Address and Closed User Group Planning

Proper address planning can greatly increase the performance of a PNNI WAN. Although a PNNI WAN can support almost any addressing scheme, an uncoordinated address scheme can cause excessive address advertisement and needless rerouting, both of which reduce network performance. A good addressing plan is one which is hierarchical in nature and thus summarizes simply and efficiently.

The PNNI Closed User Group (CUG) feature allows the network administrator to define user groups of ATM addresses. Once these user groups are defined, the administrator can control how users within the groups communicate with other group members and with those outside the group.

This chapter provides an address planning overview, a CUG planning overview, and general guidelines for creating an ATM address plan and a CUG plan.

Note All Cisco MGX and SES switch products ship with default addresses. These defaults are provided for lab evaluations of these products. Before the switch is deployed, Cisco Systems advises you to reconfigure the default addresses using the address plan guidelines in this chapter.

Address Planning Overview

Every route across a PNNI network is determined by two ATM End Station Addresses (AESAs), a source and a destination. When a connection is being established, the source PNNI routing node looks up the destination address in PNNI routing tables. If the routing tables do not contain a satisfactory predefined route, the switch uses the PNNI topology database to search for a route. Routing decisions are made based on many criteria as discussed in Chapter 4, "Planning Intermediate Route Selection." This section focuses on how proper address planning can make PNNI routing more efficient.

Note The source end of a connection is also called the master end of the connection, as the master end is responsible for initiating the connection. The destination end is also called the slave end.

PNNI provides both a routing protocol and a signaling protocol. The routing protocol is used to build a topology database and create a route table of all the reachable AESAs. The signalling protocol is used to establish calls across the PNNI network. When initiating a call, the signaling protocol refers to the routing table or topology database to locate a route to the destination ATM address.

To understand the importance of an address plan, consider how PNNI would respond if there were no plan. Consider a network with 100 non-coordinated destination ATM addresses. Assume that all addresses were chosen at random. To enable access to all destinations, PNNI has to create a separate route for each of the 100 destinations, and this has to be repeated on every switch in the network.

Furthermore, PNNI switches exchange data with all other nodes in the peer group, so lots of address information would be transmitted constantly throughout the network as PNNI monitors the network topology.

Now let's consider a more efficient example. Figure 3-1 shows a PNNI network with some simplified addresses in place of the 20-byte ATM addresses.

Figure 3-1 PNNI Addressing Example

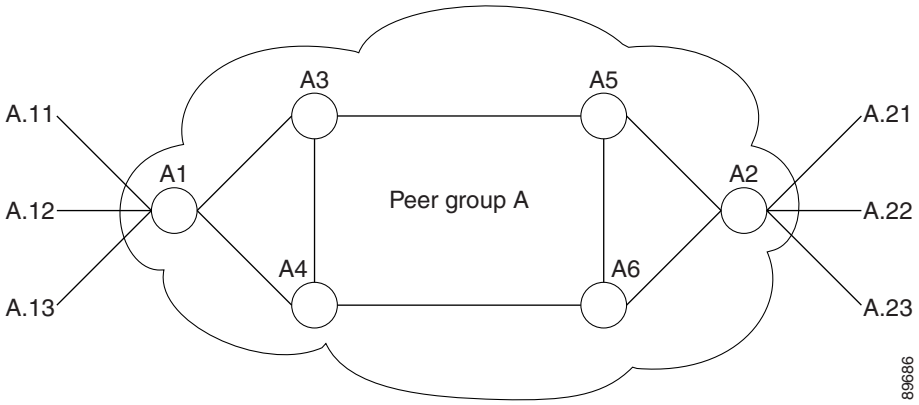

For consistency, assume that the six switches shown in Figure 3-1 connect to a total of 100 destinations. Notice that the destination addresses for the external lines connected to A.1 all use the prefix A.1, and the destination lines connected to A.2 use the prefix A.2. When you configure a common prefix for multiple addresses, you can reduce the size of the routing table and the topology database by storing routes to the address prefix, instead of routes to every destination. In this example, all nodes in Peer Group A store routes to the other switches, but there is no need to store additional routes for every destination address. The use of address prefixes is also called address summarization.

Address summarization also makes network management easier because you do not need to manually enter every AESA into the source nodes. Instead, you define a PNNI address prefix, which summarizes all destinations that share that prefix.

Address summarization does not preclude the use of non-conforming addresses. For example, if network management dictates the use of a specific non-conforming ATM address for a destination, that address can be manually entered at the switch, and PNNI will advertise a route to that device. The non-conforming address is called a foreign address. The support of foreign addresses makes PNNI more flexible, but keep in mind that excessive use of foreign addresses does impact switch performance.

Tip Chapter 4, "Planning Intermediate Route Selection," describes how up to five routes can be stored in a total of 10 route tables for each destination. To understand the impact of foreign addresses, multiply the potential of 50 routes times the number of switches in a peer group, and then multiply that number times the number of foreign addresses. Address summarization is a key component in PNNI address planning.

When a call is placed to a destination address, PNNI refers to the destination addresses and prefixes in the routing tables or topology database. After the best route is chosen to the destination switch, the destination switch selects the appropriate destination interface by searching internal address tables for the longest prefix match. When a switch and its interfaces are configured with prefixes that enable PNNI to quickly locate the destination interface, PNNI routing is most efficient.

Although address summarization does make network management easier and routing more efficient, it can be misused and make PNNI routing less efficient. Consider the case where the same address prefix is assigned to multiple nodes. This is a valid configuration, but it can lead PNNI to unnecessarily reroute

connections as it attempts to locate the correct node. A better design would use the longest possible prefix to represent all the interfaces on a node, and then a longer prefix on each interface that uniquely defines each interface.

