

cnfport

Configure Port—PXM1E

The **cnfport** command on the PXM1E lets you configure a logical port on the network interface (UNI/NNI) back card. The system does not confirm successful configuration, so use the **dsport** command to check the changes.

Restrictions Based on Port Status

This section describes modifications you can make to a port while the port is up and those modifications that must occur while the port is administratively down as a result of the **dnport** command.



Note

You cannot use **cnfport** to change the guaranteed rate and maximum rate parameters if a resource partition has been configured for the interface. You would first have to delete the resource partition.



Note

You can change the SCT ID if you first down the port by using **dnport** then using **cnfport**. After you specify a new the SCT, use the **upport** command to return the port to operation.

Syntax

```
cnfport ifNum [-min <guaranteedRate>] [-max <maxrate>] [-sct <sctID>] [-minvpi <minvpi>]
[-maxvpi <maxvpi>]
```

Syntax Description

Note that this command uses keywords (command delineator) to identify the parameters. After you identify the logical port with the *ifNum* parameter, each of the remaining parameters is optional.

<i>ifNum</i>	A logical port (interface) number. Only one logical port is allowed if the line operates as a UNI or NNI. For a virtual interface (VNNI and so on), multiple ports can exist on a line. The range for interface numbers is 1–31.
-min	Specifies the guaranteed rate on a logical port in cells per second (cps). The cumulative guaranteed rate cannot exceed the highest value in the following ranges: <ul style="list-style-type: none"> OC3: 50–353207 cps T3: 50–96000 (PLCP) or 104268 (ADM) cps E3: 50–80000 cps T1: 50–3622 cps E1: 50–4528 cps
-max	Specifies the maximum rate on a logical port in cells per second (cps). <ul style="list-style-type: none"> OC3: 50–353207 cps T3: 50–96000 (PLCP) or 104268 (ADM) cps E3: 50–80000 cps T1: 50–3622 cps E1: 50–4528 cps

-sct	Specifies the number of a service class template (SCT) for the port. The total range is 0–255. Cisco provides SCT numbers 2, 3, 4, 5, 6, 52, 53, 54, and 55. (See the addport and dspsct descriptions for details about the application of these particular SCTs). You can modify one of these SCTs through the Cisco WAN Manager application and assign a unique number in the range 101–255 to the new SCT. This range is for user-created SCTs. You can download the new SCT from the workstation by using the addset command and assign it to a port with the <i>sctID</i> parameter in the cnfport command. To see the ID of the current SCT for this port, use dspport . To see the parameters within the current SCT, use the dspportset command.
-minvpi	The minimum VPI applies to the proprietary enhanced virtual trunking. The <i>minvpi</i> cannot be greater than the <i>maxvpi</i> , <ul style="list-style-type: none"> • EVNNI range: 0 and 4095 • EVUNI range: 0 and 255
-maxvpi	The maximum VPI applies to the proprietary enhanced virtual trunking. The <i>minvpi</i> cannot be greater than the <i>maxvpi</i> , <ul style="list-style-type: none"> • EVNNI range: 0–4095 • EVUNI range: 0–255

Related Commands

addport, delpport, dspport, dspports

Attributes

Log: yes

State: active

Privilege: GROUP1

Example

For logical port 1, configure a guaranteed minimum of 10000 cps and a maximum rate of 20000 cps.

```
MGX8850.7.PXM1E.a > cnfport 1 -min 10000 max 20000
```

cnfpref

Configure Preferred Route—PXM45, PXM1E

The **cnfpref** command lets you change a preferred route, as follows:

- You can modify the definition of existing network elements (NEs) in a route.
- You can delete NEs from a route.
- You can add NEs to an existing route. This capability also lets you complete a route that began with the **addpref** command and whose complete definition involves more than 512 characters.
- You can change an NE to indicate that it is the destination node. This step is necessary to changing the number of NEs in a route. A new destination node must have the highest NE number in the route.

To complete a route that exceeds the 512-character, CLI limit, do the following:

- Specify the route ID and NE syntax.
- Specify a new destination NE position.
- Specify the NEs to complete the route.

For another example, if a route has 4 NEs and you want the route between NE2 and NE4 to become direct (and thus make a route with 3 NEs), you could run the **cnfpref** command once and change the value of the destination node from 4 to 3. In a second iteration of the command, you could remove NE3 by re-assigning the NE4 definition to NE3.

For a detailed description of the Preferred Route feature, see the **addpref** description. To see a list of all preferred routes and obtain a route ID for the **cnfpref** command, use the **dspprefs** command. To see details about a single preferred route, use the **dsppref** command.



Note

This command previously had the name **modpref**.

Syntax

```
cnfpref <routeid> <neSyntax> [-dstNEpos <NE>] [-ne1 {<node>/<port>}] [-ne2 {<node>/<port>}] ...
[-ne20 {<node>/<port>}]
```

Syntax Description

<i>routeid</i>	The preferred route identifier has a range of 1–65535.
<i>neSyntax</i>	<p>Four ways of identifying the network elements are available. If you began a route definition with more than 512 characters by using the addpref command, use the same <i>neSyntax</i>. Type one of the following:</p> <ul style="list-style-type: none"> • nodeidPnportid means the node is specified by <i>nodeId</i> and the port by <i>pnPortId</i>. • nodenamePortid means the node is specified by the node name and the port by portId. • nodeidPortid means the node is specified by the <i>nodeId</i> and the port by <i>portId</i>. • nodenamePnportid means the node is specified by the node name and the port by <i>pnPortId</i>. <p>The <i>node name</i> is specified by the cnfname command.</p> <p>The <i>nodeid</i> is the 22-octet PNNI node identifier.</p> <p>The <i>Portid</i> is the PNNI physical port ID. On a PXM1E for a narrowband service module, the format is <i>slot.port</i>. On a PXM45, the format is <i>slot:subslot.port:subport</i>. For more details, see the section, “PNNI Format,” in Chapter 1, “Introduction.”</p> <p>The <i>PnportID</i> is the PNNI logical port identifier. This form of port identifier is an integer in the range 0–4294967295.</p> <p>Default: none</p>
-dstNEpos	<p>This “destination network element position” identifies the destination NE in the preferred route. For instance, a value of 4 indicates that the fourth NE represents the destination node (including the local node).</p> <p>Range: 1–20</p> <p>Default: none</p>
-ne1, -ne2, ... -ne20	<p>Including the local node, you can specify up to 20 NEs in the preferred route.</p> <p>Each NE is defined by a pairing of a node and a port. The format of these paired elements must conform to the entry for <i>neSyntax</i>. Separate the values in the pairing by a slash and no spaces, but put a space between the keyword and NE, as follows:</p> <p style="text-align: center;">-ne(<i>n</i>) node/port</p> <p>The NE you specify as the destination node must be the highest numbered keyword, otherwise the switch rejects the command. Note that the value of the -dstNEpos parameter actually determines the last NE in the route. For cosmetic reasons or as an aid to clarity, you can make the port identifier at the destination node a 0. This 0 appears in the outputs of the display commands for preferred routes. For example, if a route has 9 NEs, the format would be:</p> <p style="text-align: center;">-ne9 node/0</p>

Related Commands

addpref, delpref, dspref, dsprefs, dspnports, dspnport, addnwnode, delnwnode, dspnwnode, dspnwnodes, cnfndirtes

Attributes

Log: yes State: active Privilege: GROUP1

Examples

Add two NEs to a preferred route 2. First, display route set 2 by using the **dsppref** command.

```
MGX8850.7.PXM.a > dsppref 2
Route Locations
ID = 2 ne1: 112/1:2.1:9,
      ne2: 34/1:1.1:3,
      ne3: 56/1:2.1.7,
      ne4: 23/0
```

Add the following two NEs to the above route and display the modified route:

```
MGX8850.7.PXM.a > cnfpref 2 -ne5 26/1:1.1:1 -ne6 33/0

MGX8850.7.PXM.a > dsppref 2
Route Locations
ID = 2 ne1: 112/1:2.1:9,
      ne2: 34/1:1.1:3,
      ne3: 56/1:2.1.7,
      ne4: 23/1:1.2:1,
      ne5: 26/1:1.1:1,
      ne6: 33/0
```

Modify two NEs in preferred route 2. In this case, NE 1 and NE 3 are the targets. First display the members of route 2 by using the **dsppref** command.

```
MGX8850.7.PXM.a > dsppref 2
Route Locations
ID = 2 ne1: 112/1:2.1:9,
      ne2: 34/1:1.1:3,
      ne3: 56/1:2.1.7,
      ne4: 23/1:1.2:1
```

Change NE 1 and route 3 in route 2, as follows:

```
MGX8850.7.PXM.a > cnfpref 2 -ne1 126/1:1.1:1 -ne3 33/1:2.2:2
```

Display route 2.

```
MGX8850.7.PXM.a > dsppref 2
Route Locations
ID = 2 ne1: 126/1:1.1:1
      ne2: 34/1:1.1:3
      ne3: 33/1:2.2:2
      ne4: 23/1:1.2:1
```

cnfpri-routing

Configure Priority Routing—PXM45, PXM1E, SES (PXM1)

The **cnfpri-routing** command lets you specify *node-level* bandwidth and time-based parameters that PNNI uses to sequence the routing of SPVCs and SPVPs.

This description provides details about the functionality of the Priority Routing feature as it operates at the level of a node, a port, and a connection. The following paragraphs introduce the Priority Routing feature. More detailed information follows the Syntax Description.

The per-connection routing priority is the primary criterion that PNNI uses to determine the point in time—relative to other connections—it routes the connection. You can assign a routing priority to a connection through the **addcon** command (except on an SES). You can modify a priority for a connection through the **cnfcon** command, and this includes the SES. The range of priorities is 0–15. The priorities are progressively lower from 0 to 15. Priority 0 is reserved for network control connections, such as RCC SVCs. For user connections, the range is 1–15.

When a SETUP message with no priority information element (PSIE) arrives at a PNNI port, the port assigns a priority based on the *port-level* priority. You can specify a port-level priority by using the **cnfportsig** command.

A connection created with older software that does not support the Priority Routing feature receives the default priority of 8 after an upgrade.

The routing order is determined by the following two criteria, in order of significance:

1. A priority number in the range 1–15, assigned to a connection through **addcon** or **cnfcon**
2. The bandwidth requirement of the connection

With Priority Routing, the routing, de-routing, and re-routing begins for connections with the highest priority number. Among connections with the same priority, PNNI first acts on the highest bandwidth connections then continues with successively lower bandwidth connections. When PNNI finishes with all connections of one priority, it starts with the next lower priority (and the highest bandwidth connections in *that* priority). Exceptions to this scheme can occur and are described in the section, “[Limitations to Priority Routing](#).”

