be from 1 to 1000 (the default is 50).

## Configuring Node-Specific Parameters

The **pncfg** command allows you to configure node-level PNNI parameters. It allows you to set the ATM address, PNNI node level, and administrative status of this node. In addition, you can configure several refresh timers for PTSE exchanges and the Hello protocol. The values you set here apply to all CSM ports in the OmniSwitch. Defaults are supplied for all parameters.

The **pncfg** command operates somewhat differently between the single peer group version of the software and the multiple-peer group version of the software. For descriptions of the multiple-peer group version of the **pncfg** command, see *Configuring Multiple-Peer Group Nodes* on page 45-43. For the single-peer group version of the **pncfg** command, see the section below.

### Configuring Single-Peer Group Nodes

You can begin configuring node-specific parameters in the single-peer group version of the software by entering

```
pncfg
```

at a system prompt. A screen similar to the following displays.

```

      ATM PNNI Node-specific Configuration
      (Will be defaulted to bracketed values if unspecified)

1) ATM address of this node(hex)           :Unspecified
   [3903488001bc9000010178aee00020da78aee000]
2) Node ID (without level and rsvd)        :Unspecified
   [3903488001bc9000010178aee00020da78aee000]
3) Local node level                        [80]:Unspecified
4) Admin status (up,down)                  [up]:Unspecified

5) Advertise Addr. summary(t,f)           [t]: Unspecified
6) Summary address for ILMI clients        : Unspecified
   [3903488001bc90000178aee0]

7) Timer configuration sub-menu

      To configure a parameter, type "item = value" (as in 3=100)
      To quit out of configuration, type "quit"
      To save the configured info, type "save"
-> 2=3903488001bc9000010178aeb10020da78aeb100
-> ?

```

To alter a parameter, enter the line number for the parameter, followed by an equal sign (=), and then the value for the parameter. For example, to enter an ATM address for the node, you might enter:

```
1=41000700040011223344556677080a1100000100
```

When you have completed configuring parameters, enter **save**. Your new values will be saved and you will exit this menu. If you want to exit this menu without saving changes, simply enter **quit**.

Line number 7 enters a submenu for configuring PTSE and Hello protocol timers. Simply enter **7** and press **<Enter>** to go to this submenu. Parameters for option 7 can be found in *Configuring PTSE and Hello Timers* on page 45-41.

### 1) ATM address of this node (hex)

The 20-byte ATM address for this node. Other nodes in the network that need to exchange PNNI protocol packets with this node will direct those packets to this address. The default ATM node address is

**3903488001bc90000101xxyzz0020daxxyzz00**

where **xxyzz** is the Alcatel-specific OID for the MPM in this chassis. This parameter can only be configured on a lowest-level node.

### 2) Node ID (without level and rsvd)

The Node ID for this node. The Node ID is a 22-octet identifier for a node within the PNNI network. Do not include the node level or other reserved characters in this specification. Include only the last 20-bytes of the Node ID. This parameter can only be configured on a lowest-level node.

### 3) Local Node Level

The level within the PNNI hierarchy where this node exists. This attribute is used to determine the default Node ID and the default Peer Group ID for this node. This value may only be set when the node's Admin Status is down. PNNI levels range from 0 to 104. The default node level is 80 in decimal notation. Lower values are higher in the hierarchy than higher values.

In a single peer group network, all nodes in the PNNI network will have the same node level.

### 4) Admin Status

Indicates the Administrative Status of this node. The default is **Up**, which means that the node can become operationally active and participate in PNNI protocol exchanges. If set to **Down**, then the node will be inactive and not participate in PNNI protocol exchanges.

### 5) Advertise Addr. summary

Indicate whether you want this node to use summarization when advertising the addresses of attached devices to other PNNI nodes. Using address summarization to advertise internal reachability speeds PNNI database searches. If a node does not support address summarization, then it will advertise the entire local address of its attached devices during PNNI exchanges. This parameter can only be configured on a lowest-level node.

### 6) Summary address for ILMI clients

Required only if you turn on address summarization in line 5. Enter the summary address that will be used to advertise all devices attached to this node. This parameter can only be configured on a lowest-level node. See *Summary Addresses* on page 45-13 for more information on summary addresses.



## Configuring PTSE and Hello Timers

Option 7 on the main **pncfg** configuration menu opens a submenu for configuring PTSE and Hello protocol timers. These submenu options set parameters for how PTSE, Database Summary, and Hello packets are transmitted, re-transmitted, and acknowledged. The submenu displays as shown below:

```

      ATM PNNI Node Timer Configuration
      (Unspecified attributes assume the value in [] during operation)

8) PTSE timers:
   a) Refresh interval (self-orig ptse)    [1800]:Unspecified
   b) Lifetime factor (multiples of a)     [2]:Unspecified
   c) Hold down (in seconds)               [1]:Unspecified
   d) Delayed ack timer                   [1]:Unspecified
9) PTSP transmit timer                    [10]:Unspecified
10) Timer to Xmit PTSPs in resp to incoming PTSE requests [1]:Unspecified
11) Db summary re-transmit time (if unack'd) [3]:Unspecified
12) Hello timers:
   a) Hello interval (seconds: 1-255)     [15]:Unspecified
   b) Hold down interval (in seconds)     [3]:Unspecified
   c) Inactivity factor (multiples of a)  [5]:Unspecified

      To configure a parameter, type "item = value" (as in 9a=2000)
      To quit out of configuration, type "quit"
      To save the configured info, type "save"
      To return to Node Config Menu, type "return"

```

### 8) PTSE Timers

#### a) Refresh interval

The time, in seconds, before a self-originated PTSE is updated. PTSEs are aged out of the database unless refreshed by the originating node. The lifetime of a PTSE is determined by multiplying the Refresh Interval by the Lifetime Factor (specified in line 9b). The range for this value is 1 to 32,767 seconds. The default value is 1800 seconds.

#### b) Lifetime factor (multiples of a)

The value for the PTSE lifetime multiplier expressed as a percentage. Valid values are integers ranging from 1 to 255. This value helps determine the initial lifetime of a PTSE in the topology database. The Lifetime Factor multiplied by the Refresh Interval (specified on line 9a) is the initial lifetime of a PTSE. The default value is 2.

#### c) Holddown (in seconds)

The minimum time, in seconds, before which this node can refresh PTSEs. A node can prevent a PTSE from aging out of the topology database by refreshing it. This Holddown value limits the node from refreshing PTSEs too often and exhausting database space too quickly. The range for this value is from 1 to 255 seconds. The default value is 1 second.

#### d) Delayed ack timer

When a node receives a PTSE from another node it sends back an Acknowledgment Packet. However, the acknowledgment is not immediate. The amount of time between the receipt of a PTSE and its acknowledgment is the value you enter here. The range is from 1 to 255 seconds.

### 9) PTSP transmit timer

PTSPs are sent until they are acknowledged by neighboring nodes. This variable is the amount of time, in seconds, between successive transmissions of PTSPs. If a PTSP is acknowledged before the time interval specified here, then a retransmission will not be sent. This value may range from 1 to 255 seconds.

### 10) Timer to Xmit PTSPs in resp to incoming ptse request

The time interval between the receipt of Database Summary packets and the sending of PTSE Request packets. Database Summary packets contain an index of the PTSEs in a node's topology database. Other nodes use Database Summary packets to request PTSEs. If a node requires a PTSE listed in a Database Summary packet, then it requests that PTSE in the form of a PTSE Request packet. This value may range from 1 to 255 seconds.

### 11) Db summary re-transmit time

The time, in seconds, before this node will re-transmit a Database Summary packet that has gone unacknowledged by another node. This value may range from 1 to 255 seconds.

### 12) Hello timers

#### a) Hello interval

The initial value for the Hello timer in seconds. In the absence of triggered Hellos, this node will send one Hello packet on each of its ports at the interval specified here. The default Hello Interval is 15 seconds. Values can range from 1 to 255 seconds.

#### b) Hold down interval

The initial value for the Hello hold down timer. This node will use this value to limit the rate at which it sends Hello messages. The default value is 3 seconds. Valid values range from 1 to 255 seconds.

#### c) Inactivity factor

The number of Hello intervals that may pass without receiving a Hello before the neighboring node is determined to have gone down. The default is 5 seconds. Valid values range from 1 to 255.

## Configuring Multiple-Peer Group Nodes

You can begin configuring node-specific parameters in the multiple-peer group version of the software by entering

**pncfg**

at a system prompt. A screen similar to the following displays.

```

      ATM PNNI Node-specific Configuration
      (Will be defaulted to bracketed values if unspecified)

1) ATM address of this node(hex)           :Unspecified
   [3903488001bc9000010178aee00020da78aee000]
2) Node ID (without level and rsvd)        :Unspecified
   [3903488001bc9000010178aee00020da78aee000]
3) Local node level                        [80]:Unspecified
4) Admin status (up,down)                  [up]:Unspecified

5) Advertise Addr. summary(t,f)           [t]: Unspecified
6) Summary address for ILMI clients        : Unspecified
   [3903488001bc90000178aee0]

7) PGL Priority                            [ 50]: Unspecified
8) Timer configuration sub-menu

      To configure a parameter, type "item = value" (as in 3=100)
      To quit out of configuration, type "quit"
      To save the configured info, type "save"
-> 2=3903488001bc9000010178aeb10020da78aeb100
-> ?

```

To alter a parameter, enter the line number for the parameter, followed by an equal sign (=), and then the value for the parameter. For example, to enter an ATM address for the node, you might enter:

```
1=41000700040011223344556677080a1100000100
```

When you have completed configuring parameters, enter **save**. Your new values will be saved and you will exit this menu. If you want to exit this menu without saving changes, simply enter **quit**.

See *Configuring Single-Peer Group Nodes* on page 45-39 for descriptions of options 1 (**ATM address of this node**) through options 6 (**Summary address for ILMI clients**). Line number 7 (**PGL Priority**) is described in *7) PGL Priority* on page 45-44. Line number 8 enters a submenu for configuring PTSE and Hello protocol timers. Simply enter **8** and press **<Enter>** to go to this submenu. Parameters for option 8 can be found in *Configuring PTSE and Hello Timers* on page 45-41.

### ◆ Note ◆

The option number for the timer configuration submenu and all of its suboptions are one (1) number higher in the multiple-peer group version of PNNI due to the addition of option 7 (**PGL Priority**), which is described on the following page.

In addition, Peer Group Leader (PGL) nodes have some additional parameters, which are described in *Configuring Peer Group Leader Nodes* on page 45-44.

### 7) PGL Priority

The Peer Group Leader (PGL) priority for this node. You can use this parameter to configure which node will become the primary leader, secondary backup leader, etc. in a peer group. Lower values have higher priority than higher values. The PGL node will commence Logical Group Node (LGN) operation. You can enter a value from 0 to 205.

Any node within a peer group can be the PGL. A border node may or may not be the PGL.

#### ◆ Note ◆

The PGL of any level must be the PGL of all descendant (child) levels.

### Configuring Peer Group Leader Nodes

There are additional parameters in the **pncfg** command when you configure a Peer Group Leader of a Logical Group Node. On these nodes, a screen similar to the following will be displayed when you execute the **pncfg** command.

#### ATM PNNI Lowest Level Node-specific Configuration (Will be defaulted to bracketed values if unspecified)

- 1) ATM address of this node(hex) :Unspecified  
[3903488001bc9000010178aee00020da78aee000]
- 2) Node ID (without level and rsvd) :Unspecified  
[3903488001bc9000010178aee00020da78aee000]
- 3) Local node level [80]:Unspecified
- 4) Admin status (up,down) [up]:Unspecified
- 5) Advertise Addr. summary(t,f) [t]: Unspecified
- 6) Summary address for ILMI clients : Unspecified  
[3903488001bc90000178aee0]
- 7) PGL Priority [ 50]: Unspecified
- 8) Complex Representation(t,f) [f]: f
- 9) Timer configuration sub-menu

To configure a parameter, type "item = value" (as in 3=100)

To quit out of configuration, type "quit"

To save the configured info, type "save"

-> 2=3903488001bc9000010178aeb10020da78aeb100

-> ?

These additional parameters are described below.

#### ◆ Note ◆

The option numbers in the timer configuration submenu and all of its suboptions are one (1) number higher due to the addition of option 8 (**Complex Representation**), which is described below.

### 8) Complex Representation(t,f)

The complex node representation is the process of representing a child peer group by a logical node in its parent peer group. Set this parameter to **t** (true) to activate complex node representation. If you set this parameter to **f** (the default), then simple node representation will be implemented where the nodal state parameter PTSEs from this node will not be used in route computations. See *Complex Representation* on page 45-9 for more information.

## 9) Timer configuration sub-menu

When you enter the timer configuration (now option 9) on a Peer Group Leader of a Logical Group Node, a screen similar to the following will be displayed:

### ATM PNNI Lowest Level Node Timer Configuration (Unspecified attributes assume the value in [] during operation)

- 10) PTSE timers:
  - a) Refresh interval (self-orig ptseqs) [1800]:Unspecified
  - b) Lifetime factor (multiples of a) [2]:Unspecified
  - c) Hold down (in seconds) [1]:Unspecified
  - d) Delayed ack timer [1]:Unspecified
- 11) PTSP transmit timer [10]:Unspecified
- 12) Timer to Xmit PTSPs in resp to incoming PTSE requests [1]:Unspecified
- 13) Db summary re-transmit time (if unack'd) [3]:Unspecified
- 14) Hello timers:
  - a) Hello interval (seconds: 1-255) [15]:Unspecified
  - b) Hold down interval (in seconds) [3]:Unspecified
  - c) Inactivity factor (multiples of a) [5]:Unspecified
- 15) LGN Specific Timers:
  - a) SVCC calling interval timer [35]:35
  - b) SVCC called interval timer [50]:50
  - c) Horizontal link inactive timer [75]:75

To configure a parameter, type "item = value" (as in 9a=2000)

To quit out of configuration, type "quit"

To save the configured info, type "save"

To return to Node Config Menu, type "return"

Options 10 (PTSE timers) through option 14 (Hello timers) are described in *Configuring PTSE and Hello Timers* on page 45-41. Option 15 (LGN Specific Timers) is described below.