At the end of this chapter, there are two worksheets (Table 3-4 and Table 3-5) into which you can enter your WAN address values. If you are familiar with designing PNNI address structures, or if a plan is already completed, you can go directly to the Address Plan Worksheet and enter the values. The procedures for configuring ATM addresses on Cisco MGX and SES switch products are described in the following guides:

- *Cisco MGX 8850 (PXM1E/PXM45), Cisco MGX 8950, Cisco MGX 8830, and Cisco MGX 8880 Configuration Guide, Release 5.1*
- *Cisco SES PNNI Controller Software Configuration Guide, Release 3*

Planning Address Configuration Settings

Use the following steps to create a WAN address plan:

Step 1 Select an ATM address format

Step 2 Select a PNNI level

Step 3 Select the PNNI peer group ID

Step 4 Select the ATM address

Step 5 Select the ILMI address prefix

Step 6 Select the SPVC address prefix

Step 7 Plan address prefixes for AINI and IISP links

Step 8 Select static addresses for UNI ports

These steps are described further in the remainder of this chapter.

Selecting an ATM Address Format

Each PNNI node must be configured for at least one ATM address format. This is an ATM requirement that must be considered when choosing PNNI addresses. To establish ATM connections, each ATM UNI end system must have at least one ATM End System Address (AESA) that uniquely identifies that ATM endpoint. This section explains the supported AESA address formats and their structures.

Caution Each node must support the address format of all its neighboring nodes.

Supported Address Formats

The Cisco MGX and SES switch products support the following standard ATM formats:

- Native E.164
- Data Country Code (DCC)
- International Code Designator (ICD)
- AESA-embedded E.164
- Local AESA

The native E.164 address specifies an Integrated Services Digital Network (ISDN) number and is used by Public Switched Telephone Networks (PSTNs). A native E.164 address has a variable length of up to 15 Binary Coded Decimal (BCD) digits. The other address formats are usually used for private networks. The default address format for the MGX 8850 is the ICD format.

In the PNNI network, native E.164 addresses are mapped to an E.164 AESA format. The native AESA is inserted as a left-justified IDI portion of the AESA, with the semi-octet Hex FFFF padded to form an integral byte at the end. This left-justified rule may be changed to right-justified via CLI if needed.

The substructures of the address formats are transparent to PNNI routing. Figure 3-2 shows the substructures of the supported ATM address formats. Table 3-1 describes the substructures shown in Figure 3-2.

Figure 3-2 Supported ATM Address Formats

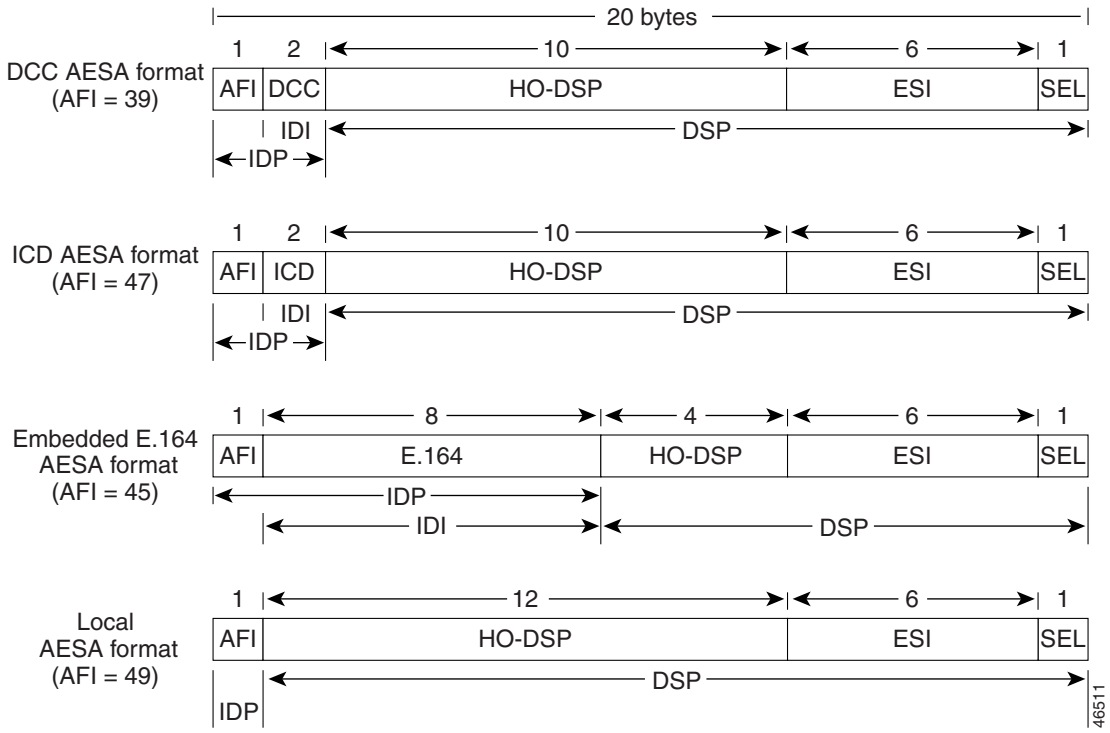

Table 3-1 ATM Address Components

ATM Field	Description	Default Values
AFI	Authority Format Identifier (1 byte).	47
ICD	International Code Designator (2 bytes) The default value is the ICD assigned to Cisco Systems.	0091
IDI	Initial Domain Identifier (8 bytes). The contents of this field vary depending on the value of the AFI. For example, with a DCC AESA (AFI=39), the IDI value of Hex 840F identifies the United States.	—
HO-DSP	High-Order Domain Specific Part (4 to 12 bytes). The meaning is defined by the address authority controlling the AESA. This component couples with AFI and IDI to route a call to the appropriate switch.	—
ESI	End System Identifier (6 bytes). This field repeats the PNNI Controller MAC address when the ATM address identifies the PNNI node. (When an ATM address identifies an ATM end system, the ESI field will be completed through ILMI registration with the end system. In this case, the ESI is typically the MAC address of the ATM CPE. The unique ESI field will distinguish that ATM end system [ATM CPE] from all other ATM end systems.)	PXM45 MAC address at first boot.
SEL	PNNI selector byte (1 byte). The selector byte is used to identify different target applications on the node.	00