Priority Routing can aid in overall resource allocation and distribution, redundancy planning, and traffic engineering. The priority routing feature operates in single-peer groups (SPGs) and multi-peer groups (MPGs). This feature is a part of the PNNI controller on the Cisco MGX 8950, MGX 8850 Release 4, MGX 8830, and Service Expansion Shelf (SES) platforms. The remainder of this description usually mentions only the MGX switches, but it should be understood that the functionality also applies to the SES.

For the node level, use the **cnfpri-routing** command to specify the following:

- Bandwidth groups: the groups of connections requiring more bandwidth are routed before groups of connections requiring less bandwidth.
- Two time-based buffers help ensure that PNNI routes connections according to priority and bandwidth requirements. One buffer functions during node startup only, and the other buffer operates in response to events that can trigger connection re-routing.

At the PNNI port level, you can do the following:

- Use the **cnfportsig** command to *modify* the routing priority for incoming calls that have no priority in the SETUP message. (For calls without a priority, PNNI assigns a default priority of 8.)
- Use the **cnfportcac** command to specify the booking factor. (Cisco recommends the default.)

For an individual connection, you can assign a routing priority at the master endpoint when you create it with the **addcon** command or CWM. You can change the routing priority through the **cnfcon** command or CWM. (Note that on an SES, you can use only **cnfcon** to assign a priority to a connection.)

**Note**

The Route Processor Module (RPM) does not support this feature. The node assigns a default routing priority of 8 to new and existing RPM connections.

Syntax

```
cnfpri-routing [-bwstart <start>] [-bwincr <incr>] [-pribuf <time>] [-nodebuf <delay>]
```

Syntax Description**Note**

Specify the routing priority for an SPVC or SPVP through the **addcon** or **cnfcon** command.

-bwstart	<p>This value is the highest cell rate in the lowest bandwidth group. The number of bandwidth groups is fixed at 50. For a detailed description, see the section, “Bandwidth Sorting and Overbooking.”</p> <p>Range: 1-500000 cps</p> <p>Default: 5000 cps</p>
-bwincr	<p>The increment for the cell rate between the upper and lower bounds of each intermediate bandwidth group. (For a detailed description, see the section, “Bandwidth Sorting and Overbooking.”) For example, an increment of 2000 means that a range starting at 1000 cps ends at 12000 cps. This increment does not apply to the following:</p> <ul style="list-style-type: none"> • The group with the lowest bandwidth requirements: for this group (bandwidth group 1), the range is determined by the value for bwstart. • The group with the highest bandwidth requirements: for this group (bandwidth group 50), the range is what remains after computations based on the following: <ul style="list-style-type: none"> – The value for bwstart – The value for bwincr <p>Range: 1–500000</p> <p>Default: 1000</p>

-pribuf The priority buffer is a time counter. It counts down to the moment when PNNI prioritizes all buffered connections for routing. (A connection is buffered due to an event that causes PNNI to re-route the connection.) For a detailed description, see the section, “[Priority Buffer and Re-Route Triggers](#).”

The timer is started when an SPVC or SPVP is released and becomes eligible for routing due to the arrival of a RELEASE message. Once this timer is running, all the SPVCS/SPVPs that become ready for routing are buffered. SPVCs can become eligible for routing due to the following:

- An interface with a master endpoint comes up.
- A routed SPVC or SPVP is released (or failed).
- An SPVC or SPVP is created.
- Route optimization begins.

Range: 0–600, in units of 0.1 seconds (0–60 seconds)

Default: 0

-nodebuf The node buffer is a time counter. It counts down the time to wait before PNNI starts routing connections. This timer is started when the first PNNI logical port comes up after a rebuild or restart. The buffer operates once, after node start-up or node reset.

Range: 0-3000, measured in units of 0.1 seconds (0–300 seconds)

Default: 0

Bandwidth Sorting and Overbooking

This section first describes bandwidth sorting within a priority then the related topics of per-connection and per-interface overbooking.

The bandwidth groups within a routing priority are the secondary criteria for routing and are node-level assignments. The number of bandwidth groups is fixed at 50, but you can specify the following:

- The range of the *lowest* bandwidth group.
- For *each* group between the highest and lowest bandwidth groups, the range of cells per second in a bandwidth group. For example, with a range is 1000 cps, you could have bandwidth groups of 5001cps–6000 cps, 6001–7000 cps, and so on.

The highest bandwidth group is a consequence of the preceding two configurations,

Because the bandwidth groups are node-level, they apply to all priorities: the same groups exist for priority 0, priority 1, priority 2, and so on down to the lowest priority. Specify the granularity of the bandwidth groupings by using the **bwstart** and **bwincr** parameters of the **cnfpri-routing** command.

Connections that require the least bandwidth are grouped at the low end of the range, and connections requiring the most bandwidth are grouped at the top end of the range. The remaining connections are progressively grouped somewhere between the upper and lower bounds.

Within the same routing priority, the connections that require more bandwidth should be routed before those requiring less bandwidth. For a list of circumstances where routing does not follow the priority routing criteria, see the section, “[Limitations to Priority Routing](#).”

**Note**

The configuration sequence contrasts with the actual operation of the feature. When you *define* the bandwidth ranges with the **cnfpri-routing** command, you start by defining the lowest range of bandwidth requirements and work up to the highest bandwidth range. The paragraphs that follow explain this approach.

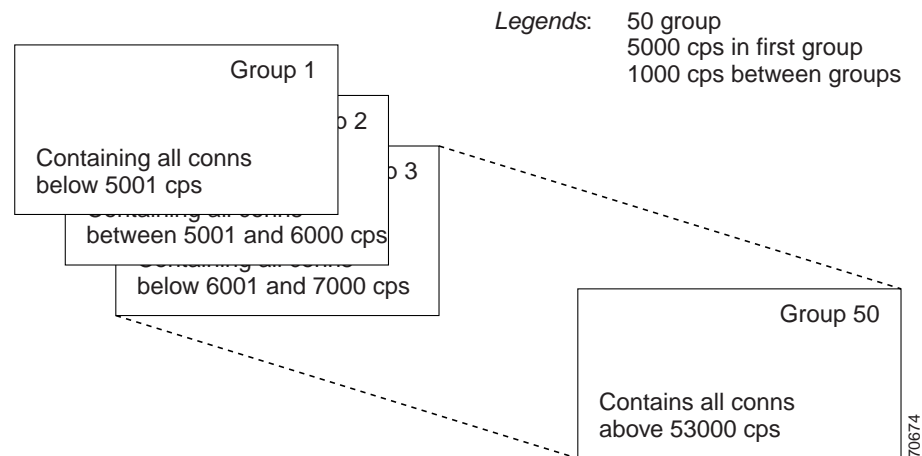
Before examining the bandwidth scheme as a whole, consider how the bandwidth grouping for a priority is done:

- You determine the lowest bandwidth group by specifying the highest rate within the range. For example, if you type 3000, the lowest bandwidth group is 0–3000 cps.
- Increment within each bandwidth group after the lowest. This increment is the same for each subsequent group up to—but excluding—the highest bandwidth group.
- The highest bandwidth group is actually what is left over after you specify the *lowest* bandwidth group and the number of *cells per second* in each subsequent group.

To illustrate, use the fixed number of bandwidth ranges (50), a maximum rate of 5000 cps for the lowest bandwidth group, and 1000 cps for the increment within each subsequent group. Refer to [Figure 2-8](#).

The lowest range accounts for 5000 cps worth of bandwidth, and 49 increments remain. The next 48 increments of 1000 cps each totals 48000 cps. The lowest range (5000 cps) plus the next 48 ranges total 53K worth of bandwidth, and one range is left—the highest. The highest range is all bandwidth above 53000 cps—the leftover after all others are computed.

Figure 2-8 Bandwidth Grouping



The other configurable items that determine the group into which a connection falls are as follows:

- The *per-connection* overbooking factor (or local percent of utilization).
- The *per-interface* overbooking factor of the UNI at the master end has an impact only at the master node because PNNI does not pass it to the via and slave-end nodes.

A connection's *load* is the rate that results from multiplying the connection's maximum cell rate by the percent of utilization. In this case, the load determines the bandwidth range into which the connection falls. For example, a connection with a maximum rate of 10000 cps that has a percent utilization of 50% falls in the range with 5000 cps.

The load of a connection varies with the service class and is calculated based on the forward cell rate requirements, as follows:

CBR =	PCR * Percentage Utilization Factors
VBR (rt and nrt) =	SCR * Percentage Utilization Factors
ABR =	MCR * Percentage Utilization Factors
UBR =	zero

The per-connection overbooking factor affects bandwidth sorting at the three points a connection can touch: the master, transit, and slave nodes. For Priority Routing, only the local percent utilization specified in **addcon** matters. This percent utilization is carried to the far end and the return path.

The per-interface overbooking factor can affect bandwidth sorting at the master node. The per-interface factor is multiplied by the per-connection factor. Thus, if the per-interface factor is less than the Cisco-recommended default of 100%, it further reduces the load. This compounded load affects only the master node because the port-level factor is not signaled to the downstream nodes. The **bookfactor** parameter for the **cnfpnportcac** command specifies the per-interface overbooking factor.

Priority Buffer and Re-Route Triggers

This section describes the mechanism for ensuring that PNNI routes connections by priority while it responds to events that can trigger re-routing.

The time-based priority buffer accumulates re-routing triggers if the priority buffer is non-zero. If the buffer is zero, PNNI routes according to priorities and without delay. Note that out-of-sequence routing can occur as described in the section, "[Limitations to Priority Routing](#)."

The routing triggers are as follows:

- Upping the interface where the master SPVC or SPVP endpoints exist
- Connection failure
- Creation of a new connection
- Route optimization (or grooming)

The priority buffer timer starts up when a connection is released (due to receipt of RELEASE message) and becomes eligible for routing. This particular connection and all other connections that become eligible for routing *while this timer is running* are buffered and are not routed until the timer expires. When the timer expires, all buffered SPVCs and SPVPs go into the priority routing queue. PNNI can then route according to priority and bandwidth requirement. A connection that has the absolute highest priority does not go in this routing event buffer because PNNI can immediately route it.

For networks that need to retain re-routing *performance* without compromising re-routing *time*, the priority buffer can remain the default of 0 seconds (empty). For networks that need to prioritize connections strictly for routing, the buffer can be set to 1.0-second worth of event arrivals. If the network requires either more or less reliance on prioritization, you can adjust the buffer value to more or less than the 1.0-second benchmark duration.

When a RELEASE message arrives for the master endpoint of a connection that already has been routed, PNNI performs the tasks for tearing down a call, then the connection either goes in the buffer or is immediately re-routed.



Note

In-progress SVCs are de-routed before routed connection regardless of routing priority.