## 15) LGN Specific Timers

### a) SVCC calling interval timer

The Switched Virtual Circuit Channel Connection (SVCC) interval timer for this node in seconds. An SVCC is a routing control channel between logical group nodes. This parameter sets the maximum amount time that a calling LGN should wait before restarting the process of establishing an SVCC-based Routing Control Channel (RCC). Values can range from 1 to 65535 seconds. (The default is 35 seconds.)

### b) SVCC called interval timer

Enter the Switched Virtual Circuit Channel Connection (SVCC) interval timer for this node. An SVCC is a routing control channel between logical group nodes. This parameter sets the maximum amount time that a called LGN should wait for an SVCC-based Routing Control Channel (RCC) to be established. Values range from 1 to 65535 seconds. (The default is 50 seconds.)

### c) Horizontal link inactive timer

The horizontal link inactive timer for this node in seconds. This parameter determines the interval a node should wait to re-establish a horizontal link with a neighboring peer group after no Hello packets have been received from that neighboring node. data between neighboring peer groups. Valid values range from 1 to 65535 seconds. (The default is 75 seconds.)

# Configuring Port Parameters

The **ppcfg** command allows you to configure port-level PNNI parameters. It allows you to enable PNNI on a port and configure metrics for each ATM traffic type supported on the port. (Before PNNI can be enabled on a port, the port must be configured as a PNNI type port through the **map** command, which is described in Chapter 42.) You can begin configuring PNNI port parameters by entering

```
ppcfg <slot>/<port>
```

at a system prompt. For example, if you wanted to configure PNNI for port 1 on the CSM module in slot 5, then you would enter:

```
ppcfg 5/1
```

You could also configure multiple ports at one time by specifying a port list after the slot number. For example, you could configure ports 1 and 2 on the CSM module in slot 5 by specifying:

```
ppcfg 5/1-2
```

**Virtual Path Tunnels.** To configure PNNI parameters for a specific virtual path tunnel, you need to include the instance number of the virtual tunnel in the **ppcfg** command. (Physical level parameters for virtual tunnels are configured through the **cvpt** command, which is described in Chapter 43.) The **ppcfg** format for virtual path tunnels is as follows:

```
ppcfg <slot>/<port>/<virtual tunnel instance>
```

where **<virtual tunnel instance>** is a unique value assigned to each virtual tunnel on a CSM module port. You can find a specific virtual tunnel instance through the **lvpt** command. If you wanted to configure PNNI parameters for the second virtual tunnel instance on third port on the CSM module in slot 4, you would specify:

```
ppcfg 4/3/2
```

The following is a sample display for the **ppcfg** command.

Modifying PNNI/AAL Operating/Advertisement Parameters for slot 5 port 1:  
1) PNNI enabled on this port (t,f) [t]:Unspecified  
2) VPI [0]:Unspecified  
3) VCI [18]:Unspecified

|            | Admin Weight  | Cell Transfer Delay | Cell Delay Variance |
|------------|---|---------------------|---------------------|
|            | (a)[5040]   | (b)                 | (c)                 |
|            | (Disable a class by assigning an admin weight to 0) |                     |                     |
| 4) CBR     | Disabled  |                     |                     |
| 5) rt-VBR  | Disabled  |                     |                     |
| 6) nrt-VBR | Disabled  |                     |                     |
| 7) ABR     | Disabled  |                     |                     |
| 8) UBR     | Disabled  |                     |                     |

To set a value, type "item=value" (as in 1=y);  
To cancel this and move on to the next port, type "next";  
To quit out of configuration, type "quit";  
To save this port information and move on to the next port, type "save".  
-> 1=t  
-> save  
Do you want these parameters to take effect immediately? (y)  
Done.

### 1) PNNI enabled on this port

Indicates whether PNNI is enabled on this port. If PNNI is enabled on all ports through the **pgcfg** command, then this value will already be set to **Yes** (enabled). This value overrides the value set through **pgcfg**. If PNNI is enabled on all ports, you can disable it on this port by setting this value to **No**. If this field is set to **No**, then all other parameters on this screen will not display.

In release 4.1 and later, you should not disable ILMI since ILMI will be running across all ATM links unless it has been explicitly disabled. In releases prior to 4.1, you should disable ILMI—through the **map** command—on a PNNI port.

### 2) VPI

The default Virtual Path Identifier (VPI) that will be used to transmit PNNI routing messages. This VPI will not be available for data connections. This Virtual Path is also referred to as the Routing Control Channel (RCC). If you are configuring a virtual tunnel, then you will need to specify this VPI value through the **cvpt** command, not in this field. Values may range from 0 to 255. The default is 0.

### 3) VCI

The default Virtual Channel Identifier (VCI) that will be used to transmit PNNI routing messages within each Virtual Path. This VCI will not be available for data connections. This Virtual Channel is also referred to as the Routing Control Channel (RCC). Values may range from 1 to 65,536. The default is 18.

## 4-8) Metrics

You can configure three metrics for each class of ATM traffic. These configurable metrics are Administrative Weight, Cell Transfer Delay (CTD), and Cell Delay Variation (CDV).

To configure these metrics, first enter the number for the Class of Service, then the letter corresponding to the metric you want to configure, an equal sign (=), and finally the value for the metric. For example, if you wanted to assign an Administrative Weight of **3200** to the CBR traffic class, then you would enter:

**4a=3200**

The three configurable metrics are described below. PNNI also uses a number of topology attributes for connections, but these attributes are not user-configurable; PNNI reads topology attributes directly from the switch fabric. See *Metrics and Attributes* on page 45-15 for further information.

### a) Admin Wt

You enable a class of traffic by assigning an Administrative Weight to it. The Administrative Weight indicates the preference of a given link relative to other links. Lower values have a higher priority than higher values.

### **b) Cell Transfer Delay**

The time it takes for cells to transmit across a link within a single peer group.

### **c) Cell Delay Var**

Also referred to as “jitter,” this metric is the change that occurs in cell spacing from the time cells leave one node and arrive at another node.

#### **◆ Note ◆**

The Maximum Transmission Unit (MTU) is set to a constant 8192 bytes and cannot be altered.



## Viewing General PNNI Information

The **pginfo** command displays several current configuration values as well as statistics on connections. When you enter

**pginfo**

a screen similar to the following displays:

### ATM PNNI General Information

|  |            |                                      |            |
|--|------------|--------------------------------------|------------|
| <b>Nodes in this switch:</b>           | <b>1</b>   | <b>Neighbors detected:</b>           | <b>5</b>   |
| <b>PNNI highest version supported:</b> | <b>1</b>   | <b>Lowest version:</b>               | <b>1</b>   |
| <b>Node routing database size:</b>     | <b>7</b>   | <b>Address database size:</b>        | <b>10</b>  |
| <b>PTSEs in database:</b>              | <b>24</b>  | <b>RCCS in database:</b>             | <b>5</b>   |
| <b>Times SPF executed:</b>             | <b>5</b>   |                                      |            |
| <b>Pt to Pt calls in progress:</b>     | <b>13</b>  | <b>Pt to MultiPt calls in prog:</b>  | <b>11</b>  |
| <b>Conns cranked back to this sys:</b> | <b>0</b>   | <b>Conns cranked from border:</b>    | <b>0</b>   |
| <b>DTL stacks in use:</b>              | <b>24</b>  | <b>DTL Stacks Free:</b>              | <b>176</b> |
| <b>Total DTL stacks originated:</b>    | <b>453</b> | <b>Total DTL borders originated:</b> | <b>0</b>   |
| <b>Alt DTLs originated:</b>            | <b>0</b>   | <b>Alt border DTLs originated:</b>   | <b>0</b>   |
| <b>Route failures:</b>                 | <b>13</b>  | <b>Border failures:</b>              | <b>0</b>   |
| <b>Route unreachable errs:</b>         | <b>0</b>   | <b>Border unreachable errs:</b>      | <b>0</b>   |
| <b>PNNI AAL Discards:</b>              | <b>0</b>   |                                      |            |
| <b>Max Concurrent Pt-Mpt Calls</b>     | <b>500</b> |                                      |            |

**Nodes in this switch.** The number of instances of the PNNI protocol within this OmniSwitch. In a single peer group network, only one instance of PNNI, or node, can exist in a single OmniSwitch.

**Neighbors detected.** The number of neighboring nodes attached to this node. Neighbor nodes are physically connected to this node via a CSM port.

**PNNI highest version supported.** The highest version of the PNNI protocol that the software in this OmniSwitch is capable of executing. The current release supports PNNI version 1.0.

**Lowest version.** The lowest version of the PNNI protocol that the software in this OmniSwitch is capable of executing.

**Node routing database size.** The current number of valid pre-calculated PNNI routes to nodes.

**Address database size.** The current number of valid PNNI routes from nodes in this PNNI routing domain to exterior ATM addresses and transit networks.

**PTSEs in database.** The total number of PNNI Topology State Elements (PTSEs) in this node's topology database. PTSEs and the topology database are discussed in *PNNI Network Initialization* on page 45-18 and *Establishing a Connection* on page 45-21.

**RCCS in database.** The number of Routing Control Channels (RCCs) currently in the database. RCCs are used to route PNNI management information, such as PNNI Topology State Packets (PTSPs) and Database Summary packets.

**Times SPF executed.** The number of times the Shortest Path First (SPF) algorithm has been executed to compute paths for call connections.

**Pt to Pt calls in progress.** The number of point-to-point calls that are being set up right now. Once a call has been set up it will not be included in this count.

**Pt to Multi-Pt calls in prog.** The number of point-to-multipoint calls that are being set up right now. Once a call has been set up it will not be included in this count.

**Conns cranked back to this sys.** The total number of connection setup messages, including DTL stacks originated by this node, that have cranked back to this node at all levels of the PNNI hierarchy. Connection setups will crankback when one node along the pre-computed path is not able to set up the connection due to bandwidth, Quality of Service, or other considerations.

**Conns cranked from border.** The total number of connection setup messages, including DTLs by this node as an entry border node, that have cranked back to this node at all levels of the PNNI hierarchy. This value does not include crankbacks for which this node was not the crankback destination. It includes only those crankbacks that were directed to this node. In a single peer group network, this value will be 0.

**DTL stacks in use.** The number of Designated Transit Lists (DTLs) that are in use right now to set up point-to-point and point-to-multipoint PNNI calls. If no calls are being set up, then this field will read zero (0).

**DTL Stacks Free.** The number of Designated Transit List (DTLs) that are available to set up PNNI calls. This value does not indicate how many DTLs are in use; it indicates the number of DTLs in the PNNI database that may be put to use for settings up calls.

**Total DTL stacks originated.** The total number of DTL stacks that this OmniSwitch has originated and placed into signalling messages. This value includes the initial DTL stacks computed by this node as well as any alternate DTL routes (second, third choice, etc.) computed in response to crankbacks.

**Total DTL borders originated.** The number of partial DTL stacks that this OmniSwitch has added into signalling messages as an entry border node. This includes the initial partial DTL stacks computed by this system as well as any alternate route (second, third choice, etc.) partial DTL stacks computed by this node in response to crankbacks. In a single peer group network, this value will be 0.

**Alt DTLs originated.** The total number of alternate DTL stacks that this node has computed and placed into signalling messages as the DTL Originator.

**Alt Border DTLs originated.** The total number of alternate partial DTL stacks that this node has computed and placed into signalling messages as an entry border node. In a single peer group network, this value will be 0.

**Route failures.** The number of times this node failed to compute a viable DTL as the originator of a call. This value further indicates the number of times a call was cleared from this node due to originator routing failure.

**Border failures.** The number of times this node failed to compute a viable partial DTL stack as an entry border node for a call. This value indicates the number of times a call was either cleared or cranked back from this node due to border routing failure. In a single peer group network, this value will be 0.

**Route unreachable errs.** The total number of times this node failed to compute a viable DTL stack as the DTL originator because the destination was unreachable (i.e., those calls that are cleared due to an unreachable transit network or destination).

**Border unreachable errs.** The number of times this node failed to compute a viable partial DTL stack as an entry border node because the target of the path calculation was unreachable (i.e., those calls that are cleared are cranked back due to an unreachable transit network or destination). In a single peer group network, this value will be 0.