Guidelines for Selecting an Address Format

It is important to select a address format plan which meets the future needs and scale of the network. Changing the node ATM address format and addresses after its initial deployment requires major service disruption, and requires complete reconfiguration of the node and all of its connections. Consider the following guidelines when selecting an address format:

- Consider whether a registered address will be required in the future. The default registered address is registered to Cisco and is part of the ATM address.
- If an address format has been chosen for the WAN, or if your WAN will consist of existing nodes for which an address format has been selected, you can select that address format.
- Both Public ATM networks (PSTNs) and Narrowband Integrated Services Digital Networks (N-ISDN) usually use E.164 numbers. PNNI allows end-system reachability to be advertised via the E.164 address prefix.
- In the Data Country Code (DCC) format, each country has a unique DCC value. If you select this address format, your value must match this standard.
- In the International Code Designator (ICD) format, the ICD identifies an organization such as a company or campus. This identification is useful when you are deploying a WAN that will be accessed by several campuses or sites.
- Native E.164 addresses can be embedded in the AESA.

Enter the address format or formats that you select into the Nodal Address Worksheet, Table 3-4, which appears at the end of this chapter.

Note Local AFIs should not be used for addressing within ATM Service Provider networks.

Address Registration Authorities

Table 3-2 lists the address registration authorities.

Table 3-2 Address Registration Authorities

Category	Type	Authorities
ATM Service Providers (ASPs)	ICD	1. US—American National Standards Institute (ANSI). 2. UK—British Standards Institution (BSI). Identifiers for Organizations for Telecommunications Addressing (IOTA) http://www.bsi-global.com/DISC/Working+Withyou/Naming+Addressing.xalter.
	DCC	1. ISO National Administrative Authority (Registration Authority). 2. List of authorities: – US—American National Standards Institute (ANSI). – Germany—Deutsche Industrie-Normen (DIN).
	E.164	International Telecommunications Union (ITU), the National Numbering Authority.
Private networks	ASP Addresses	Private ATM networks can apply to their ATM Service Provider for addresses.
	ICD	Identifiers for Organizations for Telecommunications Addressing (IOTA) http://www.bsi-global.com/DISC/Working+Withyou/Naming+Addressing.xalter.
	DCC	ISO National Administrative Authority (Registration Authority).
	Unregistered addresses	Private networks may create unregistered addresses. Note that such addresses are not globally unique. It is recommended that unregistered addresses be formed using the local AFI (49).

Selecting a PNNI Level

PNNI uses a hierarchical address scheme to define the physical topology and to create a logical hierarchy above the physical topology. Figure 3-3 shows an example of a physical topology.

Figure 3-3 PNNI Network Physical Topology

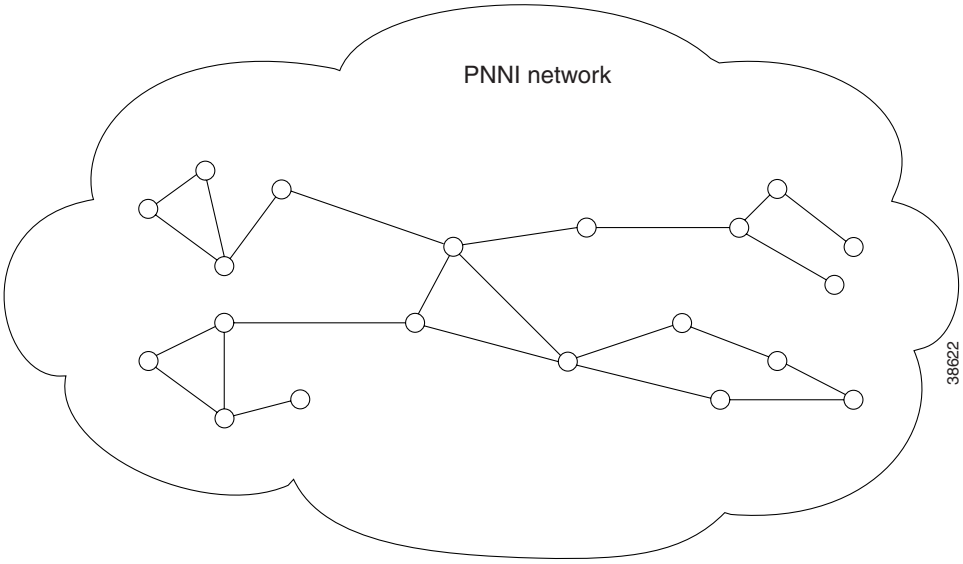

The topology shown in Figure 3-3 becomes a Single Peer Group (SPG) PNNI WAN if no hierarchy is applied. In an SPG WAN, every node stores information about every other node and the CPE that connect to it. To distribute information about all the nodes in the WAN, the PNNI switches send PNNI Topology State Element (PTSE) messages to each other on a regular basis. In a small WAN, an SPG application is appropriate. When the WAN grows beyond 100 nodes, PTSE distribution and the size of the node PNNI databases begins to affect network performance. At this point, you might want to consider creating a Multiple Peer Group (MPG) WAN.

Figure 3-4 shows an example topology of a PNNI MPG WAN.