Node Buffer

The purpose of the time-based node buffer is to ensure that routing proceeds according to priorities and bandwidth requirements. It does so by delaying the routing process long enough for all PNNI ports to come up after a node is reset. The need for this delay comes from the possibility that a port with low priority connections could come up before a port with high priority connections. In this case, PNNI might route the low priority connections before the high priority connections. The node buffer operates once—after switch start-up or reset.

When a node restarts and the node buffer is non-zero, the system waits the duration specified by **nodebuf** before it prioritizes and then routes the connections. After the waiting period, all connections from the upped interfaces are eligible for routing based on priority and bandwidth.

Locations of Priority Routing Relevance

This section describes the points where PNNI utilizes Priority Routing information.

Routing and de-routing create different requirements at different points in the network. Due to the nature of source routing, a priority has only local significance—at the originating node (where the master endpoint exists) the highest priority connections are de-routed first.

For initial routing or re-routing, downstream nodes do not use information about the priority of the connection. Nevertheless, for *de-routing* purposes, the originating node must pass the priority information in the SETUP message to the downstream nodes. The priority information travels in a standards-based priority information element (PSIE). The premise of prioritized de-routing is that highest priority connections are de-routed first and re-routed before lower priority connections.

Upon a transit point facility failure, connections are de-routed based on priority and bandwidth group. Higher priority connections are de-routed before lower priority connections. Within the same priority group, higher bandwidth connections are de-routed before lower bandwidth connections.

Another possible factor is a difference in bandwidth parameters at opposite ends of an SPVC—an asymmetrical connection. If the far end was added with different bandwidth parameters, it could fall in a different priority group than the near-end if de-routing becomes necessary.

Note that if one of the transit nodes discards (without relaying) the PSIE in the SETUP message, the downstream nodes assign the connection a priority based on each ingress PNNI interface's configured default routing priority at the downstream node. To configure the routing priority default at the ingress of a PNNI port, use the **cnfpnportsig** command and its **-svcroutingprio** parameter.

Limitations to Priority Routing

When events occur far apart in time, prioritization has no effect. Therefore, routing that is out of the prioritized sequence can occur. Numerous situations can cause the controller to route connections not strictly based on priority. The following are some of these circumstances:

- The priority routing function is limited to the local node, so a lower priority connection from one node could be routed before a higher priority connection from another node.
- If PNNI cannot route a higher priority connection (for example, during transient congestion or a routing failure condition), it suspends routing attempts for that connection while it continues to route other, lower priority connections.
- When routing triggers arrive rapidly, prioritization does not take effect in the routing process.
- If you add a high priority connection on a node, routing that connection does not supersede lower priority connections that are currently being routed.

- Even though a downstream node is capable of de-routing connections to the master node in the correct sequence by priority, the RELEASE messages may still arrive at the master node over different paths and out-of-sequence. Therefore, PNNI does not necessarily re-route according to priority if the RELEASE messages are out of sequence. The configurable buffers can alleviate the out-of-sequence arrival of RELEASE messages.
- While PNNI is busy re-routing an SPVC or SPVP, you cannot run the **cnfpri-routing** command.

SVCs and Priority Routing

To assist re-route prioritization of SPVCs at the master endpoint, PNNI releases SVCs in a sequence similar to the re-routing sequence when a transit point facility fails. (At via nodes, SPVCs are SVCs.) All SVCs are released based on priority and bandwidth grouping. This priority is signaled in a PSIE in the SETUP message. If this PSIE is missing, a port-level priority is assigned to the connection. You can modify port-level priority through the **cnfnpportsig** command. Its default is 8.

The **svcroutingpri** option for the **cnfnpportsig** command lets you assign a routing priority at the port level for SVCs or for an SPVC or SPVP that has no associated priority. The prioritized de-routing of SVCs assists in the prioritized re-routing of SPVCs and SPVPs.

Additionally, the routing priority for connections may be missing, as follows:

- Some via nodes may not support the Priority Routing feature. In this case, the port-level **svcroutingpri** parameter tags the connection with a routing priority.
- One of more SPVCs or SPVPs may lose the priority information element. In this case, the port-level **svcroutingpri** parameter tags the connection with a routing priority.

The PNNI network controller does not automatically re-route SVCs. No interaction exists between the instantaneous routing of SVC connections and the sequenced re-routing of SPVCs. An SVC initiated by a CPE or other PNNI node must be re-initiated by that entity. An SVC initiated from within the node is re-initiated by the same source.

Priority Routing and Grooming

Grooming (or route optimization) and Priority Routing can interact when multiple connections are involved. (A single connection does not require *prioritization*, but grooming *does* apply to a single connection.) Priority Routing takes precedence over grooming. If connections exist in the priority buffer when grooming begins, PNNI routes the buffered connections first. PNNI services the route optimization queue only when the standard routing priority queue is empty. See the **cnfrteopt** description for information on grooming.

Caveat for the cnfcon Command

If you use the **cnfcon** command to modify *only* the routing priority of a connection on a service module or the PXM1E, PNNI does not immediately re-route the connection. Nevertheless, if you run **dsppcon** for such a changed connection at the master endpoint, it immediately shows the changed priority *configuration* even before PNNI re-routes the connection. In contrast, at the slave-end or a transit node, **dsppncon** shows the old priority until PNNI re-routes the connection.

Related Commands

dsppri-routing, **dsppcon**, **dsppcons**, **dsppncon**, **dsppncons**, **addcon**, **cnfcon**, **cnfnpportsig**, **dsppnportsig**, **cnfnpportcac**, **dsppnportcac**, **dsppnport**

Attributes

Log: yes State: active, standby Privilege: SUPER_GP

Example

Specify the following priority routing parameters then check the configuration:

- A maximum rate of 2000 cps for the lowest bandwidth group
- An increment of 500 cps for each bandwidth group
- A priority buffer of 10 seconds
- A node buffer of 100 seconds

```
M8850_LA.8.PXM.a > cnfpri-routing -bwstart 2000 -bwincr 500 -pribuf 100 -nodebuf 1000
```

```
M8850_LA.8.PXM.a > dsppri-routing
```

```
Priority Routing Configuration
```

```
-----
```

```
Number of bandwidth groups: 50
```

```
Size of first bandwidth group (in cps): 2000
```

```
Increment between bandwidth groups (in cps): 500
```

```
Routing event buffer size (in 0.1-seconds): 100
```

```
Node startup routing delay (in 0.1-seconds): 1000
```

```
M8850_LA.8.PXM.a >
```

cnfpasswdreset

Configure Password Reset—PXM45, PXM1E

The **cnfpasswdreset** command lets you enable or disable the functionality of the series of key strokes that reset the default, Cisco password. This sequence of key strokes is as follows:

ESC CTRL-Y

- To turn on the functionality of the preceding sequence (password reset), type **cnfpasswdreset on**.
- To turn off the functionality of the preceding sequence, type **cnfpasswdreset off**.

Syntax

cnfpasswdreset <flag>

Syntax Description

<i>flag</i>	A Boolean expression to enable or disable password reset: enter “on” to enable or “off” to disable the sequence of keys that resets the password.
-------------	---

Related Commands

dsppswdreset

Attributes

Log: yes

State: active

Privilege: SERVICE_GP

Example

Enable command reset, then check its status by executing **dsppswdreset**.

```
pop20one.7.PXM.a > cnfpasswdreset on
Password Reset feature being enabled
```

```
pop20one.7.PXM.a > dsppswdreset
Password Reset feature currently enabled
```

cnfqosdefault

Configure Quality of Service Default—PXM45, PXM1E

The **cnfqosdefault** command lets you specify default, switch-level QoS values for three service classes. The applicable service classes are CBR, rt-VBR, and nrt-VBR. The switch applies these default values to an SVC or SPVC if the incoming setup message does not contain the QoS specification. For an SPVC, the values specified through **addcon** or **cnfcon** override the defaults configured through the **cnfqosdefault** command.

You can specify defaults for one service class at a time. In addition to the bandwidth parameters, you can either enable (activate) or disable the default configuration. The default state is disabled. Therefore, be sure to *enable* the configuration for each QoS if you want PNNI to use it. You can configure the parameters and leave them disabled until a suitable time (see Example).

Each bandwidth parameter is optional:

- Maximum cell transfer delay
- Peak-to-peak cell delay variation
- Maximum cell loss ratio for cells with CLP=0
- Maximum cell loss ratio for cells with CLP=1 or 0

Syntax

```
cnfqosdefault <cbr | rtvbr | nrtvbr> [<-maxctd> maxctd] [<-ppcdv> ppcdv]
[<-maxclrlp0> maxclrlp0] [<-maxclrlp01> maxclrlp01] [<-enable> {yes | no}]
```

Syntax Description

cbr, rtvbr, or nrtvbr	The service class for the current iteration of the command: type either “cbr,” “rtvbr,” or “nrtvbr.”
-maxctd	The maximum cell transfer delay has a range of 0– 65535 milliseconds. Default: (none)
-ppcdv	The peak-to-peak cell delay variation has a range of 0–16777215 microseconds. Default: (none)
-maxclrlp0	The maximum cell loss ratio for CLP0 can be an integer in the range 1–15. Default: (none)
-maxclrlp01	The maximum cell loss ratio for CLP0+1 can be an integer in the range 1–15. Default: (none)
-enable	The entry for enabling or disabling the switch level defaults for the current service class. Type “yes” to enable or “no” (or leave in the default state). Default: no

Related Commands

clrqosdefault, dspqosdefault

Attributes

Log: yes State: active Privilege: GROUP1

Example

For CBR, configure and enable the following:

- Maximum cell transfer delay of 100 milliseconds
- Maximum cell delay variation of 1000 microseconds
- Maximum cell loss ratio for CLP0 of 10
- Maximum cell loss ratio for CLP0+1 of 5

After configuring the defaults, disable them, check the configuration with **dspqosdefault**, then re-enable and re-check them. Note that the output of **dspqosdefault** has been edited to show only the CBR values.

```
8850_NY.7.PXM.a > cnfqosdefault cbr -maxctd 100 -ppcdv 1000 -maxcltclp0 10 -maxclrc1p01 5
-enable yes
```

```
8850_NY.7.PXM.a > dspqosdefault
```

```
Service Category = cbr                Qos Default Enable = yes
MaxCTD =          100                 ppCDV =              1000
MaxClrClp0 =      10                 MaxClrClp01 =        5
```

```
8850_NY.7.PXM.a > cnfqosdefault cbr -enable no
```

```
8850_NY.7.PXM.a > dspqosdefault
```

```
Service Category = cbr                Qos Default Enable = no
MaxCTD =          100                 ppCDV =              1000
MaxClrClp0 =      10                 MaxClrClp01 =        5
```

```
8850_NY.7.PXM.a > cnfqosdefault cbr -enable yes
```

```
8850_NY.7.PXM.a > dspqosdefault
```

```
Service Category = cbr                Qos Default Enable = yes
MaxCTD =          100                 ppCDV =              1000
MaxClrClp0 =      10                 MaxClrClp01 =        5
```


cnfrmrsrc

Configure Resource Monitoring—PXM45, PXM1E

The **cnfrmrsrc** command lets you configure the monitoring of various system resources, such as memory allocations, logical disk partitions, and so on. To see a list of all resource settings or a particular setting, use the **dsprmrsrcs** command or the **dsprmrsrc** command, respectively.