**PNNI AAL Discards.** An indicator of congestion-related problems. This value should always be zero (0).

**Max Concurrent Pt-Mpt Calls.** The maximum number of point-to-multipoint calls that this switch can support.

## Viewing Node-Specific Information

The **pninfo** command displays several current configuration values for this node, such as its Node ID and ATM address. For the single-peer group version of PNNI, enter

**pninfo**

at the system prompt. For the multiple-peer group version of PNNI, the syntax for this command is as follows:

**pninfo [ <node level>]**

The **<node level>** option allows you to display the current configuration for a specific node. If you do not use this option, then the **pninfo** command will display the lowest level node. For example, when you enter

**pninfo**

a screen similar to the following displays:

### ATM PNNI Node Information

```

Node ID:          50a0 3903488001bc900001020000090020da00000900
ATM address of this Node:      3903488001bc900001020000090020da00000900
Node level:       50 (80 decimal)
Peer group ID:    50 3903488001bc90000102000000

PNNI node index:   1
Administrative status: ENABLED           Operational status: UP
Lowest level node: True                  Restricted transit: False
Complex representation: False             Restricted branching: False

Address summarization: Configured (and operational)
Summary address:    3903488001bc9000010178ae0

PGL Priority:       50 (PGL Increment is a constant 50)
PGL State:          NORMAL OPERATION: I'M NOT PGL
Preferred PGL:      50a03903488001bc900001020000080020da00000800
Parent Node Id:     0000000000000000000000000000000000000000000000000000000
Parent ATM Addr:    0000000000000000000000000000000000000000000000000000000
Parent PG Id:       00 00000000000000000000000000000000
Parent PGL Node ID: 0000000000000000000000000000000000000000000000000000000

```

The fields displayed by the **pninfo** command are described below.

◆ **Note** ◆

The last set of parameters (**PGL Priority** through **Parent PGL Node ID**) will not display unless you have installed the software for multiple-peer group PNNI.

**Node ID.** The value used to identify this node within the PNNI network. If this node is a lowest-level node, then the Node ID is as follows: the first octet equals the node level within the PNNI hierarchy, the second octet equals 160, and the last 20 octets equal the node's ATM address.

If this is not a lowest-level node (i.e., it is an LGN), then the Node ID is as follows: the first octet equals the node level within the PNNI hierarchy, the next 14 octets equal the Peer Group ID for the Peer Group Leader (PGL) node connected to this LGN, the next six octets equal the End System Identifier (ESI) of the physical switch implementing this LGN functionality, and the last octet equals zero.

◆ **Note** ◆

In a single peer group network, the Node ID is computed using the lowest-level node convention.

**ATM address of this Node.** This node's ATM address. Remote systems must direct packets or calls to this address to exchange PNNI protocol packets with this node.

**Node level.** The level within the PNNI hierarchy where this node exists. This value may range from 0 to 104 with higher values indicating nodes lower in the PNNI hierarchy. This level is used to determine the default node ID and the default peer group ID for this node. The default node level is 80 decimal. In single peer group operation, all nodes will be at the same level.

**Peer group ID.** The peer group identifier of the peer group to which this node will become a member. The default value of this ID is as follows: the first octet is the level within the PNNI hierarchy where nodes in this peer group are located and the next 13 octets are the prefix for the ATM End System Address of the node.

**PNNI node index.** The value assigned to this node to identify itself to SNMP management software.

**Administrative status.** Indicates the Administrative Status of this node. **ENABLED** means that the node is allowed to become operationally active and participate in PNNI protocol exchanges. **DISABLED** means the node will be inactive and not participate in PNNI protocol exchanges.

**Operational status.** Indicates whether this node is active (**UP**) or whether it has become non-operational (**DOWN**). When **DOWN**, all state information is cleared from the node and the node is not communicating with any of its neighbor nodes.

**Lowest level node.** Indicates whether this node acts as a lowest-level node or whether it is a Logical Group Node (LGN) that becomes active when one of the other nodes in this peer group becomes a Peer Group Leader (PGL). A value of **False** indicates nodes that are capable of becoming Logical Group Nodes. In a single peer group network, all nodes will be lowest-level nodes and this value will be **True**.

**Restricted transit.** Indicates whether this node is restricted from supporting Switched Virtual Circuits (SVCs) transversing this node. **False** means this node can support ATM data links transversing this node for another destination. **True** means this node will be restricted from setting up SVCs unless overridden by another PNNI parameter.

**Complex representation.** Specifies whether this node uses complex node representation. **True** indicates complex representation is used. **False** indicates that simple node representation is used.

**Restricted branching.** Indicates whether the originating node is able to support additional multicast virtual circuit branches. **False** means that the node can support additional multicast branches. **True** means that additional branches are not supported because the maximum number of multicast virtual circuits on all modules in the node has been reached. The maximum number of multicast virtual circuits supported by a CSM-155 module is 8000, and the maximum supported by a CSM-622 module is 16,000.

**Address summarization.** Indicates whether this node uses summarization when advertising the addresses of attached devices to other PNNI nodes. Using address summarization to advertise internal reachability speeds PNNI database searches.

**Summary address.** The summary address that will be used to advertise all devices attached to this node. The summary address is the only address advertised, reducing the PNNI database size and the amount of information exchanged in PTSEs. This address prefix is also used by ILMIs when registering clients.

◆ **Note** ◆

The following fields only apply to the multiple-peer group version of PNNI.

**PGL Priority.** The leadership priority value advertised by the local node. In the election process to determine the PGL (see *Peer Group Leader (PGL) Election Algorithm* on page 45-8 for more information), lower values have higher priority than higher values. The message **PGL Increment is a constant 50** refers to the condition when a node is elected PGL, it will increment its PGL priority by 50 to prevent hung elections.

You can set this value with the **pncfg** command, which is described in *Configuring Node-Specific Parameters* on page 45-39.

**PGL State.** Indicates if this node is a lowest-level node, a Peer Group Leader (PGL) node, or a potential PGL node. The possible PGL states are described below:

**starting.** This node has begun participation in the election of a PGL.

**awaiting.** This node has sent Hello messages on at least one link but no peer has been found.

**awaitingFull.** A least one neighboring peer has been found but the database synchronization process has been completed yet.

**initialDelay.** Database synchronization has been completed with at least one neighboring peer.

**calculating.** This node is calculating its new choice for peer group leader.

**awaitUnanimity.** This node has chosen itself as PGL. It will check to see if all other nodes in the peer group have also elected it as PGL. This node will wait for a unanimous decision by all the nodes in the peer group or after a set period of time before declaring itself as the PGL.

**NORMAL OPERATION: I'M PGL.** This node is the PGL. It will continue to examine PTSEs to determine if other nodes in the peer group have a higher PGL priority.

**NORMAL OPERATION: I'M NOT PGL.** This node is not the PGL. It will continue to examine PTSEs to determine which node has the highest priority to be PGL.

**hungElection.** This node has chosen itself as PGL but it has failed to be elected PGL by at least 2/3 of the nodes in the peer group. This situation will be resolved by the node changing its preferred PGL or other nodes will accept it as the PGL.

**awaitReElection.** This node has lost its connection to the current PGL. A new election process will be started.

**Preferred PGL.** The ATM address of the node that this local believes should be or should become the preferred Peer Group Leader (PGL). If a PGL has not been chosen, then this field will display all zeroes.

**Parent Node Id.** When the local node is the Peer Group Leader (PGL), then this field will display the node ID of the parent Logical Group Node (LGN). If the local node is not the PGL of its peer group, then this field will display all zeroes.

**Parent ATM Addr.** When the local node is the Peer Group Leader (PGL), then this field will display the ATM address of the parent Logical Group Node (LGN). If the local node is not the PGL of its peer group, then this field will display all zeroes.

**Parent PG Id.** When the local node is the Peer Group Leader (PGL), then this field will display the local node's parent peer group ID. If the local node is not the PGL of its peer group, then this field will display all zeroes.

**Parent PGL Node ID.** When the local node is the Peer Group Leader (PGL), then this field will display the ID of the node elected as PGL in the parent peer group. If the local node is not the PGL of its peer group, then this field will display all zeroes.

## Viewing Timer Information

The **ptinfo** command displays current configuration values for PTSE and Hello timers. These values are configured through the **pgcfg** and **pncfg** commands. For the single-peer group version of PNNI, enter

```
ptinfo
```

at the system prompt. For the multiple-peer group version of PNNI, the syntax for this command is as follows:

```
ptinfo [ <node level>]
```

The **<node level>** option allows you to display the current configuration for a specific node. If you do not use this option, then the **ptinfo** command will display the lowest level node. For example, when you enter

```
ptinfo
```

a display similar to the following displays:

```

                PNNI Timer Information
      (All time values are seconds unless otherwise specified)

PNNI node index:                1

PTSE refresh interval:          1800
PTSE lifetime factor:           2 (refresh intervals)
PTSE hold down interval:        1
PTSE delayed acknowledgement timer: 1

PTSP transmit timer:            10
PTSE request re-transmit timer:  1
PTSE aging timer:               10
Database summary re-transmit interval: 3

Hello interval:                 15
Hello hold down interval:        3
Hello inactivity factor:         5 (hello intervals)

Link AvCR proportional multiplier: 50%
Link AvCR minimum threshold:      3%
Link CTD proportional multiplier:  50%
Link CDV proportional multiplier:  25%
```

**PNNI node index.** A value used by SNMP to identify this node in the PNNI network.

**PTSE refresh interval.** The time, in seconds, before a self-originated PTSE is updated. PTSEs are aged out of the database unless refreshed by the originating node. The lifetime of a PTSE is determined by multiplying the Refresh Interval by the Lifetime Factor. The range for this value is 1 to 32,767 seconds. The default value is 1800 seconds.

**PTSE lifetime factor.** The value for the PTSE lifetime multiplier expressed as a percentage. Valid values are integers ranging from 1 to 255. This value helps determine the initial lifetime of a PTSE. The Lifetime Factor multiplied by the PTSE Refresh Interval is the initial lifetime of a PTSE. The default value is 2.

**PTSE hold down interval.** The minimum time, in seconds, before which this node can refresh PTSEs. A node can prevent a PTSE from aging out of the topology database by refreshing it. This holddown value limits the node from refreshing PTSEs too often and exhausting database space too quickly.



**PTSE delayed acknowledgement timer.** When a node receives a PTSE from another node it sends back an Acknowledgment packet. However, this acknowledgment is not immediate. This value is the amount of time between the receipt of a PTSE and its acknowledgment.

**PTSP transmit timer.** PTSPs are sent until they are acknowledged by neighboring nodes. This variable is the amount of time, in seconds, between successive transmissions of PTSPs. If a PTSP is acknowledged before this time interval, then a retransmission will not be sent.

**PTSE request re-transmit timer.** The time interval between the receipt of Database Summaries and the sending of PTSE Request packets. Database Summary packets contain an index of the PTSEs in a node's topology database. Other nodes use Database Summary packets to request PTSEs. If a node requires a PTSE listed in a Database Summary packet, then it will request that PTSE in the form of a PTSE Request packet.

**PTSE aging timer.** The time, in seconds, before a given PTSE entry in the topology database is aged out of the database. The PTSE may be refreshed by the originating node before this entry ages out. The timer may range from 1 to 255 seconds.

**Database summary re-transmit interval.** The time, in seconds, before this node will re-transmit a Database Summary packet that has gone unacknowledged by another node.

**Hello interval.** The value for the Hello timer in seconds. In the absence of triggered Hellos, this node will send one Hello packet on each of its ports at the interval specified here. Values can range from 1 to 255 seconds.

**Hello hold down interval.** The value for the Hello hold down timer. This node will use this value to limit the rate at which it sends Hello messages. Valid values range from 1 to 255 seconds.

**Hello inactivity factor.** The number of Hello intervals that may pass without receiving a Hello before the neighboring node is determined to have gone down. Valid values range from 1 to 255.

**Link AVCR proportional multiplier.** This variable is described in *Configuring General PNNI Parameters* on page 45-31.

**Link AVCR minimum threshold.** This variable is described in *Configuring General PNNI Parameters* on page 45-31.

**Link CTD proportional multiplier.** This variable is described in *Configuring General PNNI Parameters* on page 45-31.

**Link CDV proportional multiplier.** This variable is described in *Configuring General PNNI Parameters* on page 45-31.

# Viewing PNNI Neighbor Information

The **pnbrs** command displays information on neighbor nodes connected to this OmniSwitch. It includes Node IDs of the neighboring nodes as well as statistics on PTSE communication between this node and its neighbors. For the single-peer group version of PNNI, the syntax for this command is as follows:

```
pnbrs [-s]
```

For the multiple-peer group version of PNNI, the syntax for this command is as follows:

```
pnbrs [-s] [<node level>]
```

The **-s** option provides a summary version of the **pnbrs** command (see *Summary Form of pnbrs* on page 45-59). The **<node level>** option allows you to display the current configuration for a specific node. If you do not use this option, then the **pnbrs** command will display the lowest level node. For example, when you enter

```
pnbrs
```

a screen similar to the following displays:

| PNNI Port Neighbor Information |  |           |           |                              |                 |
|--------------------------------|--|-----------|-----------|------------------------------|-----------------|
| Neighbor:                      | 50a03903488001bc900001010230e00020da0230e000 |           |           |                              | Nbr State: FULL |
| PTSP                           | PTSE ACKS                                    | PTSE REQs | DB Sums   | Port Count: 1                |                 |
| Rcvd/Xmtd                      | Rcvd/Xmtd                                    | Rcvd/Xmtd | Rcvd/Xmtd | Local Ports To This Neighbor |                 |
| <hr/>                          |  |           |           |                              |                 |
| 2/                             | 5/   | 1/        | 2/        | 5/2                          |                 |
| 5                              | 2  | 1         | 5         |                              |                 |
|                                |  |           |           |                              |                 |
| Neighbor:                      | 50a03903488001bc9000010175ee100020da75ee1000 |           |           |                              | Nbr State: FULL |
| PTSP                           | PTSE ACKS                                    | PTSE REQs | DB Sums   | Port Count: 1                |                 |
| Rcvd/Xmtd                      | Rcvd/Xmtd                                    | Rcvd/Xmtd | Rcvd/Xmtd | Local Ports To This Neighbor |                 |
| <hr/>                          |  |           |           |                              |                 |
| 2/                             | 3/   | 1/        | 2/        | 5/1                          |                 |
| 4                              | 2  | 1         | 5         |                              |                 |

**Neighbor.** The Node ID of the neighboring peer node. This command displays a separate Node ID and a separate listing for each neighbor of this node.