Figure 3-4 MPG WAN Topology

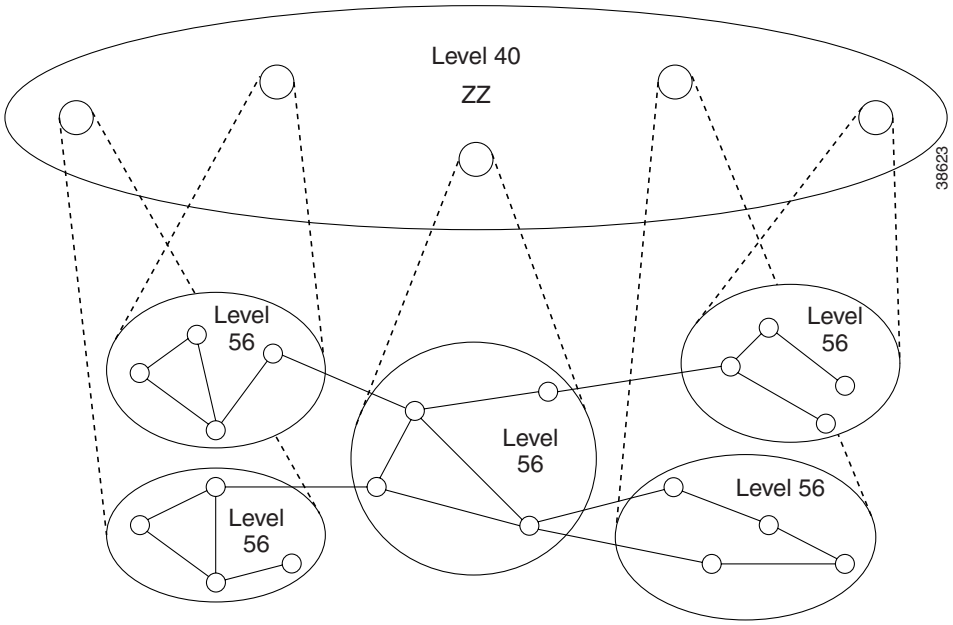

The network shown in Figure 3-4 uses the same physical topology as that shown in Figure 3-3 for an SPG WAN. The difference is that the physical network has been divided into five peer groups at level 56. The level will be explained later in this section. What is important to understand now is that the physical topology is still the same as when all nodes were in a single peer group. Dividing the physical WAN into multiple peer groups simply reduces the size of each peer group, which reduces the total number of PTSEs and the size of the PNNI database within each node. This improves PNNI network performance within each of the smaller peer groups, which leaves more bandwidth and node resources available for processing calls.

The level 40 peer group shown in Figure 3-4 is a logical peer group that has been defined to enable communications between the peer groups at the lower levels. Each of the level 56 peer groups operate more efficiently because they do not have to keep up with changes in the other level 56 peer groups. However, because the level 56 peer groups do not know the details about the other level 56 peer groups, they cannot communicate with the other groups without help from a higher level process.

The level 40 peer group shown in Figure 3-4 is created by adding a higher-level PNNI processes to one of the nodes in each level 56 peer group. Each higher level process operates as a logical group node (LGN) at this higher level, and together these nodes form a logical PNNI peer group at this level. The level 40 peer group nodes exchange PTSEs regarding the level 56 peer groups and maintain a database with routing information for communicating between the lower-level peer groups. Level 40 nodes do not store the routing details stored within the level 56 peer groups, because that information is already stored at the lower level. The level 40 nodes only store the information that the level 56 nodes need to locate and communicate with other peer groups.

If the network shown in Figure 3-4 were to grow until there were more than 100 LGNs at level 40, the level 40 peer group could be divided into multiple peer groups and a higher level could be created to enable communication between the level 40 peer groups. This process can continue until the practical maximum of 10 levels is reached. When you consider that 100 level 40 peer groups equate to approximately 10,000 level 56 nodes (100 level 40 nodes times 100 level 56 nodes), it is easy to see how adding additional layers enables PNNI to scale.

Note These calculations are based on general guidelines. Peer groups on MGX and SES nodes can support up to 160 nodes. Also, remember that these calculations are for network nodes, not CPE. The actual number of CPE and calls supported is considerably higher.

In general, when you create an SPG or MPG network, you need to select a starting level for your PNNI network, which should be the lowest level you will ever need. You can always add higher levels to an SPG or MPG network, but creating lower levels requires a significant amount of reconfiguration.

The PNNI level is mathematically related to the ATM addresses used in a PNNI network. Valid levels are 1 through 104. These numbers specify the number of ATM address bits that are used for the peer group ID, which is described in the next section. Specifically, the level identifies the number of sequential most-significant ATM address bits that define the peer group ID.

Although the PNNI specifications provide for up to 104 PNNI levels, they also limit the practical application to 10 levels. Some PNNI experts suggest that four levels will be sufficient for most PNNI networks. For these reasons, and because it easier to translate bytes of an ATM address instead of bits, Table 3-3 shows the recommended levels to use for PNNI networks.