Syntax

```
cnfrmrsrc <rsrcId> [-poll] [-loth] [-medth] [-hith]
```

Syntax Description

rsrcId In the current release of the PXM45, *rsrcId* is a number in the range 0–16 that identifies the PXM45 monitoring resource to configure. (See next parameter description entry for PXM1E.)

On the PXM45, the resources are as follows:

- 0: SSI Static Memory
 - 1: SSI dynamic Memory
 - 2: SSI SNMP Memory
 - 3: SSI IPC Small Buffer
 - 4: SSI IPC Medium Buffer
 - 5: SSI IPC Large Buffer
 - 6: SSI IPC Huge Buffer
 - 7: SSI IPC mblk Buffer
 - 8: Hard Disk Space C
 - 9: Hard Disk Space D
 - 10: Hard Disk Space E
 - 11: Hard Disk Space F
 - 12: CPU Peak Utilization
 - 13: System Memory
 - 14: SSI Timer
 - 15: SSI File Descriptor
 - 16: VxWorks File Descriptor
-

rsrcl In the current release of the PXM1E, *rsrcl* is a number in the range 0–17 that identifies the PXM1E monitoring resource to configure.

On the PXM1E, the resources are as follows:

- 0: SSI Static Memory
- 1: SSI dynamic Memory
- 2: SSI stats Memory
- 3: SSI SNMP Memory
- 4: SSI IPC Small Buffer
- 5: SSI IPC Medium Buffer
- 6: SSI IPC Large Buffer
- 7: SSI IPC Huge Buffer
- 8: SSI IPC mblk Buffer
- 9: Hard Disk Space C
- 10: Hard Disk Space D
- 11: Hard Disk Space E
- 12: Hard Disk Space F
- 13: CPU Peak Utilization
- 14: System Memory
- 15: SSI Timer
- 16: SSI File Descriptor
- 17: VxWorks File Descriptor

-poll This option is the polling period, the number of seconds from the start of one poll to the start of the next poll.

Range: 5–86400 seconds

-loth The low threshold is a percent that should be greater than 0 and less than the medium threshold.

Range: 1–100

-medth The medium threshold is a percent that should be between the low and high thresholds.

Range: 1–100

-hith The high threshold should be higher than the medium threshold but less than 100.

Range: 1–100

Related Commands

dsprmalms, dsprmrsrc, dsprmrsrcs, dsprminfo

Attributes

Log: no State: active, standby, init Privilege: CISCO_GP

Example

Check the polling timer for logical disk F. Change it to 86400 seconds, then check it again.

```
M8830_SF.2.PXM.a > dsprmrsrc 12
===== [Resource Information] =====
name      : Hard Disk Space - F: state      : OK
Maximum size      : 1000                      Cur size : 999 (MByte)
ignoreMed         : NO                        LowWaterMark : 999
Threshold Type    : percent                    poll interval: 30
Low threshold value : 100                      Medium threshold value: 150
  High threshold value : 200                    Low threshold percent : 100
Medium threshold percent: 150                    High threshold percent: 200

===== [Action Info] =====
          Send          Alarm
          Trap          Critical Major  Minor
-----
Low Action: yes          no          yes      no
Med Action: yes          no          no       yes
Ok Action : yes          no          yes      no

===== [Statistics] =====
total poll: 134246      failed poll: 0
ok to low : 0           ok to med  : 0
med to low: 0           low to ok  : 0
med to ok : 0

===== [Others] =====
```

Change the poll timer to 86400.

```
M8830_SF.2.PXM.a > cnfrmrsrc 12 -poll 86400
valPoll = 86400

M8830_SF.2.PXM.a > dsprmrsrc 12
===== [Resource Information] =====
name      : Hard Disk Space - F: state      : OK
Maximum size      : 1000                      Cur size : 999 (MByte)
ignoreMed         : NO                        LowWaterMark : 999
Threshold Type    : percent                    poll interval: 86400
Low threshold value : 100                      Medium threshold value: 150
  High threshold value : 200                    Low threshold percent : 100
Medium threshold percent: 150                    High threshold percent: 200

===== [Action Info] =====
          Send          Alarm
          Trap          Critical Major  Minor
-----
Low Action: yes          no          yes      no
Med Action: yes          no          no       yes
Ok Action : yes          no          yes      no

===== [Statistics] =====
total poll: 134252      failed poll: 0
ok to low : 0           ok to med  : 0
med to low: 0           low to ok  : 0
med to ok : 0

===== [Others] =====

M8830_SF.2.PXM.a >
```

cnfrtprm

Configure Reroute Retry Parameters—PXM45, PXM1E

The **cnfrtprm** command allows you configure the time periods that the switch waits between each reroute retry attempt.

When an SPVC fails, the system immediately attempts to reroute the connection. If the first reroute attempt fails, the switch keeps trying to reroute the connection according to the *slow retry interval* (**-slowtmr**) and the *fast retry interval base* (**-fasttmrbase**).

The *fast retry interval base* is an incremental value (in 100-millisecond units) that is incremented each time the switch attempts to reroute the connection and fails. The switch then waits the incremented amount of time before it attempts to reroute the connection again. The *fast retry interval base* continues to increment after each reroute attempt until it is equal to the *slow retry interval* value or until the reroute succeeds.

The *slow retry interval* is a fixed value (in seconds) that occurs between all subsequent reroute attempts. After the *fast retry interval base* reaches the *slow retry interval*, the switch attempts to reroute the connection at the rate of the *slow retry interval*. No limit exists for the number of reroute attempts once the *slow retry interval* begins.

For example, if the *fast retry interval base* is 50 100-millisecond intervals (5 seconds) and the *slow retry interval* is 300 seconds (5 minutes), the switch attempts to reroute the connection 5 seconds after the first attempt, 10 seconds after the second attempt, 15 seconds after the third attempt, and so on until the *fast retry interval base* equals 300 seconds (5 minutes). After that, the switch continues to attempt to reroute the connection every 5 minutes or until the reroute is successful.

Syntax

```
cnfrtprm [-slowtmr <slow retry interval>] [-fasttmrbase <fast retry interval base>]
```

Syntax Description

-slowtmr	The range for <i>slow retry interval</i> is 1–300 seconds. The default is 60 seconds. The <i>slow retry interval</i> must be greater than <i>fast retry interval base</i> .
-fasttmrbase	The <i>fast retry interval base</i> is a multiplier of 100-millisecond units. The range is 1–3000. The default is 50 100-milliseconds units (5 seconds).

Related Commands

dsprtrprm

Attributes

Log: yes State: active Privilege: SUPER_GP

Example

Set the slow timer to 300 second intervals and the fast timer base to 7 seconds (70 x 100 milliseconds). Check the results by executing **dsprtrprm**.

```
8850_NY.7.PXM.a > cnfrtprm -slowtmr 300 -fasttmrbase 70
```

```
8850_NY.7.PXM.a > dsprrtparm

Global SPVC Retry Parameters:
-----
Slow Retry Interval: 300 sec
Fast Retry Interval Base: 70 (in 100 msec)

8850_NY.7.PXM.a >
```

cnfrsoprtn

Configure Resource Partition—PXM1E

The **cnfrsoprtn** command lets you modify the bandwidth and other resource partitioning on a logical port. The entity that makes use of the resources in a partition is a network controller. For many partition parameters, you can dynamically make modifications—without administratively downing the port—by using the **cnfrsoprtn** or **cnfpart** command. However, before you can modify the minimum or maximum VPI or VCI, the port must be down.

The existing controllers are the Private Network to Network Interface (PNNI) and the Label Switch Controller (LSC). In the current release of the PXM1E, only PNNI is supported. Before you add a resource partition, have a plan for future changes, such as the support of a new controller.

A resource partition consists of:

- A guaranteed percentage of bandwidth.
- VPI and VCI ranges.
- Guaranteed minimum and maximum number of connections. Note that the *maximum* number of connections must be *greater* than 10.



Note

The **cnfpart** and **cnfrsoprtn** commands are identical and interchangeable. The name “cnfrsoprtn” is consistent with the corresponding command in Release 1 of the MGX 8850 node. Use the name that suits you. This interchangeability also applies to all the other partition commands.

Ports, Partitions, Controllers, and Interface Types

This section contains details regarding ports, partitions, and controllers that you should note before adding a partition.

On each port—regardless of the interface type—a controller can have one partition. Therefore, on a port, you can add one partition for PNNI and one for LSC (but keep in mind that the PXM1E currently uses only PNNI). This requirement applies regardless of whether an interface (specified through **addport**) is UNI, NNI, VUNI, or VNNI.

The pairing of partition ID and controller ID must be the same across all interfaces. In this situation, the interface number uniquely identifies the partition. For example, on a PXM1E-4-OC3 with two UNIs and two NNIs, you could specify:

- Logical interface 1 (on line 1), partition ID 1, controller ID 2
- Logical interface 2 (on line 2), partition ID 1, controller ID 2
- Logical interface 3 (on line 3), partition ID 1, controller ID 2
- Logical interface 4 (on line 4), partition ID 1, controller ID 2

The VPIs and VCIs you modify with the **cnfrsoprtn** command must be within the range specified when the port was created through the **addport** command,

Syntax

```
cnfrsoprtn -if <if> -id <partitionID> -ctrl <controllerID> -emin <egrMinBw> -emax <egrMaxBw>
-imin <ingMinBw> -imax <ingMaxBw> -vpin <minVpi> -vpmax <maxVpi> -vcmin <minVci>
-vcmax <maxVci> -mincon <min connections> -maxcon <max connections>
```

Syntax Description



Note

On a standard virtual trunk, where the interface type is VNNI or VUNI, the *minVpi* and *maxVpi* must be the same. See the **-vpmin** and **-vpmax** descriptions for ranges.

On an enhanced virtual trunk, where the interface type is EVNNI or EVUNI, the *minVpi* and *maxVpi* can be different. The *maxvpi* cannot be less than the *minvpi*.