**Nbr State.** The state of the neighboring node’s Peer State Machine. The neighboring Peer State Machine manages exchanges of Hello and topology state packets. In addition, it describes the state of database synchronization and flooding ongoing with the neighboring peer. Possible states are as follows:

- DOWN** No active links to the neighboring peer node.
- NEGOTIATING** The first step in setting up a link between two neighbor nodes. The master node is determined during this step.
- EXCHANGING** The node describes its topology database to its neighboring peer in the form of Database Summary packets.
- LOADING** All Database Summary packets have been exchanged and all required PTSEs have been requested, but not all PTSEs have been received.
- FULL** All PTSEs have been received from the neighboring peer node. Links to this node can now be advertised to other nodes via PTSEs.

**PTSP Rcvd/Xmtd.** The number of PNNI Topology State Packets (PTSPs) received from and transmitted to the neighboring peer node. PTSPs received are listed on the first line and PTSPs transmitted are listed on the bottom line.

**PTSE Ack Rcvd/Xmtd.** The number of PNNI Topology State Element (PTSE) Acknowledgment packets received from and transmitted to the neighboring peer node. Received PTSE Acknowledgments are listed on the first line and transmitted PTSE Acknowledgments are listed on the bottom line.

**PTSE REQ Rcvd/Xmtd.** The number of PTSE Request packets received from and transmitted to the neighboring peer node. Received PTSE Requests are listed on the first line and transmitted PTSE Acknowledgments are listed on the bottom line.

**DB Sum Recvd/Xmtd.** The number of Database Summary packets received from and transmitted to the neighboring peer node. Received Database Summaries are listed on the first line and transmitted Database Summaries are listed on the bottom line.

**Port Count.** The number of ports on this node that connect to the neighboring peer node. If the neighboring peer only communicates via an SVCC-based RCC, then the value of this variable will be zero. Otherwise it is set to the total number of ports connected to the neighboring peer that are in the Hello state, 2-WayInside. The ports included in this count are listed under the **Local Ports To This Neighbor** column.

**Local Ports To This Neighbor.** The local ports (i.e., ports on this OmniSwitch) connected to the neighboring peer node that are in the Hello state, 2-WayInside. In this Hello state, bi-directional communication between the two nodes has been achieved. The nodes are in the same peer group. Database summary packets, PTSE Request packets, PTSPs, and PTSE Acknowledgment packet can be transmitted over this link. If this node is an LGN, then this field will display **Over SVC/RCC** in this field.

## Summary Form of pnbrs

The **pnbrs** command also has a summary option that allows you to view less information on each neighbor node. Simply enter the command

```
pnbrs -s
```

to obtain a display similar to the following:

### PNNI Port Neighbor Information

```
3903488001bc9000010178aee00020da78aee000 NEGOTIATING 3/1 3/5 3/6 3/3
```

This display is a condensed form of the full **pnbrs** display. It shows the Node ID for each neighbor, the current state of the neighbor, and the local port connected to that neighbor. Full definitions of each of these pieces of information can be found in the above explanations.

## Viewing Port Information

The **ppinfo** command displays information on all CSM ports (about which PNNI is aware) in this OmniSwitch. Information includes the port type, the port's current Hello State, and the Administrative Weight of the attached link. When you enter

**ppinfo**

a screen similar to the following displays:

### PNNI Port Information

| PNNI<br>SI/Prt<br>PortId | VPI//VCI<br>Type | VP<br>Cap | Neighbor<br>State | Advertised<br>Max CR/<br>Avail CR | Admin Weight            |          |                          |
|--------------------------|------------------|-----------|-------------------|-----------------------------------|-------------------------|----------|--------------------------|
| =====                    | =====            | ==        | =====             | =====                             | =====                   | =====    | =====                    |
| 5/ 2<br>(257)            | 0/18<br>OC3      | n         | RCC Unavail       | 350000                            | CBR:5040<br>rt-VBR:5040 | ABR:5040 | UBR:5040<br>nrt-VBR:5040 |
| 7/ 1<br>(384)            | 0/18<br>OC3      | n         | LOADING           | 350000                            | CBR:5040<br>rt-VBR:5040 | ABR:5040 | UBR:5040<br>nrt-VBR:5040 |

**PNNI SI/Prt PortId.** The slot and port for this CSM port. The first number is the slot number for the CSM module and the second number (after the slash) is the port number on the CSM module. The number in parentheses is the internal identification for this port.

**Vpi/Vci.** The Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI) for the Routing Control Channel (RCC) defined here. The VPI is listed first, followed by a slash (/), and then the VCI.

The port type is listed under the VPI/VCI. This information indicates whether this is an OC-3c/STM-1c port connection (155 Mbps) or an OC-12c/STM-4c (622 Mbps) port connection. If this entry is describing the internal port on an FCSM module or an ATM uplink port on an ASM module, then this column will read **ASM**.

**VpCap.** Indicates whether this port advertises to other nodes that it supports the establishment of Virtual Paths. **Y** means the port advertises that Virtual Paths can be established on this physical link. **N** means that this port does not advertise that it is capable of setting up Virtual Paths.

**Neighbor state.** Indicates the state of the neighboring node. Possible states are as follows:

|                    |   |
|--------------------|---|
| <b>DOWN</b>        | No active links to the neighboring peer.  |
| <b>NEGOTIATING</b> | The first step in setting up a link between two neighbors.  |
| <b>EXCHANGING</b>  | The node describes its topology database to its neighboring peer in the form of Database Summary packets.                             |
| <b>LOADING</b>     | All Database Summary packets have been exchanged and all required PTSEs have been requested, but not all PTSEs have been received.    |
| <b>FULL</b>        | All PTSEs have been received from the neighboring peer. Links to the neighboring peer can now be advertised to other nodes via PTSEs. |
| <b>RCC Unavail</b> | The Routing Control Channel is not available on his link. The link may not be supported by PNNI.                                      |
| <b>Border Node</b> | The neighboring node is a border node.  |

**Advertised Max CR/Avail CR.** The first line is the maximum cell rate (in cells per second) for each QoS on this port. This value is determined by hardware. The second line is the currently available bandwidth for each QoS; this value may be less than the maximum cell rate due to usage on the port. The maximum cell rate, in cells per second, for each CSM port type is as follows:

- OC-3 ports                      350,000
- OC-12 ports                    1,400,000
- ATM 25 Mbps ports        50,000

**AdminWt.** The Administrative Weight assigned to each Class of Service on this port. Administrative Weight is assigned through the **ppcfg** command.

## Summary Form of ppinfo

The **ppinfo** command also has a summary option that allows you to view less information on each PNNI port. Simply enter the command

**ppinfo -s**

to obtain a display similar to the following:

### PNNI Port Information Summary

| SI/Port<br>[/Inst] | Phys<br>Port | State       | SI/Port<br>[/Inst] | Phys<br>Port | State |
|--------------------|--------------|-------------|--------------------|--------------|-------|
| =====              | =====        | =====       | =====              | =====        | ===== |
| 3/ 1               | (128)        | FULL        | 3/ 3               | (144)        | FULL  |
| 3/ 5               | (160)        | FULL        | 3/ 6               | (168)        | FULL  |
| 3/ 2/1             | (701)        | RCC UNAVAIL |                    |              |       |

This display is a condensed form of the full **ppinfo** display. It shows physical port information and the state of the neighbor currently attached to the local port. Full definitions of each of these pieces of information can be found in the above explanations.

## Viewing Link Information

The **plink** command contains information on the logical links attached to this OmniSwitch and the relationship with the neighbor nodes on the other end of the links. For the single-peer group version of PNNI, the syntax for this command is as follows:

**plink [-s]**

For the multiple-peer group version of PNNI, the syntax for this command is as follows:

**plink [-s] [ <node level>]**

The **-s** option provides a summary version of the **plink** command (see *Summary Form of plink* on page 45-64). The **<node level>** option allows you to display the current configuration for a specific node. If you do not use this option, then the **plink** command will display the lowest level node. For example, when you enter

**plink**

a screen similar to the following displays:

| PNNI Link Table   |                  |                   |  |                      |                            |
|-------------------|------------------|-------------------|--|----------------------|----------------------------|
| Lcl/Rmt<br>PortId | S/P/I<br>IfIndex | State/<br>Version | Remote Node Id                               | Hellos<br>Xmtd/Recvd | Link Type                  |
| 128/<br>128       | 3/1<br>30100     | 2 WAY INSIDE<br>1 | 3903488001bc90000101<br>7a1bd00020da7a1bd000 | 1295/<br>1287        | Lowest Level<br>Horizontal |
| 144/<br>144       | 3/3<br>30300     | 2 WAY INSIDE<br>1 | 3903488001bc90000101<br>7a1bd00020da7a1bd000 | 1295/<br>1286        | Lowest Level<br>Horizontal |
| 160/<br>160       | 3/5<br>30500     | 2 WAY INSIDE<br>1 | 3903488001bc90000101<br>7a1bd00020da7a1bd000 | 1295/<br>1286        | Lowest Level<br>Horizontal |

**Lcl/Rmt PortId.** The first line in each row of this column is the PNNI Port Identifier of the local port on this link as selected by the local node. This value is relevant to the local switch only. Different ATM switches may use different formulas to derive this value.

The second line in each row of this column is the Port Identifier at the remote end of this link as assigned by the remote node. If the Link Type is **Outside Uplink**, then this ID is assigned by the lowest-level neighbor node to identify the outside link. If the Remote Port ID is unknown or the link type is **Uplink**, then this value is set to zero.

**S/P/I If Index.** The first line in each row of this column is the CSM slot and port for link. The second line is the IfIndex for this interface. For horizontal and outside links between the lowest-level nodes and for links of an unknown type, the IfIndex identifies the physical interface to which this logical link corresponds. For all other links, this value is zero.

**State.** Indicates the state of the Hello protocol exchange over this link. Possible values are as follows:

|                     |  |
|---------------------|--|
| <b>DOWN</b>         | No PNNI Routing packets sent or received over the link.  |
| <b>ATTEMPTING</b>   | Attempts were made to contact the neighboring node with Hello messages, but no valid Hellos have been received from the neighbor.  |
| <b>1 WAY INSIDE</b> | A Hello has been received from the neighboring node. Both nodes are members of the same peer group, but the node and port identifiers from the neighbor are set to zero. |

|                       |   |
|-----------------------|---|
| <b>2 WAY INSIDE</b>   | Bi-directional communication between the two nodes on this link has been achieved. The nodes are in the same peer group. Database summary packets, PTSE Request packets, PTSPs, and PTSE Acknowledgment packets can be transmitted over this link.  |
| <b>1 WAY OUTSIDE</b>  | A Hello has been received from the neighboring node. The nodes are members of different peer groups. The node and port identifiers from the neighbor are set to zero. This node will now search for a common peer group that contains both nodes.   |
| <b>2 WAY OUTSIDE</b>  | A Hello has been received from the neighboring node. The nodes are members of different peer groups. Valid node and port identifiers have been received, but a common peer group between the two nodes has not been identified. This node will now search for a common peer group that contains both nodes. |
| <b>COMMON_OUTSIDE</b> | Bi-directional communication between the two nodes on this link has been achieved. This link can now be advertised through PTSEs.   |

**Version.** Indicates the version of the PNNI Routing protocol used to exchange information over this link. If communication with the neighbor node has not yet been established, then this value is set to **Unknown**. The version of PNNI supported in this release is 1.0.

**Remote Node Id.** Indicates the Node ID of the remote (neighboring) node on the other end of the link. If the Link Type is **Outside Uplink**, then this is the Node ID of the lowest-level neighbor node on the other end of the outside link. If the remote node ID is unknown or if the Link Type is **Uplink**, then this variable is set to all zeros.

**Hellos: Transmitted and Received.** For horizontal links between lowest-level nodes, these values indicate the number of Hello packets transmitted or received over this link.

**Type.** The type of link described in this display. The following types are possible:

|                                |   |
|--------------------------------|---|
| <b>Unknown</b>                 | Unknown type of link.   |
| <b>Lowest Level Horizontal</b> | Lowest-level node, horizontal link.   |
| <b>Lowest Level Outside</b>    | Lowest-level node connected via an outside link.  |
| <b>Hor link to LGN</b>         | Horizontal link to a Logical Group Node. (Not supported in a single peer group network).              |
| <b>Uplink</b>                  | Uplink to a Logical Group Node in another peer group. (Not supported in a single peer group network.) |
| <b>Outside uplink</b>          | An outside link that is also an uplink. (Not supported In a single peer group network.)               |

Summary Form of plink

The **plink** command also has a summary option that allows you to view less information on the logical links connected to this node. Simply enter the command

```
plink -s
```

to obtain a display similar to the following:

| PNNI Link Table |            |              |                             |
|-----------------|------------|--------------|-----------------------------|
| Lcl PortId      | Rmt PortId | State        | Link Type                   |
| =====           | =====      | =====        | =====                       |
| 128             | 128        | 2 WAY INSIDE | Lowest Level and Horizontal |
| 136             | 136        | 2 WAY INSIDE | Lowest Level and Horizontal |

This display is a condensed form of the full plink display. It shows the local port ID, the remote port ID, the state of the Hello protocol exchange over this link, and the type of link. Full definitions of each of these pieces of information can be found in the above explanations.