Table 3-3 Recommended PNNI Level Values

Level	Peer Group ID Portion of ATM Address	Peer Group ID Length (Bytes)
8	11 xx xx xx xx xx xx xx xx xx xx xx xx xx xx xx xx xx xx	1
16	11 22 xx xx xx xx xx xx xx xx xx xx xx xx xx xx xx xx xx	2
24	11 22 33 xx xx xx xx xx xx xx xx xx xx xx xx xx xx xx xx	3
32	11 22 33 44 xx xx xx xx xx xx xx xx xx xx xx xx xx xx xx	4
40	11 22 33 44 55 xx xx xx xx xx xx xx xx xx xx xx xx xx xx	5
48	11 22 33 44 55 66 xx xx xx xx xx xx xx xx xx xx xx xx xx	6
56	11 22 33 44 55 66 77 xx xx xx xx xx xx xx xx xx xx xx xx	7
64	11 22 33 44 55 66 77 88 xx xx xx xx xx xx xx xx xx xx xx	8
72	11 22 33 44 55 66 77 88 99 xx xx xx xx xx xx xx xx xx xx	9
80	11 22 33 44 55 66 77 88 99 AA xx xx xx xx xx xx xx xx xx	10
88	11 22 33 44 55 66 77 88 99 AA BB xx xx xx xx xx xx xx xx	11
96	11 22 33 44 55 66 77 88 99 AA BB CC xx xx xx xx xx xx xx	12
104	11 22 33 44 55 66 77 88 99 AA BB CC DD xx xx xx xx xx xx	13

The default PNNI level for MGX 8850 is 56, which is the midpoint of the recommended values. If this is the lowest level that you expect to need, you can accept the default. If you anticipate needing lower levels in the future, you should select the lowest level that you think you will need now, and enter the level number in the Nodal Address Worksheet, Table 3-4, which appears at the end of this chapter. If you are planning to create higher PNNI levels, you can also note these in the worksheet.

Note If your ATM network connects to a public network, or if you want to conform to public network ATM address rules for future expansion, you cannot create more than one peer group at PNNI levels 1 through 8. This is because the first 8 bits (byte 1) are reserved for the AFI and must be set to a fixed value. Also, if the address format you choose is DCC AESA or ICD AESA, you can create only one peer group for each DCC or ICD.

Selecting the PNNI Peer Group ID

As described in the previous section, the PNNI level selects the number of ATM address bits that are unique within the peer group ID. After you select a PNNI level for a peer group, you need to define the peer group ID using the PNNI level, the number of address bits defined by the PNNI level, and trailing zeros. Figure 3-5 shows the format of the default peer group ID for the MGX 8850.

Figure 3-5 Default Peer Group ID

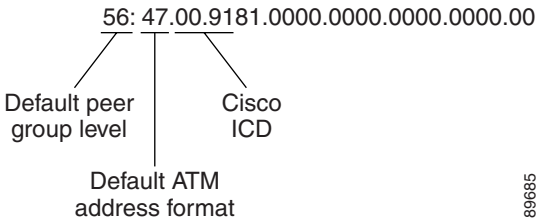

As Figure 3-5 shows, the peer group ID begins with the PNNI level, followed by a colon. The unique portion of the peer group ID follows next. The unique portion of the ID, which is the first 7 bytes by default, corresponds with the left-most or most-significant bytes of the ATM address. The Cisco default PNNI level is 56, so the first 7 bytes of the default ATM address make up the unique portion of the peer group ID: 47.009181000000.

The total length of a peer group ID is 14 bytes, so the bytes that follow the unique portion of the peer group ID are all set to 0. Therefore, the complete default peer group ID for all MGX 8850 is: 56:47.00.9181.0000.0000.0000.0000.00. The periods within the peer group ID are used to make it easier to read the peer group ID. To create a second peer group at the same default level, you must modify the unique portion of the peer group ID. For example: 56:47.00.9181.0000.01.

Note Only the unique portion of the peer group ID, which is defined by the PNNI level, is used to identify the peer group. In the example of the default level 56, the first 7 bytes of the ATM address define the peer group ID. Although up to 13 bytes can be used for the peer group ID, all bytes beyond what is specified by the PNNI level are ignored with respect to the peer group ID. Although the nonunique bits in the first 13 bytes appear as zeros in the peer group ID display, they do not have to be set to 0 for ATM addresses.

The peer group ID is used to identify ATM addresses that are part of the same PNNI peer group. For example, the following PNNI addresses are all in the same default PNNI peer group:

- 47.009181000000112233445566.778899101112.01
- 47.009181000000112233445566.778899101113.01
- 47.009181000000112233445566.778899101114.01
- 47.009181000000778899101112.112233445566.01

The above addresses are all in the same peer group because the PNNI level for all addresses is the default level (56 bits or 7 bytes) and the first 7 bytes of all these addresses are the same.

When planning peer group IDs for your WAN, consider the following:

- All peer group IDs within a peer group must be identical.
- Each peer group must have its own ID that is unique within the WAN.

- If you change the address format, you need to change the peer group ID.
- If you change any of the identifiers within the unique portion of the peer group ID (for example, the ICD), you must change the peer group ID.

Enter the peer group ID into the Nodal Address Worksheet, Table 3-4, which appears at the end of this chapter.

Selecting the ATM Address

The node ATM address must be unique on the WAN and conform to the selections you have made for the following parameters:

- Address format
- Peer group ID

Figure 3-6 shows the default ATM address for the MGX 8850 switch.

Figure 3-6 20-byte Node Address

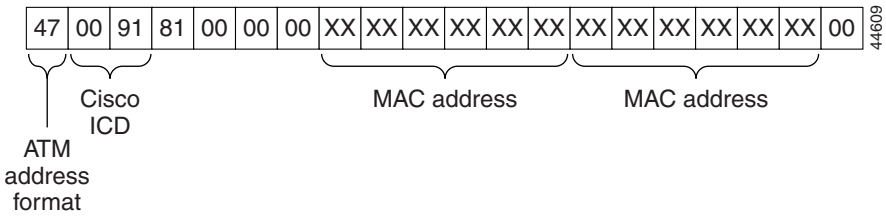

The first byte (47) of the default address identifies the address as an International Code Designator (ICD) ATM End Station Address (AESA). The second and third bytes (0091) define the globally unique ICD assigned to Cisco, and the next four bytes (81000000) are identical for all MGX 8850. The unique portion of the default node address is the 6-byte MAC address, which is used in bytes 8 through 13 and again in bytes 14 through 19. Byte 20, which is the selector byte, is set to 00 by default.