-if	The range for logical interface (port) numbers is follows 1–31.
-id	The range for partition ID numbers is 1–20.
-ctlr	<p>The <i>controllerID</i> is a number that identifies a network controller. The PXM1E supports only the PNNI controller—option 2. (The range for <i>reserved</i> controller IDs is 1–3.)</p> <p>The reserved controller IDs are as follows:</p> <ul style="list-style-type: none"> • 1 = PAR (Portable AutoRoute)—currently not used on the PXM1E or PXM45 • 2 = PNNI • 3 = LSC (Label Switch Controller, also known as MPLS for Multiprotocol Label Switching) is not supported on the PXM1E <p>(The absolute range for the PXM1E is 1–254.)</p>
-emin	A guaranteed percentage of egress bandwidth. Each unit of <i>egrMinBw</i> is 0.00001 of the total bandwidth on the port. (An <i>egrMinBw</i> of 1000000 = 100%.) This approach provides a high level of granularity.
-emax	A maximum percentage of the bandwidth. Each unit of <i>egrMaxBw</i> is 0.00001 of the total bandwidth available to the port. (An <i>egrMaxBw</i> of 1000000 = 100%.) The resulting bandwidth must be at least 50 cps.
-imin	A guaranteed percentage of the ingress bandwidth. Each unit of <i>ingMinBw</i> is 0.00001 of the total bandwidth available to the port. For example, an <i>ingMinBw</i> of 1000000 = 100%.
-imax	A maximum percentage of the ingress bandwidth. Each increment of <i>ingMaxBw</i> is 0.00001 of the total bandwidth on the port. For example, an <i>ingMaxBw</i> of 1000000 = 100%. Note that the maximum ingress bandwidth must be at least 50 cps.
-vpmin	<p>The minimum VPI applies to the proprietary enhanced virtual trunking. The <i>minvpi</i> cannot be greater than the <i>maxvpi</i>,</p> <ul style="list-style-type: none"> • EVNNI range: 0 and 4095 • EVUNI range: 0 and 255
-vpmax	<p>The maximum VPI applies to the proprietary enhanced virtual trunking. The <i>minvpi</i> cannot be greater than the <i>maxvpi</i>,</p> <ul style="list-style-type: none"> • EVNNI range: 0–4095 • EVUNI range: 0–255
-vcmin	Minimum VCI in the range 32–65535.
-vcmax	Maximum VCI in the range 32–65535.

-mincon	<p>Specifies the guaranteed number of connections. On the PXM1E UNI/NNI, the ranges vary according to the line types, as follows:</p> <ul style="list-style-type: none"> • For OC3, T3, and E3 lines, the range is 10–27000. • For T1 and E1 lines, the range is 10–13500. <p>(On the AXSM series of cards, the range is 10 through the maximum number of connections in the port group. See dspecd for information about port groups. On narrowband service modules, the range varies: see the CLI of individual cards.)</p>
-maxcon	<p>Specifies the guaranteed number of connections. On the PXM1E UNI/NNI, the ranges vary according to the line types, as follows:</p> <ul style="list-style-type: none"> • For OC3, T3, and E3 lines, the range is 10–27000. • For T1 and E1 lines, the range is 10–13500. <p><i>maxConns</i> cannot be less than <i>minConns</i>.</p> <p>(On the AXSM series of cards, the range is 10 through the maximum number of connections in the port group. See dspecd for information about port groups. On narrowband service modules, the range varies: see the CLI of individual cards.)</p>

Related Commands

addrscrtn, delrscrtn, dsprscrtns, dsprscrtn

Attributes

Log: yes State: active Privilege: GROUP1

Example

Configure the following:

- The logical port (*ifNum*) is 1.
- The partition number is 1.
- The controller is PNNI (number 2).
- The ingress and egress each have a minimum of 10% and a maximum of 15% of the bandwidth.
- VPI range is 20–100.
- VCI range is 1–32767.
- Minimum guaranteed number of connections is 1000.
- Maximum number of connections is 2000.

```
MGX8850.7.PXM1E.a > cnfrscrtn -if 1 -id 1 -ctlr 2 -emin 100000 -emax 150000 -imin 10000
-imax 15000 -vpmin 20 -vpmax 100 -vcmin 1 -vcmax 32767 -mincon 1000 -maxcon 2000
```


cnfrteopt

Configure Route Optimization—PXM45, PXM1E

The **cnfrteopt** command lets you configure automatic, periodic route optimization to improve bandwidth utilization. This optimization is also called *connection grooming*. (To force *immediate* optimization, use **rrtcon** for a single connection or **optrte** for a range of connections.) The load created by grooming is very small and does not cause congestion.



Note

If a connection with a preferred route has been taken off its preferred route, it can regain its preferred route through grooming. Re-routing can occur by scheduled grooming or through forced optimization.



Note

The **cnfrteopt**, **cnfrteoptthld**, and **optrte** commands rely on the criterion of route cost, but PNNI re-routes point-to-multipoint connections based on branching criteria, so these commands do not apply to P2MP connections. Use **rrtcon** to trigger a P2MP re-route.

You can choose a time for optimization so that disruption is minimal. Grooming is a background process and does not attempt to optimize all connections at once. For example, you could specify one hour of grooming for a range of connections starting at 1:00 AM.

The nature of SPVCs and SPVPs provides a reason for periodic grooming: during the course of daily operation, better routes may become available. Factors in determining a better route are as follows:

- The maximum cost (*maxcost*) for the connection. See the **addcon** description for details on *maxcost*.
- The cost-per-link (known as *administrative weight* or AW) that PNNI uses for routing. For information on AW, see the **cnfpnni-intf** description.
- The routing policy for on-demand routing can affect whether a connection receives the optimal route. Very infrequently, a connection may not have the optimal route after an optimization pass. Such a connection receives the optimal route on the next pass. In the **cnfpnni-routing-policy** command, the choice of “bestfit” for the **-onDemand** parameter ensures that the connection receives the optimal route in the course of grooming. See the **cnfpnni-routing-policy** description.



Note

If you do not specify a *maxcost* with either the **addcon** or **cnfcon** command, the routing protocol uses the AW on only the forward links to calculate a new route for the connection. If the connection has a specified *maxcost*, the routing protocol calculates possible routes by using the AW in both directions.

Usage Guidelines

Note the following characteristics of route optimization:

- Within a range of connections, the **cnfrteopt** command applies to only the master endpoints. The slave endpoints are not processed by **cnfrteopt**.
- Route optimization applies to routed connections only. The switch fabric performs load comparison between the routing cost of a connection's current route and the new, potential, best route.

By default, PNNI calculates that a route is better if its routing cost is 30% less than the current cost. You can change this cost threshold through the **cnfrteoptthld** command. You can also configure a specific route cost for each service type. See the **cnfrteoptthld** description for details.

The following briefly characterizes the defaults for **cnfrteopt**:

- The default state is disabled.
- If you do not specify a range, all connections on the port are subject to optimization.
- If you do not specify an interval, optimization begins every 60 minutes.
- If you do not specify a TOD, the default is any time during the day (but still subject to the interval).

Syntax

```
cnfrteopt <portid> [enable / disable] [-range <starting-vpi/vci..ending-vpi/vci>]
[-interval <interval>] [-tod <start-time..end-time>]
```

Syntax Description

<i>portid</i>	<p>The format of the PNNI physical port identifier can vary, as follows:</p> <ul style="list-style-type: none"> • On a PXM45: <i>slot:subslot.port:subport</i> • On a PXM1E for UNI/NNI back card: <i>slot:subslot.port:subport</i>. On the UNI/NNI back card, the subslot is always 2, but the <i>slot</i> depends on the chassis, as follows: <ul style="list-style-type: none"> – In an MGX 8850 chassis, <i>slot</i> is always the logical slot 7. – In an MGX 8830 chassis, <i>slot</i> is always the logical slot 1. • On a PXM1E for a narrowband service module (NBSM): <i>slot.port</i>. <p>For more details, see the section, “PNNI Format,” in Chapter 1, “Introduction.”</p>
enable / disable	Enables or disables route optimization. The default is disabled, but if grooming is operational and you want to disable it, run cnfrteopt and enter “disable.”
-range	<p>The range of connections for grooming. The VPI of the starting SPVC must be less than the ending VPI, and the starting VCI must be less than the ending VCI.</p> <p>Use the notation as it appears on the syntax line: type a slash between the VPI and VCI and two dots with no spaces between the starting and ending values. For example, 100/1000..200/10000 is a valid entry. The ranges are as follows:</p> <ul style="list-style-type: none"> • VPI: 0–4095 • VCI: 32–65535 <p>In the variable parameter <i>starting-vpi/vci..ending-vpi/vci</i>, the starting and ending VPIs are independent of the starting and ending VCIs. For example, -range 3/40..5/50 means the following:</p> <p style="padding-left: 40px;">Optimize all SPVCs that have a VPI in the range 3–5 and a VCI in the range 40–50.</p> <p>This example range could include 3/40, 3/45, 4/45, 5/45 and 5/50 but not 4/60.</p> <p>Note The default for -range is <i>all</i> connections on the PNNI port specified by <i>portid</i>. Therefore, to groom all connections on the port, leave out the -range parameter.</p>

-interval	<p>Keyword that specifies the frequency at which grooming begins. The units of measure are minutes. The range is 10–10000. The default is 60. Counting starts at one of two moments:</p> <ul style="list-style-type: none"> • The moment you run cnfrteopt • The starting time specified by TOD in cnfrteopt <p>If the interval is less than half the amount of time specified by the <i>start-time..end-time</i> parameter, route optimization may begin more than once during the time period. For example, if the periods of optimization are two hours beginning at midnight and 4:00 AM and the interval is one hour, route optimization could occur two to four times per day.</p>
-tod	<p>Keyword that specifies the time to start and stop grooming. The format is a 24-hour clock: 00:00–23:59. The default for both <i>start</i> and <i>end-time</i> is 00:00. If you run cnfrteopt during the time specified by tod, the optimization cycle begins during the next time interval.</p> <p>If the time for the node changes (by way of the cnftime command, for example), the node might skip one optimization cycle.</p> <p>Note Use the notation in the Syntax section: type two dots with no spaces between starting and ending times.</p>

Related Commands

cnfrteoptthld, oprtrt, dsprrteoptcnf, dsprrteoptstat

Attributes

Log: yes State: active Privilege: GROUP1

Examples

For logical port 2 on the lower bay of the service module in slot 1, configure 1 hour of connection grooming starting between 1:00 and 3:00 AM local time. The range of SPVCs is 100.1000 through 100.10000.

```
cnfrteopt 1:2.1:2 enable -range 100/1000..100/10000 -interval 60 -tod 01:00..03:00
```

cnfrteoptthld

Configure Route Optimization Threshold—PXM45, PXM1E

The **cnfrteoptthld** command lets you set node-level thresholds that PNNI uses during route optimization to trigger re-routing. Re-routing occurs if a new route for a connection provides sufficient improvement. For example, if PNNI calculates during route optimization that a new route for a particular connection gives a 30% reduction in routing cost, the connection gets that new route.



Note

Route optimization (or grooming) is enabled at the PNNI port level by the **cnfrteopt** command. (The **cnfrteopt** command also lets specify a time and VPI/VCI range for connection grooming.) For a port where grooming is disabled, the **cnfrteoptthld** command has no effect.