## Viewing the PTSE Database

The **pptse** command displays the kind of information contained in a node's topology database. A separate entry is provided for each PNNI Topology State Element (PTSE) in the database. Output from this command can become lengthy depending on the complexity of the network. The number of PTSEs is dependent on the number of nodes in the network as well as the number of End Systems, routes, and topology metrics.

This command also has a verbose mode that displays a hexadecimal dump of the topology database contents. This hex dump displays actual database entries as they reside in the topology database. Examples of the standard and verbose mode output from this command are given below.

### Standard Output

For the single-peer group version of PNNI, the syntax for this command is as follows:

```
pptse [-s | -v]
```

For the multiple-peer group version of PNNI, the syntax for this command is as follows:

```
pptse [-s | -v] [<node level>]
```

The **-s** option provides a summary version of the **pptse** command (see *Standard Output* on page 45-65) and the **-v** option provides a verbose option (see *Verbose Mode Output* on page 45-66). The **<node level>** option allows you to display the current configuration for a specific node. If you do not use this option, then the **pptse** command will display the lowest level node. For example, when you enter

```
pptse
```

a screen similar to the following displays:

#### PNNI PTSE Database Summary

```
Node ID:50a03903488001bc900001010230e00020da0230e000    PTSE:  1
Checksum: 1a81    Remaining lifetime: 3129 seconds        Seq #: 1
Information Group Carried: Nodal IG
=====
Node ID:50a03903488001bc900001010230e00020da0230e000    PTSE:  3
Checksum: 58da    Remaining lifetime: 3519 seconds        Seq #: 1
Information Group Carried: Horizontal Links IG
=====
Node ID:50a03903488001bc9000010175ee100020da75ee1000    PTSE:  1
Checksum: 5408    Remaining lifetime: 2888 seconds        Seq #: 1
Information Group Carried: Nodal IG
=====
Node ID:50a03903488001bc9000010175ee100020da75ee1000    PTSE:  2
Checksum: 5a49    Remaining lifetime: 3168 seconds        Seq #: 1
Information Group Carried: Horizontal Links IG
=====
Node ID:50a03903488001bc90000101761c900020da761c9000    PTSE :  1
Checksum: 53ab    Remaining lifetime: 3490 seconds        Seq #: 1
Information Group Carried: Nodal IG
=====
Node ID:50a03903488001bc90000101761c900020da761c9000    PTSE :  2
Checksum: d8db    Remaining lifetime: 3520 seconds        Seq #: 1
Information Group Carried: Horizontal Links IG
=====
```

**Node ID.** The Node Identifier of the node that originated the PTSE.

**PTSE.** The PTSE Identifier assigned to this PTSE by the node that originated the PTSE.

**Checksum.** The value of the PTSE checksum as it appears in the local topology database.

**Remaining lifetime.** The remaining lifetime, in seconds, for this PTSE in the local node's topology database.

**Seq #.** The sequence number for this PTSE as it appears in this node's local topology database. This value differs from the PTSE ID as it is defined by the recipient of the PTSE (i.e., this node). The PTSE ID is assigned by the sender of the PTSE (i.e., the originating neighbor node).

**INFORMATION GROUPS CARRIED.** The type of information contained in this PTSE.

### Verbose Mode Output

You can obtain the actual contents for each PTSE in the topology database through the **pptse** verbose mode. Note that output can become quite lengthy.

When you enter

```
pptse -v
```

a screen similar to the following displays:

#### PNNI PTSE Database Summary

Node ID:50a03903488001bc900001010230e00020da0230e000 PTSE: 1

Checksum: 1a81 Remaining lifetime: 3019 seconds Seq #: 1

Information Group Carried: Nodal IG

PTSE contents:

```
0000: 00 02 00 70 01 01 01 00 60 a0 39 03 48 80 01 bc
0010: 90 00 01 01 72 b1 b0 00 20 da 72 b1 b0 00 50 39
0020: 03 48 80 01 bc 90 00 01 01 01 00 00 00 40 00 44
0030: 00 61 00 00 00 00 01 00 00 00 01 1a 81 0b cb
0040: 00 61 00 50 39 03 48 80 01bc 90 00 01 00 01 00
0050: 01 00 20 da 72 b1 b0 00 00 00 00 00 00 00 00
0060: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

Node ID:50a03903488001bc900001010230e00020da0230e000 PTSE: 3

Checksum: 58da Remaining lifetime: 3409 seconds Seq #: 1

Information Group Carried: Horizontal Links IG

PTSE contents:

```
0000: 00 02 00 c0 01 01 01 00 50 a0 39 03 48 80 01bc
0010: 90 00 01 00 01 00 01 00 20 da 72 b1 b0 00 50 39
0020: 03 48 80 01 bc 90 00 01 01 01 00 00 00 40 00 94
0030: 01 20 00 00 00 00 03 00 00 00 01 58 da 0d 51
0040: 01 20 00 80 80 00 50 a0 39 03 48 80 01bc 90 00
0050: 01 01 01 00 01 00 20 da 76 1c 90 00 00 00 01 80
0060: 00 00 00 c0 00 00 00 01 00 80 00 2c 18 00 00 00
0070: 00 00 13 b0 00 05 57 30 00 04 93 e0 00 00 00 0a
0080: 00 00 00 02 00 08 00 00 00 a0 00 0c 00 00 00 64
0090: 00 00 00 02 00 80 00 2c e0 01 00 00 00 00 13 b0
00a0: 00 05 57 30 00 04 93 e0 00 00 00 0a 00 00 00 02
00b0: 00 08 00 00 00 a0 00 0c 00 00 00 64 00 00 00 02
```

A separate listing is provided for each PTSE. The top portion of each listing is the same information provided through the non-verbose **pptse** command. The additional information is a hexadecimal dump of the actual PTSE data.

### Verbose Information on a Single PTSE

You can also obtain the verbose mode output for a single PTSE. This display may be useful when you require detailed PTSE information, but not for every PTSE in the PNNI database. To request detailed information for a PTSE you will need to know its Node ID and PTSE ID. You can get these two pieces of information via the standard **pptse** output. After obtaining the Node ID and PTSE enter you **pptse** command as follows:

```
pptse -v <22-byte Node ID> -p <PTSE ID>
```

The following is an example of a command line specification:

```
pptse -v 50a03903488001bc900001010230e00020da0230e000 -p 3
```

This command line would display the same output as second PTSE in the sample verbose mode output shown above.

## Summary Mode Output

You can also obtain a summary display of all PTSEs associated with a given Node ID through the **pptse** summary mode. When you enter

**pptse -s**

a screen similar to the following displays:

```

                                PNNI PTSE Summary Table
                                (PTSE lifetime is displayed in seconds)

Node Id: 50a03903488001bc900001016cdb200020da6cdb2000
PG ID:   503903488001bc90000101000000
  PTSE:   1 (seq 512): Nodal IG                2868
  PTSE:   2 (seq 513): IReach ATM addresses IG  2868
  PTSE:   3 (seq 512): HLink                    2918
  PTSE:   4 (seq 512): HLink                    2918
  PTSE:   5 (seq 510): HLink                    1688
Node Id: 50a03903488001bc900001017def000020da7def0000
PG ID:   503903488001bc90000101000000
  PTSE:   1 (seq 557): Nodal IG                3459
  PTSE:   2 (seq 558): IReach ATM addresses IG  3459
  PTSE:   5 (seq 480): HLink                    3529
  PTSE:   6 (seq 479): HLink                    3529
Node Id: 50a03903488001bc900001017f6bc00020da7f6bc000
PG ID:   503903488001bc90000101000000
  PTSE:   1 (seq 668): Nodal IG                2820
  PTSE:   2 (seq 669): IReach ATM addresses IG  2820
  PTSE:   9 (seq 574): HLink                    2830
  PTSE:  10 (seq 574): HLink                    2830
  PTSE:  11 (seq 574): HLink                    2830
Node Id: 50a03903488001bc900001018519000020da85190000 (MYSELF)
PG ID:   503903488001bc90000101000000
  PTSE:   1 (seq 557): Nodal IG                3468
  PTSE:   2 (seq 558): IReach ATM addresses IG  3468
  PTSE:   3 (seq 557): HLink                    3518
  PTSE:   4 (seq 557): HLink                    3518
  PTSE:   5 (seq 557): HLink                    3518
  PTSE:   6 (seq 555): HLink                    2608
Node Id: 50a03903488001bc900001018519b00020da8519b000
PG ID:   503903488001bc90000101000000
  PTSE:   1 (seq 514): Nodal IG                3539
  PTSE:   2 (seq 515): IReach ATM addresses IG  3539
  PTSE:   5 (seq 37): HLink                     3549
  PTSE:   7 (seq 37): HLink                     3549
Node Id: 50a03903488001bc900001019238700020da23870000
PG ID:   503903488001bc90000101000000
  PTSE:   1 (seq 557): Nodal IG                3459
  PTSE:   2 (seq 558): IReach ATM addresses IG  3459
  PTSE:   3 (seq 557): HLink                    3519

```

This summary display groups PTSEs with their associated Node IDs. Each PTSE entry shows the number of the PTSE on that node, the sequence number for the PTSE, the information groups in the PTSE, and then the remaining lifetime (in seconds) for the PTSE. See *Standard Output* on page 45-65 for a more complete description of each of these values.

## Viewing End-Point Adjacencies

The **padj** command displays adjacencies to this node. An adjacency is an End System or other ATM destination that is attached to a CSM port on this OmniSwitch node. Adjacencies are learned via ILMI. An adjacency differs from a neighbor in that a neighbor is an active participant in PNNI exchanges and is normally another node or switch.

When you enter

**padj**

a screen similar to the following displays:

### PNNI Adjacency Table

**Client:** 4700790000000000000000000000a03e00000100    **Advertised:** TRUE  
**Learnt:** TUE JAN 20 11:44:59 1998    **Slot/Port/Inst:** 3/8 (port 184)

The following adjacencies are summarized by this node's summary address  
(which is 3903488001bc900001017a1bd0):

| Native Address                           | PNNI Port     | Learned at Time     |
|--|---------------|---------------------|
| =====                                    | =====         | =====               |
| 3903488001bc900001017a1bd000041347561000 | 3/ 8    (184) | TUE JAN 20 11:44:59 |

**Client.** The 20-octet End System ATM address for the adjacency described by this entry.

**Learnt.** The date and time at which this adjacency was learned by this node. Node adjacencies are normally learned through ILMI.

**Advertised.** Since adjacencies are learned via ILMI, they can be learned at any time, including when PNNI is not operational in the OmniSwitch. If PNNI is not operational or if it has not been updated with this adjacency information, then this information is not advertised by PNNI throughout the network. In other words, information on this adjacency is not being passed via PTSEs to other nodes and this field would be False (**F**). If PNNI is operational and information on this adjacency is being passed to other nodes, then this field will be True (**T**).

Also, note that adjacencies are advertised on internally reachable destinations in PTSE IREACH frames.

**Slot/Port/Inst.** The CSM slot and port to which this adjacency is attached. The slot number is listed first, followed by a slash (/), and then the port number on the CSM module.

Some adjacencies may use address summarization to advertise their status. Those addresses that use summarization are listed with the following column descriptions:

**Native Address.** The full non-summarized 20-octet ATM address used to describe this adjacency.

**PNNI Port.** The physical slot and port number and the internal port number for this adjacency.

**Learned at Time.** The date and time at which the PNNI database learned about this adjacency.

## Configuring PNNI Scope Mapping Parameters

The **psmap** command is used to view and configure how the UNI/ILMI scope is mapped to the PNNI scope. UNI 4.0/ILMI 4.0 clients have a scope associated with an address. During ILMI address registration, ILMI will pass this scope along with an address to PNNI. In order to determine the scope of address advertisement, PNNI needs to map this ILMI scope to a PNNI scope. The default UNI/ILMI scope value is 15 (global) when ILMI scope is absent.

### ◆ Note ◆

You *must* install the software for multiple group operation to use this command.

For example, say an operational PNNI ATM network has three (3) levels (scope): 96, 80, 72. When ILMI registers address A.1.1 with scope 1 (which maps to PNNI scope 96), PNNI will only advertise this address in the PNNI level 96 network. When ILMI registers address A.1.2 with scope 6 (which maps to PNNI scope 72), then PNNI will advertise this address in level 72 (which is higher in the in the PNNI hierarchy than level 96), this address will known by more nodes.

To use the **psmap** command, enter