Note Cisco recommends that you change the Cisco ICD portion of the address (0091). This number is registered to Cisco and using it will cause conflicts if the network you create is ever connected to a network to which Cisco connects, or to a network that is using Cisco equipment with the default parameters.

You do not have to change the default ATM address for MGX 8850 if the combination of the peer group ID and the MAC address is acceptable. If you want to create a custom ATM address for the switch, enter the address into the Nodal Address Worksheet, Table 3-4, which appears at the end of this chapter.

Caution The default ATM address is created using the primary PXM45 or PXM1E card MAC address. If the default address is being used and the primary PXM card is replaced, the ATM address of the switch changes. After replacing a PXM card, check the switch ATM address and reconfigure it if necessary. To avoid this problem entirely, configure a unique ATM address for the switch.

Selecting the ILMI Address Prefix

Although ILMI is not part of the PNNI specification, ILMI addressing should be coordinated with PNNI addressing to minimize the number of PNNI advertised ATM addresses. The MGX 8850 support ILMI dynamic addressing on UNI ports. When dynamic addressing is enabled, one or more ILMI prefixes can be used to generate ATM addresses for CPE as follows:

1. The CPE retrieves the 13-byte ILMI prefix from the switch.
2. The CPE appends its 7 bytes with the 13-byte prefix to form its AESA.
3. The ILMI running on the switch registers the constructed AESA on the switch.

The default ILMI prefix is the first 13-bytes of the default ATM address, which consists of the 7-byte peer group ID (0x47 0091 8100 0000) plus the unique 6-byte MAC address. If you change the peer group ID for the switch, you should also change the ILMI address prefix so that the bytes that correspond to the peer group ID match the corresponding bytes in the ILMI prefix.

When ILMI is enabled on a UNI port, you can add up to 16 address prefixes for that port. The same ILMI prefix can be assigned to multiple ports. These ILMI prefixes are advertised by PNNI to enable switched virtual circuit (SVC) routing to CPE that use these prefixes.

Enter the ILMI prefixes you plan to use into the Port Address Worksheet, Table 3-5, which appears at the end of this chapter.

Selecting the SPVC Address Prefix

If you set up soft permanent virtual connections (SPVCs), the port at each end of the connection must have a globally unique SPVC address. This address is generated by the switch when the connection is defined and consists of the SPVC prefix and an internally generated number that identifies the port.

The default SPVC prefix is the first 13-bytes of the default ATM address, which consists of the 7-byte peer group ID (0x47 0091 8100 0000) plus the unique 6-byte MAC address. If you change the peer group ID for the switch, you should also change the SPVC address prefix so that the bytes that correspond to the peer group ID match the corresponding bytes in the SPVC prefix.

When planning the SPVC prefix for your WAN, consider the following:

- The SPVC prefix and the ILMI prefix can be the same, or they can be different.
- There can be just one SPVC prefix for each node.
- Once you create a connection using an SPVC prefix, you cannot change the SPVC prefix until all SPVCs have been deleted.

Enter the SPVC prefix into the Nodal Address Worksheet, Table 3-4, which appears at the end of this chapter.

Planning Address Prefixes for AINI and IISP Links

ATM Inter-Network Interface (AINI) and Interim Inter-Switch Protocol (IISP) are two protocols that are used for connecting private PNNI networks to public PNNI networks or to other private PNNI networks. These links enable communications between separately managed networks without exposing the internal structure of each independent network to the other. For example, when an AINI or IISP link is properly configured, a CPE on one independent network can communicate with a CPE on another independent network. However, PTSEs are not transmitted across these links, so the independent networks only have access to ATM addresses that are deliberately shared during configuration.

To enable communications over AINI and IISP links, static addresses must be configured on the end of each link as described in the following guides:

- *Cisco MGX 8850 (PXM1E/PXM45), Cisco MGX 8950, Cisco MGX 8830, and Cisco MGX 8880 Configuration Guide, Release 5.1*
- *Cisco SES PNNI Controller Software Configuration Guide, Release 3*

There is no default prefix for AINI and IISP links, and because these protocols are used on separate link types (not PNNI links), there is no requirement to configure prefixes for AINI and IISP links. However, the PNNI database within each network does store the static addresses, so if there are multiple static addresses that have the same prefix, you can improve PNNI routing efficiency and save configuration time by configuring a summary address prefix that covers multiple ATM addresses. The summary address prefix is a partial ATM address and represents all destinations for which the most significant bytes of the ATM address match the summary address.

When planning AINI and IISP prefixes for your WAN, consider the following possibilities:

- If you are connecting to a network managed by another authority, that authority will probably issue the destination addresses to you.
- The same address or summary address can be configured on more than one port, and multiple addresses can be configured on each port.
- Because the destination devices are not part of the PNNI network, IISP address prefixes do not have to conform to the PNNI level, peer group ID, or node prefix.

Enter any AINI or IISP prefixes into the Port Address Worksheet, Table 3-5, which appears at the end of this chapter.

Selecting Static Addresses for UNI Ports

When CPE devices do not support ILMI, they cannot automatically gain an ATM address from the node, so you must configure a static ATM address on the port that leads to the CPE. You can add up to 255 static addresses on each port, if this number remains within the maximum addresses per node limit.

Multiple ports can be configured with the same static address, but there should be just one CPE that uses each address. When a port leads to multiple CPE that use a common prefix, you can use a summary address to create a single entry that routes to multiple CPE.

Enter the static addresses or summary addresses into the Port Address Worksheet, Table 3-5, which appears at the end of this chapter.