The criteria for selecting a new route are thresholds in the form of a percent of improvement (or *reduction*) in routing cost and a specific—or “absolute”—routing cost. The percent of cost reduction applies to all service classes. The absolute cost is configurable for each service class. These criteria are configurable for each service class.

The controller does a logical AND of the criteria of percent of improvement and absolute route cost. A connection of groomed if *both* the following are true:

The new route cost is less than or equal to the old route cost *minus* the absolute grooming threshold.

$$\text{NEW} \leq (\text{OLD} - \text{ABSOLUTE})$$

The new route cost is less than or equal to the old route cost *multiplied* by the *result* of 1 minus the percent of improvement (1 – % of improvement).

$$\text{NEW} \leq \text{OLD} \times (1 - \% \text{ of improvement})$$



Note

If you want the controller to consider only the percent of improvement for a particular service class, leave the absolute threshold for that service class at the default of 0.

Characteristics of Grooming Thresholds

Because the thresholds are node-level settings, grooming relies on the same value regardless of the current cost.

The number of hops has no bearing on the grooming threshold.

In the current release, consideration of the absolute cost ignores the routing metrics of administrative weight, cell transfer delay, and cell delay variation.

Syntax

```
cnfrteoptthld [-percent <p-value>] [-abscbr <a-value>] [-absrtvbr <a-value>]
[-absnrtvbr <a-value>] [-absabr <a-value>] [-absubr <a-value>]
```

Syntax Description

-percent	The <i>p-value</i> is the percent of reduction in routing cost that triggers re-routing.
Note	This percent can cause significant changes to network performance. If the percent of change is too small, PNNI concentrates on flooding insignificant changes and so has less time to manage the passing of customer data.
	<ul style="list-style-type: none"> • Range: 5–100 • Default: 30
-absabr	For each of the five service classes, you can specify a particular reduction in the cost for the connections in that service class.
-absrtvbr	<ul style="list-style-type: none"> • Range: 0–4294967295 (or $2^{32}-1$)
-absnrtvbr	<ul style="list-style-type: none"> • Default: 0
-absabr	
-absubr	

Related Commands

cnfrteopt, oprtrte, dsprteoptcnf, dsprteoptstat

Attributes

Log: yes State: active Privilege: GROUP1

Example

Change the re-routing threshold to a 50% reduction in the route cost, and create the following absolute thresholds for the indicated service types:

- CBR: 34567
- ABR: 12000
- UBR: 1000

```
pinnacle.7.PXM> cnfrteoptthld -percent 50 -absabr 34567 -absabr 12000 -absubr 1000
```

cnfsct

Configure Service Class Template—PXM45, PXM1E

The **cnfsct** command lets you write over an existing SCT in the appropriate F:SCT sub-directory with a different *minor* version of the same SCT. Most often, a minor change comes from Cisco Systems and applies to one of the SCT it provides. Less frequently, if a company has developed its own set of SCTs through Cisco WAN Manager, it can modify an SCT and give it a new minor version number. If Cisco provides the minor change, it also delivers the checksum. If you make the minor change through CWM, CWM provides the checksum. This latter course of action might occur if Cisco issues a new major version but you define a need to change certain values in that SCT.



Note

Only Cisco Systems can generate a new major version of a file. It does so when adding a new parameter to a MIB. A minor change occurs when a parameter *value* is changed.



Note

The PXM1E uses only port SCTs.

If you want to broadcast the new version throughout the network, you can use CWM. CWM takes care of putting the modified file in the TEMP directory and verifying its integrity through the checksum then moves the file to the appropriate F:/SCT subdirectory. If you want to register the new version on individual nodes, do the following:

1. Use FTP to transfer the modified file to the C:/SCT/TEMP directory
2. Use the **cnfsct** command to move the modified SCT to the F:/SCT directory.
3. At an appropriate time, use the **setsctver** command on the CLI of a particular card to set up the new version for application to that card.

See the **addset** description for more details on major and minor versions of the SCTs.



Note

The **cnfsct** command does not cause a new SCT to become active on the card type you specify with this command. To activate the new SCT on an individual card, you must reset the standby card in a redundant pair or the active card in a non-redundant configuration.

Syntax

```
cnfsct <cardtype> <scttype> <sctId> <majorversion> <checksum>
```

Syntax Description

<i>cardtype</i>	The <i>cardtype</i> is the card whose SCT you want to make available to the card by installing the SCT in the appropriate directory. 1: AXSM 2: AXSM-E 3: PXM1E 4: FRSM12 Default: None
<i>sctype</i>	The <i>sctype</i> parameter identifies either a port-level or a card-level SCT, as follows: 1: Port-level SCT 2: Card-level SCT (not applicable when this command runs on the PXM1E) Default: None
<i>sctId</i>	The SCT ID refers to a specific service class template. The SCT is either provided by Cisco or created through CWM.
<i>majorversion</i>	Major version number of a file. The minor version is embedded in the file, and the system includes it when it writes over the previous version.
<i>checksum</i>	Each SCT that Cisco provides comes with a checksum that is published in the release notes. When you use CWM to create a new SCT from an existing SCT, CWM generates a checksum for that new SCT. The range is 0–0xffffffff.
<i>description</i>	The description is an ASCII string that provides information about the SCT. It can be from 1 to 132 characters and cannot include spaces. If necessary, you could use underscores, for example: test_modified_version_of_Cisco_SCT_4.

Related Commands

delsct, addset, dspsects, setsetver, addport, cnfport, dspport, cnfcdsct, dspportset, dspcdsct, dspset

Attributes

Log: yes State: active Privilege: GROUP1

Example

For AXSM card types, configure port-level SCT number 4, major release 1. The checksum is ac102f.

```
D1.7.PXM.a > cnfsct 1 1 4 1 ac102f
```

The **cnfsct** command does not cause a new SCT to become active on the card type you specify with this command. To activate the new SCT on an individual card, you must reset the standby card in a redundant pair or the active card in a non-redundant configuration.
Do you want to proceed (Yes/No)? **n**
(command not executed)

cnfserialif

Configure Serial Interface—PXM45, PXM1E

The **cnfserialif** command lets you change the data rate on a serial interface on the PXM UI-S3 back card. The two types of serial ports are the console port and the maintenance port. These ports provide user-access for controlling the switch. The default speed on a serial interface is 9600 bits per second, but higher speed terminals are frequently available.

Each port connects to a different type of terminal implementation. Refer to the switch software configuration guide, for a description of how to use these physical ports for switch control.

Syntax

cnfserialif <port#> <speed>

Syntax Description

<i>port#</i>	Specifies the physical port: <ul style="list-style-type: none"> • 1=maintenance port. • 2=console port.
<i>speed</i>	Specifies a data rate in bits per second. Valid entries are 1200, 2400, 4800, 9600, 19200, and 38400.

Related Commands

delserialif, **dspserialif**

Attributes

Log: no State: active Privilege: SUPER_GP

Example

Configure the console port to have a data rate of 19200 bits per second.

```
node19.8.PXM.a > cnfserialif 2 19200
```


cnfsig

Configure Signaling—PXM45, PXM1E

The **cnfsig** command lets you configure signaling timers for a port whether the port is up or down. The new configuration applies to new incoming calls while existing calls remain intact. In addition to standard timers, this command lets specify the maximum number of crankbacks that PNNI can attempt at a port. You must specify at least one of the optional parameters.

Syntax

```
cnfsig <portid> -t301 t301-timer ] -t303 t303-timer ] -t308 t308-timer ] -t310 t310-timer ]
-t316 t316-timer ] -t317 t317-timer ] -t322 t322-timer ] -t397 t397-timer ] -t398 t398-timer ]
-t399 t399-timer ] -maxcrbk value ]
```

Syntax Description

<i>portid</i>	The format of the PNNI physical port identifier can vary, as follows: <ul style="list-style-type: none"> On a PXM45: <i>slot:subslot.port:subport</i> On a PXM1E for UNI/NNI back card: <i>slot:subslot.port:subport</i>. On the UNI/NNI back card, the subslot is always 2, but the <i>slot</i> depends on the chassis, as follows: <ul style="list-style-type: none"> In an MGX 8850 chassis, <i>slot</i> is always the logical slot 7. In an MGX 8830 chassis, <i>slot</i> is always the logical slot 1. On a PXM1E for a narrowband service module (NBSM): <i>slot.port</i>. For more details, see the section, “PNNI Format,” in Chapter 1, “Introduction.”
-t301	The T301 timer. Range: 150–240 seconds Default: 180
-t303	The T303 timer. Range: 4–8 seconds Default: 4
-t308	The T308 timer. Range: 20–45 seconds Default: 30
-t310	The T310 timer. Range: <ul style="list-style-type: none"> 10–20 seconds for UNI. 30–120 seconds for PNNI. The range you can <i>specify</i> for PNNI is 30–120. If you do not specify a T310 timer value for PNNI, it remains the default of 10 seconds for PNNI. Default: 10
-t316	The <i>t316-timer</i> : Set the T316 timer. Range: 90–300 seconds Default: 90

-t317	The <i>t317-timer</i> : Set the T317 timer. Range: 60–300 seconds Default: 60
-t322	The <i>t322-timer</i> : Set the T322 timer. Range: 4–20 seconds. Default: 4
-t397	The <i>t397-timer</i> : Set the T397 timer. Range: 180–240 seconds Default: 180
-t398	The <i>t398-timer</i> : Set the T398 timer. Range: 4–20 seconds Default: 4
-t399	The <i>t399-timer</i> : Set the T399 timer. Range: 14–28 seconds for UNI 3.0 and 3.1, and 34–124 for UNI 4.0. Default: 14 and 34, respectively.
-maxcrbk	The maximum number of crankback attempts allowed on the port. Range: 0–10 Default: 3

Related Commands

dspsig

Attributes

Log: yes State: active Privilege: SUPER_GP

Example

Configure port 1:1.1:1 to have the maximum crankback count of 5. Check the results with the **dspsig** command. Note the default values in the **dspsig** output.

```
pop20one.7.PXM.a > cnfsig 1:1.1:1 -maxcrbk 5
```

```
pop20one.7.PXM.a > dspsig 1:1.1:1
Signaling Timers for port: 1:1.1:1
```

```
Timer:          Value(secs)
-----          -
t301             180
t303              4
t308             30
t310             10
t316             90
t317             60
t322              4
t397            180
t398              4
t399             14
```

```
Max Crankback: 5
```

cnfsigdiag

Configure Signaling Diagnostic—PXM45, PXM1E

The **cnfsigdiag** command lets you enable signaling diagnostics and create filters for these diagnostics,



Note

The **cnfsigdiag** command primarily is intended for use by Cisco engineers. The more suitable command for end-users is the **dbgsig** command.