```
psmap
```

at the system prompt. A screen similar to the following will be displayed.

**PNNI Scope Mapping Table**

| Scope Name              | UNI/ILMI Scope | PNNI Scope/Level |
|-------------------------|----------------|------------------|
| =====                   |                |                  |
| 1) LocalNetwork         | 1              | 96               |
| 2) LocalNetworkPlusOne  | 2              | 96               |
| 3) LocalNetworkPlusTwo  | 3              | 96               |
| 4) SiteMinusOne         | 4              | 80               |
| 5) IntraSite            | 5              | 80               |
| 6) SitePlusOne          | 6              | 72               |
| 7) OrganizationMinusOne | 7              | 72               |
| 8) IntraOrganization    | 8              | 64               |
| 9) OrganizationPlusOne  | 9              | 64               |
| 10) CommunityMinusOne   | 10             | 64               |
| 11) IntraCommunity      | 11             | 48               |
| 12) CommunityPlusOne    | 12             | 48               |
| 13) Regional            | 13             | 32               |
| 14) InterRegional       | 14             | 32               |
| 15) Global              | 15             | 0                |

To configure a parameter, type "item = value" (as in 1=96)

To quit out of configuration, type "quit"

To save the configured info, type "save"

:

The PNNI mapping parameters listed by the **psmap** command are displayed in ascending order by hierarchy level. Each parameter is identified under the headings **Scope Name**, **UNI/ILMI Scope**, and **PNNI Scope/Level**. The **Scope Name** is a text description of the scope. The **UNI/ILMI Scope** is the UNI 4.0/ILMI 4.0 organizational scope. The lowest level is 1 and the highest level is 15. And the **PNNI Scope/Level** is the configurable PNNI routing level.

Higher values in the PNNI routing level indicate lower levels in the PNNI hierarchy. A logical node will only advertise its reachability to an organizational level that is greater than or equal to the node ID. Therefore, setting the PNNI routing level of an organizational level lower than the node's ID will prevent the node's address from being advertised at that level. Setting this value to 0 will prevent the node from advertising its reachability to this level.

All parameters are set to default values listed in ATM Forum PNNI Specification Version 1.0. To alter a parameter, enter the line number for the parameter, followed by an equal sign (=), and then the value for the parameter. For example, to change **Regional** from 32 to 48, enter:

**13=48**

at the prompt.

You can redisplay the parameters by entering a question mark (?). When you have completed configuring parameters, enter **save**. Your new values will be saved and you will exit this menu. If you want to exit this menu without saving changes, simply enter **quit**. If you have made any changes, then the following prompt will be displayed.

**Do you want these changes to take effect immediately? (y)**

Press **<Return>** to implement your changes immediately or press **n** to implement your changes at the next reboot.

The PNNI scope mapping parameters displayed by the **psmap** command are described below.

### 1) LocalNetwork

This level corresponds to the concept of a physical network. The default value is 96.

### 2) LocalNetworkPlusOne

This level corresponds to the concept of an ATM subnet that does *not* use inter-building or wide-area links. This level could consist of a peer group and its neighboring peer group, for example. The default value is 96.

### 3) LocalNetworkPlusTwo

This level corresponds to the concept of an ATM subnet that does *not* use inter-building or wide-area links. This level could consist of two or more peer groups, for example. The default value is 96.

### 4) SiteMinusOne

This level corresponds to the concept of an ATM subnet that does *not* use inter-building or wide-area links. This level could consist of the majority of peer groups in a building, for example. The default value is 80.

### 5) IntraSite

This level identifies the inclusive routing hierarchy of nodes that are not geographically separated. You can use this parameter to confine ATM traffic with a local location and thereby avoid using inter-building and wide-area links. The default value is 80.

### **6) SitePlusOne**

This level identifies ATM networks that may use inter-building and wide-area links. This level could consist of ATM networks in two different buildings, for example. The default value is 72.

### **7) OrganizationMinusOne**

This level identifies ATM networks that may use inter-building and wide-area links. This level could consist of ATM networks in several different buildings, for example. The default value is 72.

### **8) IntraOrganization**

This level identifies the inclusive routing of an autonomous organization, which is defined as the organization that has administrative authority of the network. This level may use inter-building and wide-area links. The default value is 64.

### **9) OrganizationPlusOne**

This level identifies the union of at least two autonomous organizations. The default value is 64.

### **10) CommunityMinusOne**

This level identifies the union of two or more autonomous organizations. The default value is 64.

### **11) IntraCommunity**

This level identifies a collection of autonomous organization that are organized by a provider or organizational partnership. The default value is 48.

### **12) CommunityPlusOne**

This level identifies a collection of autonomous organization that are organized by a provider or organizational partnership. The default value is 48.

### **13) Regional**

This level identifies a collection of autonomous organization that are organized by a provider or organizational partnership. The default value is 32.

### **14) InterRegional**

This level identifies a collection of autonomous organization that are organized by a provider or organizational partnership. The default value is 48.

### **15) Global**

This level represents all autonomous organizations That form a connected private ATM network. The default value is 0.



## Viewing the PNNI Map Table

The **pmap** command displays the PNNI map table, which contains information on all links and nodes in the PNNI hierarchy from the perspective of a local node. For the single-peer group version of PNNI, the syntax for this command is as follows:

**pmap [-s]**

For the multiple-peer group version of PNNI, the syntax for this command is as follows:

**pmap [-s] [<node level>]**

The **-s** option provides a summary version of the **pmap** command (see *Summary Form of pmap* on page 45-74). The **<node level>** option allows you to display the current configuration for a specific node. If you do not use this option, then the **pmap** command will display the lowest level node. For example, when you enter

**pmap**

a screen similar to the following displays:

| PNNI Map Table |                        |  |                               |
|----------------|------------------------|--|-------------------------------|
| 1)             | <b>Orig Node Id:</b>   | 50a03903488001bc900001010230e00020da0230e000 |                               |
|                | <b>Remote Node Id:</b> | 50a03903488001bc90000101761c900020da761c9000 |                               |
|                | <b>PGID:</b>           | 50 3903488001bc90000101000000                |                               |
|                | <b>Orig Port ID:</b>   | 192  | <b>Remote Port ID:</b> 384    |
|                | <b>Map entry type:</b> | Horizontal Link                              | <b>Derived Aggr Token:</b> 1  |
|                | <b>PTSE-Id:</b>        | 3  | <b>VP Capability:</b> Enabled |

Map table entries are numbered with a node index number in the leftmost column. This index identifies each connection. A single node may have multiple entries for multiple connections.

**Orig Node Id.** The node identifier of the node originating the PTSE. If the **Map entry type** is **Node**, then this value is also set to zero.

**Remote Node Id.** For horizontal links, this value is the node identifier of the node at the other end of the link from the node originating the PTSE. If the link is unknown, then PNNI sets this value to all zeros. If the **Map entry type** is **Node**, then this value is also set to zero.

**PGID.** The peer group ID of the originating node.

**Orig Port ID.** The Port Identifier assigned to this port by its node. The first number is the CSM slot number and the second number (after the slash) is the port number on the module.

**Remote Port Id.** For horizontal links, this value is the Port Identifier of the port at the remote end of this link. If the remote port is unknown, then PNNI sets this value to zero. For Nodes, this value is the Port Identifier of the remote port connected to the originating port. This value is only relevant to the local switch as ATM switches calculate this value differently.

**Map entry type.** The type of PNNI entity described in this entry. The PNNI type will either be a Horizontal Link (**HORIZ LINK**) or a Node (**NODE**).

**Derived Aggr Token.** The aggregation token for this port on the remote node. This variable is configured through the **ppcfg** command. The aggregation token allows links from the same switch to be advertised separately and contain independent topology metrics. If two links contain different aggregation tokens, then they will be viewed as distinct links by the peer group. If two links have the same aggregation token value, then they will be viewed as the same link by the peer group.

**PTSE-Id.** The PTSE Identifier for the PTSE sent by the originating node that contains the information group(s) describing the PNNI entity. Each PNNI entity (node or link) or aspect of a PNNI entity (such as a node's peer group) is completely described by a single PTSE.

**VP Capability.** Indicates whether this port advertises to other nodes that it supports the establishment of Virtual Path Connections (VPCs). **Enabled** means the port advertises that Virtual Paths can be established on this physical link. **Disabled** means that this port will not advertise that it is capable of setting up Virtual Paths.

### Summary Form of pmap

You can also obtain a summary output of the **pmap** command that organizes all map table information into two tables. When you enter

```
pmap -s
```

A screen similar to the following displays:

| PNNI Map Summary                                 |   |  |  |  |  |
|--|---|--|--|--|--|
| Within Peer Group: 50 3903488001bc90000101000000 |   |  |  |  |  |
| Node Index                                       | Node ID   |  |  |  |  |
| 1  | 50a03903488001bc900001016cdb200020da6cdb2000          |  |  |  |  |
| 2  | 50a03903488001bc900001017defc00020da7defc000          |  |  |  |  |
| 3  | 50a03903488001bc900001017f6bc00020da7f6bc000          |  |  |  |  |
| 4  | 50a03903488001bc900001018519000020da85190000 (MYSELF) |  |  |  |  |
| 5  | 50a03903488001bc900001018519b00020da8519b000          |  |  |  |  |
| 6  | 50a03903488001bc900001019238700020da23870000          |  |  |  |  |

| Orig Node Index | Orig Port Id | Link Type         | Rem Port Id | Rem Node Index |             |
|-----------------|--------------|-------------------|-------------|----------------|-------------|
| 1               | (192)        | <== HLINK ==>     | (192)       | 2              |             |
| 1               | (200)        | <== HLINK ==>     | (416)       | 3              |             |
| 1               | (208)        | <== HLINK ==> 4/2 | (200)       | 4              |             |
| 1               | (216)        | <== HLINK ==> 4/4 | (216)       | 4              |             |
| 2               | (200)        | <== HLINK ==> 7/5 | (416)       | 4              |             |
| 3               | (464)        | <== HLINK ==> 8/3 | (464)       | 4              |             |
| 3               | (200)        | <== HLINK ==> 4/3 | (208)       | 5              |             |
| 4               | 7/8 (440)    | <== HLINK ==>     | (200)       | 5              |             |
| 5               | (216)        | == HLINK ==>      | (200)       | 6              | (Unusable)* |


\* - Unusable links may be so noted if their PTSE is being re-originated due to lifetime expiration.

The first table lists all the node indexes in the PNNI map table along with their associated Node IDs. The second table lists attributes of each node index. Descriptions of these variables can be found in the section for the standard **pmap** output. Note that one of the links is considered “unusable” because PNNI cannot forward on links that are not bi-directional.

A network diagram of this display is shown on the following page.

## Network Diagram of pmap Summary Display

As an illustration, the summary **pmap** display on page 45-74 is generated from a network that is configured as follows. The port numbers shown on each node correspond to internal PNNI port numbers.



Network Configuration Used for pmap Display

In the **pmap** UI display, each connection is represented by a pair of port ids: **Orig Port Id** and **Rem Port Id**. Each port id also has a corresponding node index. If you compare the above diagram with the **pmap** summary display, you will find each port-to-port connection in the display.

The one variant connection is the one listed last in the **pmap** display. This link is listed as **Unusable** in the display. The **pmap** display provides a clue as to why this connection is not usable.

Note that under the **Link Type** column, the arrow points in only one direction—from port 216 to port 200. All other connections in the table show the arrow pointing in both directions, which is symbolic of a bi-directional traffic flow. In the unusable link, traffic can flow from Node 5 (port 216) to Node 6 (port 200), but it cannot flow from Node 6 back to Node 5. When a link is not bi-directional, PNNI deems it unusable.

# Viewing the PNNI Nodal Map Table

The **pnmap** command displays information learned by the local node from nodal information PTSEs. For the single-peer group version of PNNI, the syntax for this command is as follows:

```
pnmap [-s]
```

For the multiple-peer group version of PNNI, the syntax for this command is as follows:

```
pnmap [-s] [ <node level>]
```

The **-s** option provides a summary version of the **pnmap** command (see *Summary Form of pnmap* on page 45-78). The **<node level>** option allows you to display the current configuration for a specific node. If you do not use this option, then the **pnmap** command will display the lowest level node. For example, when you enter