Additional Guidelines for Creating an Address Plan

The following are guidelines for creating an address plan:

- Select the lowest-level PNNI level that will enable future expansion.
- Use PNNI levels that fall on 8-byte boundaries. This improves scalability and makes the PNNI level number easier to work with.
- Use default values for as many entities as possible.
- Use AINI and IISP links to connect to public WANs.
- Do not use Cisco node address defaults in public WANs.

- Confirm that each reconfigured node ID and node address are unique. The switch software does not detect configuration errors caused by duplicate ATM addresses.
- Use summary port prefixes wherever possible to reduce overhead.

Closed User Group Overview

The PNNI Closed User Group (CUG) feature allows network users to form a closed community within a PNNI network. Figure 3-7 shows an example of closed user groups in a network.

Figure 3-7 Closed User Group Example

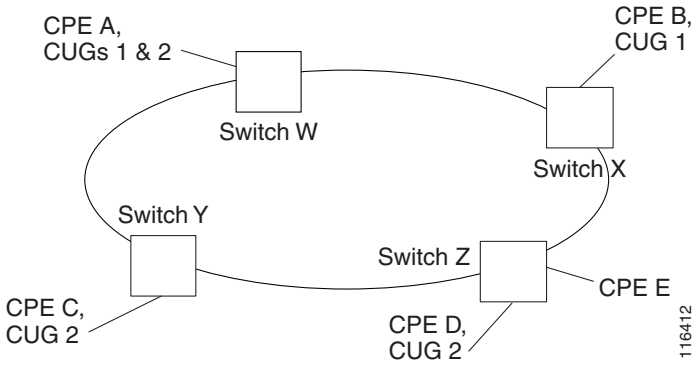

A network user may be associated with one, multiple, or no CUGs. In Figure 3-7, CPE A is a CUG member of CUGs 1 and 2. CPE B, C, and D are members of either CUG 1 or CUG 2. Members of a specific CUG can communicate typically among themselves, but in general not with network users outside of the CUG. In the example, CPE A can communicate with CPE B, C, and D because it is a member of both CUGs. This section will also show how CPE A can be enabled for communications with CPE E.

Specific network users can have additional restrictions preventing them from originating calls to, or receiving calls from, network users of the same CUG. For example, CPE B can be configured so that it cannot originate calls to other CUG 1 members, but it can accept calls from other members.

Configuration options allow a network user to be further restricted when originating calls to, or receiving calls from, network users outside of any CUG membership defined for the network user. In Figure 3-7, CPE E is not a member of any CUG and the default configuration for CUG members will prevent communications between CPE E and the other CPE. Using configuration options, however, CPE1 can be allowed to originate calls to CPE E, and CPE D can be configured to accept calls from CPE E.

The user within a CUG is actually a UNI ATM End Station Address (AESA) or an ATM address prefix, and this address or address prefix can be assigned to more than one interface on a switch. When an ATM address is assigned to more than one CUG, the CPE that use that address must specify the CUG for a connection or accept a configured default CUG called the preferential CUG.

CUG membership is evaluated only when setting up connections. CUG membership is an independent feature and does not interoperate with the address filtering feature.

The CUG feature follows the ITU-T Q.2955.1 recommendation and supports point-to-point and point-to-multipoint connections.

CUGs are managed with the switch CLI. Cisco MGX switches and Cisco LS1010 switches can participate in CUGs. The Cisco WAN Manager (CWM) program does not currently support CUGs.

CUG membership is supported as follows:

- An ATM address or ATM address prefix can be a member of up to 100 CUGs.
- CUGs can be provisioned on up to 200 ATM addresses or prefixes.
- The maximum number of CUGs is 65535.
- An ATM address to which a CUG is assigned can use either the NSAP or E.164 address format.

Note The CUG feature is *not* supported on nodes which are configured with right-justified E.164 addresses using the **cnfe164justify** command.

Planning CUG Configuration Settings

The following sections introduce and provide planning guidelines for the CUG configuration parameters.

Selecting an Interlock Code

A CUG is established by assigning the same 24-byte *interlock code* to two or more prefixes or AESAs on a PNNI network. All prefixes and addresses that share the same interlock code are considered part of the same CUG and can establish connections amongst themselves, unless these connections are blocked by configuration options.

When selecting an interlock code, consider the following guidelines:

- The interlock codes should be managed so that an existing interlock code is not selected for a new CUG.
- Other than being unique to a single CUG, there is no requirement on the contents of the interlock code.

When planning a CUG, consider using the CUG Configuration Worksheet, Table 3-6, which appears at the end of this chapter. Each CUG uses only one interlock code, so place that code in the first row of the worksheet.

Selecting an Index

A CUG index is a number that the administrator specifies when making an address or prefix part of a CUG. The CUG index is mapped to the appropriate interlock code within the switch. During CPE configuration, the appropriate CUG index is configured on the CPE to match the index already defined on the node. When the CPE requests a call, it supplies the index, which is used by the switch to identify the appropriate CUG.

When specifying an index, consider the following guidelines:

- Within each switch, the same, unique index number should be used for all CUG assignments that share the same interlock code.
- The interlock code is not used outside of the switch. Once the index number is mapped to the interlock code, the interlock code is used for all network communications.
- The CPE cannot use the index without further configuration.

The CPE must be configured to specify a particular CUG index during call setup when any of the following conditions exist:

- One or more CUGs are defined for the CPE prefix or address and no preferential CUG is defined.
- Multiple CUGs are defined for the prefix or address and the CPE intends to use a CUG other than the preferential CUG.

If a CPE AESA is a member of only one CUG and that CUG is defined as the *preferential* CUG (see "Specifying a Preferential CUG," which appears later in this chapter), the CPE does not need to be configured to use a particular CUG. The preferential CUG serves as the *implicit* CUG, and is used whenever a CUG index is not specified by the CPE.