The ATM signaling diagnostics are tools for troubleshooting call failures in the network and should not be enabled while the switch is operating.

Syntax

```
cnfsigdiag {[enable | disable | index]} [-cldaddr nsap-address] [-clgaddr nsap-address]
[-cldaddrmask atm-address-mask] [-clgaddrmask atm-address-mask] [-casttype {all | p2p | p2mp}]
[-clrcause clear-cause-code] [-connctgy {all | svc | svp | swvc | swvp}] [-inport portid]
[-outport portid] [-maxrec max-num-records] [-scope {all | ext | int}]
[-servctgy {all | cbr | rtvbr | nrtvbr |ubr | abr}] [-status {active | inactive}]
```

Syntax Description

enable, disable, or index	<p>Enable or disable signaling diagnostics or configure an index.</p> <p>Specify the diagnostics index number for the filter table and enter the diagnostics configuration mode. If you do not specify an index, the enable or disable condition globally applies to all signaling diagnostics.</p> <p>Range for <i>index</i>: 1–50.</p> <p>Default: disable</p>
-cldaddr	<p>The <i>nsap-address</i> is the filter for ATM signaling call failures against this called address.</p> <p>Default: NULL</p>
-clgaddr	<p>The <i>nsap-address</i> is the filter for ATM signaling call failures against this calling address.</p> <p>Default: NULL</p>
-cldaddrmask	<p>The the <i>atm-address-mask</i>: Address mask for identifying valid bits of the called NSAP address field (ff.ff.ff, for example). To match this selection criterion, a failed connect setup must have a called party address value equal to the configured called party address for all bits that are 1 in the specified mask.</p> <p>Default: NULL. NULL means the rejected call matches the filter criteria for any called address in the rejected call.</p>
-clgaddrmask	<p>The <i>atm-address-mask</i>: Address mask for identifying valid bits of the calling NSAP address field. (ff.ff.ff, for example). To match this selection criteria, a failed connect setup must have a calling party address value equal to the configured calling party address for all bits that are 1 in the specified mask.</p> <p>Default: NULL means the call matches the filter criteria for any calling address in the rejected call.</p>

-casttype	Filtering by connection type. The types are point-to-point (p2p), point-to-multipoint (p2mp —currently not supported), or both (all). Default: all
-clrcause	The <i>clear-cause-code</i> : Filters ATM signaling call failures by the release cause code (a decimal number) as specified in the ATM Forum UNI 3.1 specification. Default: 0, meaning the cause code is not considered during filtering.
-connctgy	Filters ATM signaling call failures by virtual circuit category (SPVC, SPVP, SVC, SVP, or all of these circuit categories). Default: all
-inport	The <i>portid</i> : filters ATM signaling call failures based on the incoming port of the call. Default: 0, meaning the incoming port is not considered during filtering.
-outport	The <i>portid</i> : filters ATM signaling call failures based on the outgoing port of the call. Default: 0, meaning the outgoing port is not considered during filtering.
-maxrec	The <i>max-num-records</i> : the maximum number of records collected for a particular signaling diagnostics filter table entry. When the maximum value is reached, the older records are deleted. If this field is set to -1, the records are not overwritten. Setting this field to -1 increases memory usage for call failure records and can lead to shortages of available system memory. Range: -1 through 214783647 Default: 20
-scope	The filtering scope choices are within the switch (int), on other switches (ext), or both (all). Default: all
-servctgy	Filters ATM signaling call failures by service category (service type): valid entries are: all (for all service types), cbr, rtvbr, nrtvbr, ubr, or abr. Default: all
-status	The status of the entry for the signaling diagnostics filter table. Type active to begin filtering failed connections or inactive to stop filtering failed connections. The inactive specification causes the node to delete all the records associated with the filter entry. Default: inactive

Related Commands

delsigdiag, dspsigdiag, dspsigstats, clrstats

Attributes

Log: yes State: active Privilege: SERVICE_GP

cnfsnmp

Configure SNMP Strings—PXM45, PXM1E

Configure the SNMP strings. The three strings are *community*, *contact*, and *system location*. You can configure only one of these strings with a single execution of **cnfsnmp**.

Syntax

cnfsnmp community *string* [**ro** | **rw**>]

cnfsnmp contact *string*

cnfsnmp location *string*



Note

You can specify only one keyword per iteration of the **cnfsnmp** command. (The keywords are **community**, **contact**, and **location**.) Also, unlike most keyword-driven parameters, no dash character (-) precedes the keyword.

Syntax Description

community	<p>A community string provides an authentication mechanism to access MIB objects by using SNMPv1 protocol. One read-only (ro) access string and one read-write (rw) access string are supported. You can specify your own rw and ro strings (up to 32 characters each) or keep one or both of the reserved strings. You cannot set ro and rw to the same string value. If you specify <i>no</i> community strings, “private rw” is assumed. Note that <i>password</i> cannot contain a blank space, an “@,” or a quote (“) character. Reserved strings are:</p> <ul style="list-style-type: none"> • Community rw string: “private” • Community ro: “public” <p>The <i>string</i> acts like a password to permit access to the SNMP Protocol. Further, the access privilege of either ro or rw determines allowable operations on MIB Objects. The setting can be either “ro” for read-only or “rw” for read-write. With read-only access, a management station is allowed only to <i>retrieve</i> information. With read-write access, authorized management stations are able to <i>retrieve and modify</i> MIB objects.</p>
contact	<p>Specifies the contact string for sysContact MIB object in MIB-II. The contact is a printable string of 0–255 characters. The string in this case is text that identifies the contact. For example, the contact could be an administrator’s email address. Spaces are allowed between character strings. You can reset the contact string to no text by entering the cnfsnmp command and contact keyword then press Enter or Return with no other text. See Example.</p> <p>Default: no text</p>
location	<p>Specifies the location of the system. The location is a printable string of 0–255 characters. The system location string is used for sysLocation MIB object in MIB-II. Spaces are allowed between character strings. You can reset the location string to no text by entering the cnfsnmp command and location keyword then press Enter or Return with no other text. See Example.</p> <p>Default: no text</p>

Attributes

Log: yes State: active Privilege: SUPER_GP

Related Commands

dspsnmp

Example

Do the following:

1. Configure the rw string as “toplevel.”
2. Enter the contact “davids@be.com”
3. Enter the location “Bldg J, room 619.”
4. Display the SNMP settings.
5. Reset the contact and location strings to no text.
6. Display the SNMP settings.

```
M8850_NY.7.PXM.a > cnfsnmp community toplevel rw
```

```
M8850_NY.7.PXM.a > cnfsnmp contact davids@be.com
```

```
M8850_NY.7.PXM.a > cnfsnmp location Bldg J, room 619
```

```
M8850_NY.7.PXM.a > dspsnmp
```

```
M8850_NY                                      System Rev: 02.01      Nov. 29, 2001 12:49:03 PST
MGX8850                                                              Node Alarm: none
```

```
Community (rw):                      toplevel
Community (ro):                      public
System Location:                      Bldg J, room 619
System Contact:                      davids@be.com
```

```
M8850_NY.7.PXM.a > cnfsnmp location
```

```
M8850_NY.7.PXM.a > cnfsnmp contact
```

```
M8850_NY.7.PXM.a > dspsnmp
```

```
M8850_NY                                      System Rev: 02.01      Nov. 29, 2001 12:50:40 PST
MGX8850                                                              Node Alarm: none
```

```
Community (rw):                      toplevel
Community (ro):                      public
System Location:
System Contact:
```

```
M8850_NY.7.PXM.a >
```

cnfsntp

Configure Simple Network Time Protocol—PXM45, PXM1E

The node-level **cnfsntp** command lets you configure the current switch to be a server, a client, or both of these for time-of-day (TOD) synchronization through the support of the Simple Network Time Protocol (SNTP). You can also configure certain timers for a client or a stratum level for a server.



Note

The common server for TOD synchronization is a workstation with CWM, but the **cnfsntp** command lets you specify a switch to operate as a primary server or a secondary server. The command also lets you specify certain behaviors of the client switch (regardless of the type of server).

To register a workstation or a switch as a server for a client-node, use the **addsntprrmtsvr** command on the client. See the Example section for this part of TOD synchronization.

SNTP propagates a standardized TOD from the current SNTP server to all the switches you register as SNTP clients. (If you do not specify a switch as a client, it maintains its own TOD—unsynchronized with other switches.) An SNTP timestamp is a 64-bit, unsigned, fixed-point number and is measured in seconds relative to 0-hour on 1 January, 1900. The integer part is in the first 32 bits, and the fraction part is in the last 32 bits. A timestamp is useful for such applications as event logging or billing applications.

Note that SNTP can also be used to sync the date and time of any RPM card in the node to that node's date and time. (The other cards automatically receive their date/time from the PXM during bootup.)

An SNTP *client* can be any of the following platforms:

- MGX 8950, MGX 8850 (Release 4), and MGX 8830 switches

An SNTP client has the following characteristics:

- You can configure a client to synchronize with a primary NTP/SNTP server and up to three secondary NTP/SNTP servers.
- The client periodically requests TOD from the current server. The polling timer is configurable.
- An SNTP client runs over UDP/IP using traditional IP or the IP connectivity feature provided by the Cisco MGX Release 4 platform.
- If a client detects the current server not reachable, it switches to next available server.
- The client rolls back to the primary server when it can be reached.
- Only unicast mode is supported.

An SNTP *server* can be any of the following platforms:

- A CWM workstation
- BPX 8600-series/SES, MGX 8950, MGX 8850 (Release 4), and MGX 8830 switches

An SNTP server has the following characteristics:

- The server listens to the SNTP UDP port, processes requests, and provides the TOD information.
- The server provides loopback detection to prevent a synchronization loop.
- A server can support multiple clients.
- Only unicast mode is supported.

The primary server should have the highest quality and most accurate TOD source. This designation should correspond to the highest quality clock source (see the section, "[The Role of the Stratum](#)"). A secondary server often has a lower stratum than the primary server.

A *secondary* server is a backup in case the primary server cannot provide TOD updates at specified intervals. If a client cannot reach the primary server after three polling attempts, the client switches to a secondary server. After the primary server becomes reliable, the client switches back to the primary server (see Syntax Description for the configurable criteria).

In the area of control card redundancy, note the following:

- SNTP runs on the active controller only.
- After switchover, the newly active card starts SNTP for client or server based on the redundant database configuration.

The Role of the Stratum

The purpose of the stratum hierarchy is to prevent time-loops. To fulfill this purpose, assign a stratum to each server-node but not the client. A client reads the stratum of the server then adds a 1 to that stratum, thus forcing the client to have a lower stratum (and prevent a time-loop).