```
pnmap
```

a screen similar to the following displays:

| PNNI Nodal Map Table |  |                            |       |
|----------------------|--|----------------------------|-------|
| 1) Node ID:          | 50a03903488001bc900001020000080020da00000800 * |                            |       |
| Peer Group Id:       | 50   | 3903488001bc90000102000000 |       |
| ATM Address:         | 3903488001bc900001020000080020da00000800       |                            |       |
| Restr. Transit:      | False  | Restr Branching:           | False |
| Complex Rep:         | False  | DB Overloaded:             | False |
| Is Peer Grp Leader:  | False  | Leadership Priority:       | 100   |
|                      |  |                            |       |
| Preferred PGL:       | 50a03903488001bc900001020000080020da00000800   |                            |       |
| Parent Node Id:      | 00     |                            |       |
| Parent ATM Addr:     | 00       |                            |       |
| Parent PG Id:        | 00 00000000000000000000000000000000            |                            |       |
| Parent PGL Node ID:  | 00       |                            |       |

◆ Note ◆

The last set of parameters (**Preferred PGL** through **Parent PGL Node ID**) will not display unless you have installed the software for multiple-peer group PNNI.

- Node ID.** The node identifier for the node described in this entry.
- Peer Group Id.** The peer group of the originating node.
- ATM Address.** The ATM End System (ES) address of the originating node.
- Restr. Transit.** Indicates whether this node is restricted from supporting Switched Virtual Circuits (SVCs) transiting this node. **False** means the port can support ATM transit data links. **True** means the port will be restricted from setting up SVCs unless overridden by another PNNI parameter; only SVCs originating and terminating at this node are supported.
- Restr Branching.** Indicates whether the originating node is able to support additional multicast virtual circuit branches. **False** means that the node can support additional multicast branches. **True** means that additional branches are not supported because the maximum number of multicast virtual circuits on all modules in the node has been reached. The maximum number of multicast virtual circuits supported by a CSM-155 module is 8000, and the maximum supported by a CSM-622 module is 16,000.
- Complex Rep.** Indicates whether the originating nodes use complex node representation. **True** indicates complex representation is used. **False** indicates that simple node representation is used.

**DB Overloaded.** Indicates whether the originating node is currently operating in topology database overload state. If the node is in overload state, then you may want to increase PNNI operating limits through the **pgcfg** command (described in *Configuring General PNNI Parameters* on page 45-31).

**Is Peer Grp Leader.** Indicates whether the originating node claims to be peer group leader of its peer group.

**Leadership Priority.** The leadership priority advertised by the originating node.

◆ **Note** ◆

The following fields only apply to nodes in multiple peer groups.

**Preferred PGL.** The ATM address of the node that this local believes should be or should become the preferred Peer Group Leader (PGL). If a PGL has not been chosen, then this field will display all zeroes.

**Parent Node Id.** When the local node is the Peer Group Leader (PGL), then this field will display the node ID of the parent Logical Group Node (LGN). If the local node is not the PGL of its peer group, then this field will display all zeroes.

**Parent ATM Addr.** When the local node is the Peer Group Leader (PGL), then this field will display the ATM address of the parent Logical Group Node (LGN). If the local node is not the PGL of its peer group, then this field will display all zeroes.

**Parent PG Id.** When the local node is the Peer Group Leader (PGL), then this field will display the local node's parent peer group ID. If the local node is not the PGL of its peer group, then this field will display all zeroes.

**Parent PGL Node ID.** When the local node is the Peer Group Leader (PGL), then this field will display the ID of the node elected as PGL in the parent peer group. If the local node is not the PGL of its peer group, then this field will display all zeroes.

### Summary Form of **pnmap**

The **pnmap** command also has a summary option that allows you to view less information on each node. Simply enter the command

**pnmap -s**

to obtain a display similar to the following:

#### **PNNI Nodal Map Summary Table**

- 1) Node ID: 50a03903488001bc90000101784e500020da784e5000
- 2) Node ID: 50a03903488001bc9000010178aee00020da78aee000
- 3) Node ID: 50a03903488001bc900001017a1bd00020da7a1bd000 (MYSELF)
- 4) Node ID: 50a03903488001bc900001017f6e900020da7f6e9000

This display is a condensed form of the full **pnmap** display. It simply lists the Node IDs for nodes currently included in the PNNI topology database for this network.

## Viewing Current PNNI Calls

The **pcalls** command displays the current in-progress PNNI calls. The number of current calls is listed in the **pginfo** command. The **pcalls** command provides more detail on each of these calls. The table includes current Point-to-Point and Point-to-Multipoint calls. The table for each call type is the same.

### ◆ Note ◆

PNNI calls are opened and closed continuously. Therefore, the calls listed in **pcalls** output are probably complete by the time you view information in this display.

The following is an example of the **pcalls** output. Note that only information for Point-to-Multipoint calls is displayed in this sample. If there were Point-to-Point calls present, then information on those calls would display before the Point-to-Multipoint table.

The Point-to-Point Call Table is empty.

PNNI Point-to-Multipoint Call Table

| CallRef  | Call Id  | Dtl   | CallRef  | Call Id  | Dtl   |
|----------|----------|-------|----------|----------|-------|
| =====    | =====    | ===== | =====    | =====    | ===== |
| 4b200010 | 00000001 | 1     | 4b200010 | 00000002 | 1     |
| 4b201870 | 00000004 | 1     | 4b201c80 | 00000005 | 1     |
| 4b202ab8 | 00000007 | 1     | 4b202cc0 | 00000008 | 1     |
| 4b2030d0 | 0000000a | 1     | 4b2034e0 | 0000000b | 1     |
| 4b203f08 | 0000000c | 1     | 4b204318 | 0000000d | 1     |
| 4b200010 | 0000000e | 1     | 4b200010 | 0000000f | 1     |
| 4b200010 | 00000012 | 1     | 4b209458 | 00000014 | 1     |

**CallRef.** An internal reference to this call used by the PNNI protocol. This number is useful when using a network analyzer to view call activity.

**Call Id.** An internal reference to this call used by signaling software. The value is similar to the **CallRef** variable except it is used by signaling software rather than PNNI to identify this call.

**Dtl.** The Designated Transit List (DTL) or correct path associated with this call.

## Viewing Current DTLs

The **pdtl** command allows you to view a list of all current Designated Transit Lists (DTLs) in the PNNI database. The list will be limited to 200 DTL entries. The following is a sample of the output shown from the **pdtl** command.

**PNNI Designated Transit List Table**

(Note that RefCount 0 DTLs are free DTL entries.)

| DTL Index | Ref Count | Hop Count | DTL (Node ID + Port ID)                           |
|-----------|-----------|-----------|---|
| =====     | =====     | =====     | =====   |
| 1         | 0         | 0         | <No DTL>  |
| 2         | 4         | 3         | 50a03903488001bc9000010178aee00020da:78aee000 208 |
|           |           |           | 50a03903488001bc900001010230e00020da:0230e000 209 |
|           |           |           | 50a03903488001bc900001016542fg0020da:6542fg00 210 |
| 3         | 5         | 3         | 50a03903488001bc9000010178aee00020da:78aee000 208 |
|           |           |           | 50a03903488001bc900001010230e00020da:0230e000 209 |
|           |           |           | 50a03903488001bc900001010967df0020da:0967df00 201 |
| 4         | 1         | 2         | 50a03903488001bc9000010178aee00020da:78aee000 208 |
|           |           |           | 50a03903488001bc900001016542fg0020da:6542fg00 210 |
| 6         |           | 2         | 50a03903488001bc900001010230e00020da:0230e000 209 |
|           |           |           | 50a03903488001bc900001016542fg0020da:6542fg00 210 |
| 7         | 10        | 2         | 50a03903488001bc9000010178aee00020da:78aee000 208 |
|           |           |           | 50a03903488001bc900001016542fg0020da:6542fg00 210 |
| 8         | 20        | 2         | 50a03903488001bc9000010178aee00020da:78aee000 208 |
|           |           |           | 50a03903488001bc9000010178yu650020da:78yu6500 249 |
| 9         | 15        | 4         | 50a03903488001bc9000010178aee00020da:78aee000 208 |
|           |           |           | 50a03903488001bc900001010230e00020da:0230e000 209 |
|           |           |           | 50a03903488001bc9000010178yu650020da:78yu6500 249 |
|           |           |           | 50a03903488001bc90000101135jh40020da:135jh400 235 |
| 10        | 6         | 3         | 50a03903488001bc9000010178aee00020da:78aee000 208 |
|           |           |           | 50a03903488001bc9000010178yu650020da:78yu6500 249 |
|           |           |           | 50a03903488001bc90000101135jh40020da:135jh400 235 |

**DTL Index.** The internal reference number used by PNNI to identify this Designated Transit List.

**Ref Count.** The number of calls that have used this DTL to as a path to reach their destination.

**Hop Count.** The number of hops, or PNNI nodes, on this DTL. This value will be zero (0) until the DTL is actually used by a call. Therefore, if the **Ref Count** column displays a positive number, then the number of hops on this DTL should be listed in this column.

**DTL (Node ID + Port ID).** A description of the hops, or nodes, that comprise this DTL. Each DTL is described as a number of node and port IDs. Each row describes one node/port in the DTL. The number of node/port pairs should correspond to the number of hops listed in the **Hop Count** column.

### Summary Form of pdtl

You can also obtain a summary output of the **pdtl** command that displays the number of DTLs and the number of DTLs that are currently being used. When you enter

```
pdtl -s
```

A screen similar to the following displays:

**From a total of 200 DTLs, 5 are currently used.**



## Viewing Basic Port Statistics

The **pgstats** command displays statistics for the number of Hello, PTSE, and Database Summary packets sent and received on each CSM port in the OmniSwitch. For the multiple-peer group version of PNNI, enter

**pgstats**

at the system prompt. For the multiple-peer group version of PNNI, the syntax for this command is as follows:

**pgstats [ <node level>]**

The **<node level>** option allows you to display the current configuration for a specific node. If you do not use this option, then the **pgstats** command will display the lowest level node. For example, when you enter

**pgstats**

a screen similar to the following displays:

### PNNI Port Basic Statistical Information

| Neighbor                                     | Intf | Hellos<br>Xmtd/Rcvd | PTSPs<br>Xmtd/Rcvd | Dbase Sum Pdus<br>Xmtd/Rcvd |
|--|------|---------------------|--------------------|-----------------------------|
| 3903488001bc90000101<br>75ee100020da75ee1000 | 5/ 2 | 7<br>6              | 7<br>4             | 2<br>3                      |
| 3903488001bc90000101<br>72b1b00020da72b1b000 | 7/ 1 | 7<br>6              | 3<br>1             | 5<br>3                      |

**Neighbor.** The node identifier of the neighboring peer node.

**Intf.** The CSM slot and port for which these statistics are compiled. The slot is listed first, followed by a slash (/), and then the port number.

**Hello Xmtd/Rcvd.** The number of Hello packets transmitted (top value) or received (bottom value) over his link. For links other than those between lowest-level nodes in this peer group, this value will be zero.

**PTSPs Xmtd/Rcvd.** The number of PTSPs transmitted/retransmitted to (top value) or received from (bottom value) the neighboring peer node. Topology database information in the form of PNNI Topology State Elements (PTSEs) is encoded in PTSPs.

**DBase Sum Pdus Xmtd/Rcvd.** The number of Database Summary packets transmitted/retransmitted to (top value) or received from (bottom value) the neighboring peer. Database Summary packets contain identifiers of the PTSE data available from the neighboring node's topology database, but these summary packets do not contain the actual PTSE data. To obtain PTSE data, a node issues a PTSE Request packet after receiving a Database Summary packet.

## Viewing Port Error Statistics

The **pestats** command displays statistics for the number of errors (i.e., PDU errors and cell discards) sent and received on each CSM port in the OmniSwitch. For the single-peer group version of PNNI, enter

```
pestats
```

at the system prompt. For the multiple-peer group version of PNNI, the syntax for this command is as follows:

```
pestats [ <node level>]
```

The **<node level>** option allows you to display the current configuration for a specific node. If you do not use this option, then the **pestats** command will display the lowest level node. For example, when you enter