Note When a CPE requests a specific CUG during call setup, this is called an *explicit* CUG request.

When planning a CUG using the CUG Configuration Worksheet, Table 3-6, enter the index in the second row of the worksheet.

Selecting CPE Addresses

To add a CPE to a CUG, the configuration process assigns a CPE address or prefix to a CUG interlock code and index. For each CUG assignment, you must specify the following:

- The ATM address or prefix of a local UNI interface.
- The length of the ATM address.
- The ATM address plan, which is either NSAP or E164.

This information is required so that the switch interprets the address or prefix correctly.

If the prefix or address you are assigning to a CUG uses the NSAP format, specify the address length in bits. A full AESA is 160 bits (20 bytes times 8 bits). A shorter address length indicates an ATM address prefix, which assigns all addresses with that prefix to the CUG you specify.

If the prefix or address you are assigning to a CUG uses the E.164 format, specify the prefix or address length in digits.

When planning a CUG using the CUG Configuration Worksheet, Table 3-6, use one worksheet row to identify the CUG configuration for each CUG member. The first column identifies the address or prefix for the CUG member, and the rest of the columns specify the address information, access information, and preferential CUG status.

Selecting Internal CUG Access Options

Internal access options control communications between a specific CUG member and the rest of the CUG. In the CLI, this is expressed in terms of *calls barred*. If you want to block outgoing calls from one CUG member to other CUG members, write the word *outgoing* in the row for the CUG member address in the CUG Configuration Worksheet, Table 3-6. To block calls from other CUG members to a CUG member, write in the word *incoming*.

Note The network administrator can set the internal access during the initial CUG member configuration, or change the configuration later. There is no option to simultaneously block incoming and outgoing communications. If the administrator needs to block incoming and outgoing communications, the member should be removed from the CUG.

Selecting External CUG Access Options

External access options control communications between a specific CUG member and all destinations outside of the CUG. By default, CUG members cannot access destinations outside of the CUG. In the CLI, the external access options are divided into incoming and outgoing controls. The CUG Configuration Worksheet, Table 3-6 provides separate columns where you can enter the incoming and outgoing external access options for each CUG member.

There are two controls for managing incoming communications to CUG members from external sources: *disallowed* and *allowed*.

There are three controls for managing outgoing communications from a CUG member to external destinations: *disallowed*, *per call*, and *permanent*. The *disallowed* control does what its name implies. The *per call* control enables outgoing calls when an outgoing CPE call specifically requests outside access, and the *permanent* control permanently enables outgoing connections as if they were CUG membership connections.

Note The network administrator can set the external access during the initial CUG member configuration, or change the configuration later.

Specifying a Preferential CUG

A preferential CUG is a configuration definition that specifies which CUG membership applies when the CUG member (CPE) does not specify a CUG index during call set up. There can be just one preferential CUG for each CUG member. A preferential CUG assignment is ignored during call setup when the CPE explicitly requests a CUG (using a CUG index).

If a preferential CUG is not assigned to a user and the CPE originates a call without a CUG index, the call is treated as a normal call that is not part of any CUG. Normal calls cannot be established with CUG members unless those members have been configured to communicate outside the CUG.

Note If outgoing calls to the CUG are barred for the user, the CUG cannot be defined as the preferential CUG.

When planning a CUG using the CUG Configuration Worksheet, Table 3-6, use the last table column to indicate if the CUG index and interlock code specified at the top of the table should be the default CUG for this CUG member.

Selecting a Default CUG Address

A default CUG address is a default address that is assigned to a switch to be used for CUG validation when the connected CPE does not signal a calling party ATM address. The default CUG address does not have to match any addresses or prefixes assigned on the switch. It is not used for PNNI routing. It is simply a default address to which a CUG can be assigned.

When a default CUG address is configured, all calls originating and terminating at the switch are treated as CUG calls, regardless of the ATM address. If the CPE does not signal an ATM address, CUG validation uses the default CUG address and evaluates the call based on the CUG membership assigned to the default CUG address. If multiple CUGs are assigned to the default CUG address, it is a good plan to specify one CUG as a preferential CUG.

When planning a default CUG address using the CUG Configuration Worksheet, Table 3-6, remember that there can only be one default CUG address per switch. If you want to assign more than one CUG to the default CUG address, copy Table 3-6 for each CUG assignment and remember that the default CUG address or prefix must be the same in all copies planned for the same switch.

Worksheets

This section provides the configuration worksheets that are described earlier in this chapter.

Table 3-4 is a worksheet that you can use to write down ATM address planning information that applies to the switches in your WAN.

Table 3-4 Nodal Address Worksheet

Node Name	Address Format	Lowest PNNI Level	Peer Group ID	ATM Address	Additional PNNI Levels	SPVC Prefix

Table 3-5 is a worksheet that you can use to write down ATM address planning information that applies to the ports on a single switch. To complete an address plan, complete one Nodal Address Worksheet for the WAN and an individual Port Address Worksheet for each switch in the WAN.

Table 3-5 Port Address Worksheet

Port	ILMI Prefixes	AINI and IISP Prefixes and Addresses	UNI Addresses

Table 3-6 is a worksheet for planning a single closed user group on a single switch. Use a copy of this table for each CUG on a switch. Remember that only one address or prefix can serve as the default CUG address on a switch, and there can be only one preferred CUG per address or prefix.

Table 3-6 CUG Configuration Worksheet

CUG Interlock Code							
CUG Index							
			Internal Access (calls barred)	External Access		Preferred?	Default CUG Address?
Address or prefix	Length	Plan		Outgoing Access	Incoming Access		