```
pestats
```

a screen similar to the following displays:

| PNNI Port Error Statistical Information      |      |          |          |          |          |
|--|------|----------|----------|----------|----------|
| Neighbor                                     | Intf | Errors   |          | Discards |          |
|  |      | Incoming | Outgoing | Incoming | Outgoing |
| 3903488001bc90000101<br>75ee100020da75ee1000 | 5/ 2 | 0        | 0        | 0        | 0        |
| 3903488001bc90000101<br>72b1b00020da72b1b000 | 7/ 1 | 0        | 0        | 0        | 0        |

**Neighbor.** The ATM address for the node on which error statistics are provided.

**Intf.** The CSM Slot and Port for which error statistics are provided. The slot is listed first, followed by a slash (/), and then the port number.

**Errors Incoming.** The number of PNNI Protocol Data Unit (PDU) errors that have been received on this port. This figure includes only malformed frames from remote destinations (i.e., outside this peer group). These errors will not occur between neighboring nodes in the same peer group; errors between such nodes will be resolved by the Hello protocol. For example, if there are errors in Hello messages between two nodes in the same peer group, then the link with the neighboring node will simply not be set up and no PDUs will be exchanged.

**Errors Outgoing.** The number of PNNI Protocol Data Unit (PDU) errors that have been transmitted on this port. These errors occur when the OmniSwitch software driver fails to transmit a frame successfully.

**Discards Incoming.** The number of PNNI Protocol Data Units (PDUs) that have been discarded on the receive side of this port. Received frames are discarded if they are corrupt, received on a user-disabled CSM port, or received on a port where the Hello state has gone Down.

**Discards Outgoing.** The number of PNNI Protocol Data Units (PDUs) that have been discarded on the transmit side of this port. Transmit frames are discarded when the OmniSwitch software driver fails to transmit a frame successfully.

## Viewing Port PTSE Statistics

The **ppstats** command displays statistics on the number of PTSE requests and acknowledgments sent and received on each CSM port in the OmniSwitch. For the single-peer group version of PNNI, enter

**ppstats**

at the system prompt. For the multiple-peer group version of PNNI, the syntax for this command is as follows:

**ppstats [ <node level>]**

The **<node level>** option allows you to display the current configuration for a specific node. If you do not use this option, then the **ppstats** command will display the lowest level node. For example, when you enter

**ppstats**

a screen similar to the following displays:

| PNNI Port PTSE Statistical Information       |      |              |      |                       |      |
|--|------|--------------|------|-----------------------|------|
| Neighbor                                     | Intf | PTSE Request |      | PTSE Acknowledgements |      |
|  |      | Xmtd         | Rcvd | Xmtd                  | Rcvd |
| 3903488001bc90000101<br>75ee100020da75ee1000 | 5/ 2 | 1            | 0    | 2                     | 6    |
| 3903488001bc90000101<br>72b1b00020da72b1b000 | 7/ 1 | 0            | 1    | 1                     | 2    |

**Neighbor.** The node identifier for the neighboring peer node.

**Intf.** The CSM slot and port for which these statistics are compiled. The slot is listed first, followed by a slash (/), and then the port number.

**PTSE Request Xmtd.** The number of PTSE Request frames transmitted via this CSM port to neighboring peer nodes. PTSE Requests are sent by this node in response to Database Summary packets from neighboring nodes. When this node finds a PTSE for which it requires more information, it sends a PTSE Request to obtain the actual PTSE data referred to in the Database Summary packet.

**PTSE Request Rcvd.** The number of PTSE Request frames received on this CSM port from neighboring peer nodes. PTSE Requests are sent by neighboring nodes in response to Database Summary packets from this node. When a neighboring node finds a PTSE for which it requires more information it sends a PTSE Request to obtain the actual PTSE data referred to in the Database Summary packet.

**PTSE Acknowledgement Xmtd.** The number of PTSE Acknowledgment frames transmitted via this CSM port to neighboring peer nodes. PTSE Acknowledgments are sent in response to received PTSEs. When this node receives a PTSE it requested through a PTSE Request packet, it acknowledges this receipt via a PTSE Acknowledgment packet.

**PTSE Acknowledgement Rcvd.** The number of PTSE Acknowledgment frames received on this CSM port from neighboring peer nodes. PTSE Acknowledgments are sent in response to PTSE Requests. When a node receives a PTSE it requested through a PTSE Request packet, it acknowledges this receipt via a PTSE Acknowledgment packet.

# Halting PNNI Operations

The **phalt** command disables the PNNI protocol entity currently running in this OmniSwitch. This command clears out the topology database and all port and node configuration data on this switch. Using **phalt** could result in a loss of network connectivity for this node and could affect databases of PNNI nodes in the same peer group. (The **phalt** command, however, does not unload the PNNI module from system memory.) This command should only be used during off-peak network periods. The **phalt** command must be used before using the **prestart** command (described in *Restarting PNNI* on page 45-85).

For the single-peer group version of PNNI, enter

**phalt**

at the system prompt. For the multiple-peer group version of PNNI, the syntax for this command is as follows:

**phalt [ <node level>]**

The **<node level>** option allows you to halt PNNI operation at the level indicated. If you do not use this option, then the **phalt** command will halt all PNNI operation in an OmniSwitch at every level. For example, to halt the PNNI protocol at every level, enter

**phalt**

The following warning message and confirmation prompt will display:

**Shutting down the PNNI Operation currently executing in this node could cause temporary ATM network routing problems such as network partitioning. Which, in turn, could result in a loss of connectivity. Also, remote ATM PNNI nodes may experience temporal disinformation as they expire this node's PTSEs and neighbor information.**

**This typically will not result in a loss of ongoing connections through this switch. It will prevent new connections from temporarily being routed through this node however. Also, ongoing connections which crankback will be discarded, causing them to timeout at the originator.**

**Are you absolutely sure you wish to do this? (n) y**

If you want to continue with the PNNI halt, enter a **Y**. Otherwise, press **<Enter>** and the command will be exited. If you enter **Y**, the following status messages will display as the PNNI entity is brought down:

|   |    |
|---|----|
| Disabling all available PNNI ports...         | OK |
| Removing all end systems directly attached... | OK |
| Removing all statically configured routes...  | OK |
| Unbinding with Upper Call Control...          | OK |
| Unconfiguring this PNNI Node's basic info...  | OK |
| Freeing buffers...                            | OK |
| Unconfiguring Dynamic Routing Tables...       | OK |
| Generating Final Mgmt Event...                | OK |
| Done.   |    |

## Restarting PNNI

The PNNI protocol entity on this OmniSwitch can be restarted after you halt it via the **phalt** command. The **prestart** command brings the PNNI protocol back up after a halt. Database and configuration information will be set to system defaults upon this restart.

For the single-peer group version of PNNI, enter

**prestart**

at the system prompt. For the multiple-peer group version of PNNI, the syntax for this command is as follows:

**prestart [ <node level>]**

The **<node level>** option allows you to restart PNNI operation at the level indicated. If you do not use this option, then the **prestart** command will restart all PNNI operation in an OmniSwitch at every level. For example, to restart the PNNI protocol entity, enter:

**prestart**

Status messages will display, updating you on the current status of the restart:

|  |    |
|--|----|
| Generating Mgmt Event...                       | OK |
| Configuring Dynamic Routing Tables...          | OK |
| Configuring this PNNI Node's basic info...     | OK |
| Configuring Call Control service...            | OK |
| Inserting all end systems directly attached... | OK |
| Inserting all statically configured routes...  | OK |
| Enabling all available PNNI ports...           | OK |
| Done.  |    |

Note that you cannot restart PNNI before disabling it through the **phalt** command. If you do not use the **phalt** command before **prestart**, the following message will display:

**You must first shut down pnni using phalt**

# Resetting PNNI Statistics Counters

The **preset** command resets PNNI port statistics counters for this OmniSwitch or for selected PNNI ports. It affects the results of commands on the **Pstats** submenu (**pgstats**, **pestats**, and **ppstats**).

For the single-peer group version of PNNI, the syntax for this command is as follows:

```
preset all | <slot/port>[/<instance>] | <slot/port-list>[/<instance>]
```

The **[/<instance>]** option allows you to specify a specific virtual path tunnel. To reset statistics counters for all PNNI ports in this OmniSwitch, for example, enter the following:

```
preset all
```

the following message will display:

```
PNNI statistics for all interface's have been cleared.
```

To reset statistics for a specific CSM port, enter the **preset** command followed by the slot and port number for the port. For example to reset counters for port 1 on the CSM module in slot 3, enter:

```
preset 3/1
```

You can also, enter multiple port numbers to reset several ports on a single CSM module. For example, to reset counters for ports 1 through 4 on the CSM module in slot 3, you would enter:

```
preset 3/1-4
```

The following message would display in response to this command:

```
3/1-4 cleared.
```

In addition, you can also reset the PNNI port statistics counters for a specific virtual path tunnel on a single port or on a range of ports on a CSM module. For example, to reset the PNNI port statistics counter for virtual path tunnel 1 on CSM port 3 in slot 3, enter

```
preset 3/3/1
```

at the prompt.

## Viewing PNNI Configuration Information

The **pvcfg** command displays the type of PNNI information that has been user-configured. This information may have been configured through UI commands or through SNMP-based network management software. You can configure general PNNI parameters, node-specific parameters, and PNNI port-specific parameters. In addition, you can configure IISP route and route property information.

You can see the types of PNNI information that have been configured by entering **pvcfg**. A screen similar to the following displays:

```
pgcfg info
pncfg info
ppcfg 4/4 info
Route property info - property 1
Route instance 0 info for property 1
```

This display shows that general PNNI parameters (**pgcfg info**), node-specific parameters (**pncfg info**), and PNNI port parameters for port 4 on the CSM module in slot 4 (**ppcfg 4/4 info**) has been configured. In addition, a route property template has been set up (**Route property info - property 1**) and one route has been configured using the route property template (**Route instance 0 info for property 1**).

## Removing PNNI Configuration Information

When you configure PNNI through UI commands or via SNMP-based network management software, configuration data is stored in the **mpm.cnf** file. The PNNI-specific information in this file may be removed by using the **prmcfg** command. Using **prmcfg** does affect any other non-PNNI information in the **mpm.cnf** file.

The syntax for this command is as follows:

```
prmcfg [-i]
```

The **-i** option will cause the **prmcfg** command to run “interactively.” In other words, you will be prompted before any configuration information is removed.

To remove PNNI configuration information (without prompts), enter

```
prmcfg
```

at a system prompt. PNNI unconfigures any static PNNI information. Run-time PNNI configuration information is not affected. However, when you reset the switch, the factory default configurations for all PNNI parameters will be restored. If you set a port to the PNNI type through the **map** command, then that port will remain a PNNI port after using **prmcfg** and rebooting.

# Verifying Routes

Since each PNNI node learns the topology for the whole network, each node has the ability to reach all known destinations. The **prtst** command helps to verify PNNI reachability to a destination. When you issue the command, the system displays how PNNI will attempt to route a call through the network to a destination ATM address you specify.

To verify a route, enter **prtst** followed by the ATM address for which you want to check the route. For example, to check the route to ATM address **4222334455667788990011223344556677889900**, you would enter the following:

```
prtst 4222334455667788990011223344556677889900
```

The system displays a screen showing the Node ID and ports used to reach the address.

**PNNI created the following Designated Transit List to setup this connection:**  
**(Note that the first node in the DTL is displayed first, while the last node is displayed last. For reference, if the downstream switch is an OmniSwitch, its OC3 slot/port is displayed after the logical port id conveyed in the DTL)**

| Logical Node ID   | Logical Port | XCell Slot/Port |
|---|--------------|-----------------|
| 60a0 3903488001bc900001178aee010020da:78aee000          | 208          | (4/3)           |
| current->60a0 3903488001bc900001178db3010020da:78db3000 | 232          | (4/6)           |

The “current “node is the local OmniSwitch and the other node is the one where the specified ATM address is attached. If the ATM address is not known or if it is attached to the local switch, then a message will be returned informing of the unknown or local address.



## Operating PNNI with Redundant MPMs

If you have an OmniSwitch with redundant MPMs, PNNI node configuration parameters can be lost during failover. For example, if you do not specify the switch's node ID, it will change if the primary MPM goes down and the secondary MPM becomes primary. To prevent the switch's PNNI node parameters from changing after a failover, perform the steps described in the subsection below.

As an option, you can verify that the node information is consistent on both MPMs by performing the steps described in *Verifying PNNI Node Information on Redundant MPMs* on page 45-90.

### Configuring Node Information on Redundant MPMs

Perform the steps below to configure the same node information on both MPMs.

1. Enter

**pncfg**

at the system prompt. (See *Configuring Node-Specific Parameters* on page 45-39 for more information on the **pncfg** command.)

2. Enter **1=** followed by the ATM address of the switch (to change it from **Unspecified**). For example, if the ATM address is **3903488001bc9000010178aee00020da78aee000**, enter

**1=3903488001bc9000010178aee00020da78aee000**

at the prompt.

3. Enter **2=** followed by the node ID of the switch (to change it from **Unspecified**). For example, if the node ID is **3903488001bc9000010178aee00020da78aee000**, enter

**2=3903488001bc9000010178aee00020da78aee000**

at the prompt.

4. Enter **3=** followed by the node level of the switch (to change it from **Unspecified**). For example, if the node ID is **80** (decimal), enter

**3=80**

at the prompt.

5. Enter **6=** followed by the summary address for ILMI clients (to change it from **Unspecified**). For example, if the summary address is **3903488001bc90000178aee0**, enter

**6=3903488001bc90000178aee0**

at the prompt.

6. If you have the single-peer group version of the software, proceed to step 7. If you have the multiple-peer group version of the software, enter **7=** followed by the Peer Group Leader (PGL) priority of the node (to change it from **Unspecified**). For example, if the PGL priority is **50** (decimal), enter

**7=50**

at the prompt.

7. Enter

**save**

at the prompt to save your settings. The following prompt will display.

**Do you want these changes to take effect immediately? (y)**

8. Press **<Enter>** to have your changes take effect immediately.

**save**

at the prompt to save your settings.

9. Enter

**configsync**

at the system prompt to synchronize the configuration file on both MPMs. (See Chapter 10, “Configuring Management Processor Modules,” for more information on the **configsync** command.) Messages similar to the following will be displayed.

**Syncing Config file**  
**Config files are currently synchronized.**

## Verifying PNNI Node Information on Redundant MPMs

To verify that the PNNI node configurations are the same on both MPMs, perform the following steps.

1. Enter

**pncfg**

at the system prompt. (See *Configuring Node-Specific Parameters* on page 45-39 for more information on the **pncfg** command.)

2. Copy the ATM address, node ID, node level, and summary address for ILMI displayed in fields 1, 2, 3, and 6, respectively. If you have the multiple-peer group version of the software, copy the PGL priority displayed in field 7.

3. Enter

**quit**

at the prompt to exit the command.

4. Enter

**renounce**

at the prompt to make the secondary MPM the primary one. (See Chapter 10, “Configuring Management Processor Modules,” for more information on the **renounce** command.)

5. Log into the new primary MPM.

6. Enter

**pncfg**

at the system prompt.

7. Compare the data displayed in fields 1, 2, 3, and 6 (and field 7 if you have the multiple-peer group version of PNNI) with the data you copied in Step 2. If the data are not consistent, perform the steps described in *Configuring Node Information on Redundant MPMs* on page 45-89.

## FCSM I PNNI Frame Size Guidelines


PNNI packet (frame) sizes vary from 16 to 8192 bytes, depending on what type of information is being sent between PNNI nodes. For example, hello packets are approximately 100 bytes long. However, database summaries and PNNI Topology State Elements (PTSEs) exchanged between switches can run from 64 to 8192 bytes long.

As frame sizes increase, the total number of CSM ports you can have decreases on OmniSwitches with an FCSM I. In all-Alcatel network, this is not a problem since Alcatel switches do not generate frames greater than 2048 bytes. However, in multi-vendor networks, frames can be as large as 8192 bytes.

◆ **Note** ◆

This restriction only applies to OmniSwitches with an FCSM I; it does not apply to the FCSM II or MPM-C.

The figure below shows the relationship between PNNI frame size and the maximum number of CSM ports you can have an OmniSwitch with an FCSM I. For example, if you need to support PNNI frames up to 8192 bytes, the total number of CSM ports that you can have in the chassis is limited to 21 CSM ports.



**PNNI Frame Size vs. Maximum Number of CSM Ports**

You can configure the maximum frame size on the FCSM I with the **map** command, which is described in Chapter 35, “Managing ATM Access Modules.” (You configure the frame size on the ASM side of the FCSM I and not the CSM side.) For example, if the FCSM I is in slot 2, you would enter

**map 2/2**

at the system prompt.

When you configure the frame size, you should follow these two guidelines;

- Small-to-medium networks (i.e., up to 20 PNNI nodes ) : 2048 to 6144 bytes
- Large networks (i.e., more than 20 nodes): 8192 bytes

In addition, the maximum number of PNNI neighbors per PNNI node should be less than, or equal to, 30 neighbors. More than 30 neighbors will trigger higher CPU utilization on the MPM, unless the hello timers are re-configured using the **pncfg** command with higher values. (See *Configuring Node-Specific Parameters* on page 45-39 for more information on the **pncfg** command.)