The stratum in the SNTP context applies to the TOD source whether that source is a workstation or a switch. Although SNTP does not interact with the switch's clock manager and the stratum does not directly refer to the nodal clock source, the server's stratum should reflect the quality of the clock source. For example, if the server node is the master clock source with a stratum 1 BITS device, this switch is a likely candidate to be the primary SNTP server and be configured through the **cnfsntp** command to have a stratum of 1.

The range for stratum values is 0–15. The highest stratum is 1, and the lowest is 15. A default stratum of 0 means that no stratum configuration exists.

A Typical SNTP Configuration

This section illustrates a small network that has the SNTP configuration in [Figure 2-9](#). The figure also shows a sequence in which a client switches to a secondary server then switches back to the primary server. The configuration tasks that result in this setup appear in the Example section.



Note

When you enable a switch as a secondary server, you must also configure it as a client in a separate iteration of the **cnfsntp** command (presuming you want the client to synchronize with the primary server). In this case, the stratum assignment applies to the server role of the switch. (See the Example section of this description.)

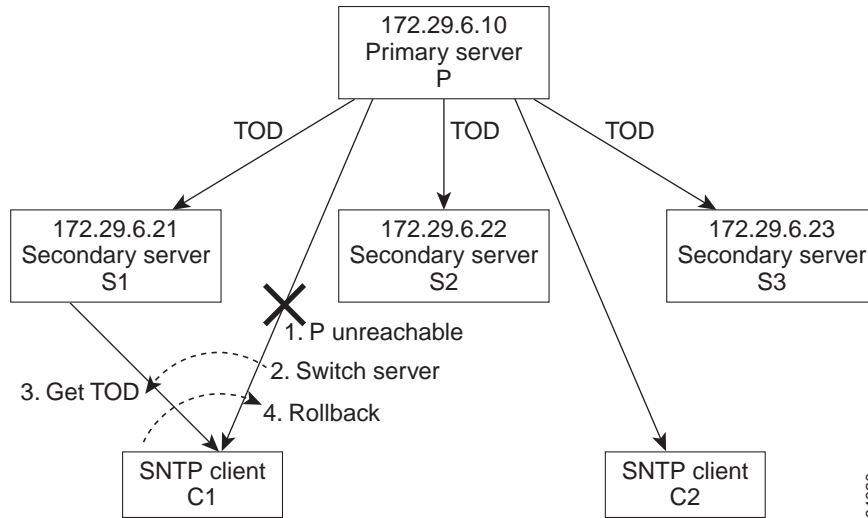
The setup in [Figure 2-9](#) consists of the following configuration:

- The primary server is “P” and provides TOD to S1, S2, S3, C1, and C2.
- The secondary servers are S1, S2, and S3. During normal operation—while P is the *current* server—these switches operate as clients.
- The switches configured solely as clients are C1 and C2.

In [Figure 2-9](#), the following events take place:

1. The link between P and C1 goes down.
2. C1 switches to S1 for TOD synchronization.
3. S1 provides TOD to C1 while C1 polls P at intervals specified by the *rollback* timer.
4. After three TOD responses from P, C1 switches back to P.

Figure 2-9 An SNTP Client and Server Configuration



SNTP Timers

The SNTP timers are as follows:

- The *polling* timer counts the seconds until the client sends a request for TOD from the current server.
- The *waiting* timer counts down the seconds to wait for the TOD response. If the waiting timer expires three times without a TOD from the server, the client sends requests to a secondary server.
- The *rollback* timer applies only when the SNTP client has switched to a secondary server. This timer specifies a time the client waits to poll the primary server to determine if the primary has become available. After the client receives three TOD responses from the primary server, the client returns to requesting TOD updates from the primary server. To avoid large traffic volume, the rollback timer should have greater value than the polling timer.

The following example illustrates the time for a client to switch over to a secondary server if the default timer values are used:

$(3 \text{ retries} + 1) * \text{polling time} = 256 \text{ seconds (typical)}$ for the client to switch from the primary to secondary server

An additional value that SNTP monitors is the *round trip delay*. The round trip delay is the total time from the request to the arrival of the TOD response and has a fixed maximum of 400 milliseconds. The fixed value of 400 milliseconds means that SNTP ensures a maximum margin of error of 200 milliseconds in the TOD. If the round trip takes more than 400 milliseconds, the client discards the TOD. If the 400-millisecond delay is exceeded three times, the client switches to a secondary server.

Operational Characteristics

This section describes characteristics you should consider before using the **cnfsntp** command.

- The maximum number of clients should be less than 200. This limit is based on the default UDP queue size.
- SNTP does not synchronize the timezone throughout the network. You should specify the timezone as needed at each node by using the **cnfmzn** (and possibly **cnfmzngmt**) command.
- To identify the servers to the clients, run the **addsntprrmtsvr** command on each client. Run this command once to identify the primary server, and run it as needed once for each secondary server.

You must use separate iterations of the **cnfsntp** command to configure different parts of SNTP functionality. The following example illustrates this requirement in the case where a switch is configured as a client and as a secondary server. This local node is configured to be a server with a stratum of 2. The example also shows that the node with a disk IP address of 172.29.51.87 is configured to be the primary remote server, so the node is configured to operate as a client, so the switch really is intended to be a secondary server. Note that, without the step that configures the client to be enabled on the secondary server, the primary server would not synchronize with the secondary server.

1. **cnfsntp -stratum 2**
2. **cnfsntp -server on**
3. **cnfsntp -client on**
4. **addsntprrmtsvr 172.29.51.87 -primary yes**

Syntax

```
cnfsntp {[-polling insecond] [-waiting insecond] [-rb insecond] [-client on | off] [-server on | off] [-stratum 0-15]}
```

Syntax Description

The **-polling**, **-waiting**, and **-rb** parameters apply only to the client. Also, no mandatory parameters exist for this command. Therefore, you could, for example, modify a polling timer on a client node by entering the command, the keyword, and the timer value.

-polling	The polling timer specifies the number of seconds to wait before the client switch requests a TOD update from the server. Range: 64–1024 seconds Default: 64 seconds
-waiting	The waiting timer is the number of seconds the client waits for a TOD response. If this timer expires three times before a TOD response arrives, the client switches to a secondary server. Range: 1–10 seconds Default 5 seconds
-rb	If a client switches to a secondary server, it starts the rollback timer and polls the primary server after a period equal to the value of the rollback timer. After the client receives three legitimate TOD responses from the primary server, it switches back to the primary. Range: 64–2048 seconds Default: 1024 seconds
-client	This parameter lets you specify whether the current switch is an SNTP client. Type “on” to make it a client or “off” to remove it from the list of clients. Default: off

-server	This parameter lets you specify whether the current switch is an SNTP server. Type “on” to make it a server or “off” to remove it from the list of servers. Default: off
-stratum	The range for stratum is 0–15. The highest stratum level is 1, and the lowest is 15. A stratum of 0 on a client node means it can take its TOD from any server. The application of the stratum varies on client and server, as follows: <ul style="list-style-type: none"> • On a server, the stratum is that of the local node. • On a client, the stratum is the lowest level stratum of a server that this client can use for TOD. Default: 0 (no stratum configured)

Related Commands

addsntprmtsvr, cnfsntprmtsvr, delsntprmtsvr, dpsntp, dpsntprmtsvr, dbgsntp, clrntpstats, dpsntpstats, dpsntp-dbg

Attributes

Log: yes State: active Privilege: GROUP1

Example 1

The first example consists of the steps required to configure SNTP as shown in [Figure 2-9](#).

Determining the Primary Server

Specify that switch P has a stratum of 1 and is a server. Leave the timers with default values.

```
> cnfsntp -stratum 1
> cnfsntp -server on
```

Configuring S1

S1 configuration consists of two tasks: client configuration and server configuration.

Three CLI entries create S1’s client role: use the **addsntprmtsvr** command to identify P as the primary server for S1, then identify S2 as S1’s secondary server. Enable S1 as a client with default timer values.

```
> addsntprmtsvr 172.29.6.10 -primary yes
> addsntprmtsvr 172.29.6.22 -primary no
> cnfsntp -client on
```

For S1’s role as a secondary server, enter the following sequence on the CLI of S1:

```
> cnfsntp -server on
```

Configuring S2

S2 configuration consists of two tasks: client configuration and server configuration.

Three CLI entries create S2’s client role: use the **addsntprmtsvr** command to identify P as the primary server for S2. Identify S3 as the secondary server for S2. Enable S2 as a client with default timer values.

```
> addsntprmtsvr 172.29.6.10 -primary yes
> addsntprmtsvr 172.29.6.23 -primary no
> cnfsntp -client on
```

For S2's role as a secondary server, enter the following sequence on the CLI of S2:

```
> cnfsntp -server on
```

Configuring S3

S3 configuration consists of two tasks: client configuration and server configuration.

Two steps exist for S3's client role: use the **addsntprmtsvr** command to identify P as the primary server for S2. Identify S3 as the secondary server for S2. Enable S3 as a client with default timer values.

```
> addsntprmtsvr 172.29.6.10 -primary yes
> addsntprmtsvr 172.29.6.23 -primary no
> cnfsntp -client on
```

For S3's role as a secondary server, enter the following sequence on the CLI of S3:

```
> cnfsntp -server on
```

Configuring C1

Three CLI entries create C1's client role: use the **addsntprmtsvr** command to identify P as the primary server for C1, then identify S1 as C1's secondary server. Enable C1 as a client with default timer values.

```
> addsntprmtsvr 172.29.6.10 -primary yes
> addsntprmtsvr 172.29.6.21 -primary no
> cnfsntp -client on
```

Configuring C2

Three CLI entries create C2's client role: use the **addsntprmtsvr** command to identify P as the primary server for C1, then identify S3 as C2's secondary server. Enable C2 as a client with default timer values.

```
> addsntprmtsvr 172.29.6.10 -primary yes
> addsntprmtsvr 172.29.6.23 -primary no
> cnfsntp -client on
```

Example 2

This example uses two switches. Initially, neither switch has an SNTP configuration. The intended primary server is M8850_NY because it has the master clock source for the network. The SNTP client is PXM1E_SJ.

Configure M8850_NY to be the primary SNTP server with a stratum of 1, then display the SNTP state. Note that the server function has been enabled, and both the "default" (configured) stratum and current stratum are now 1.

```
M8850_NY.7.PXM.a > cnfsntp -stratum 1
```

```
M8850_NY.7.PXM.a > cnfsntp -server on
```

```
M8850_NY.7.PXM.a > dspsntp
```

```
client: no
server: yes
```

```
polling: 64
waiting: 5
rollback: 1024
stratum(default): 1
stratum(current): 1
sync: no
```

Display the date and time on the primary server, then display the date and time on PXM1E_SJ. The date on the latter node is 6 December, 1999, so it appears not to have been configured.

```
M8850_NY.7.PXM.a > dspdate
Dec 27 2002 02:41:09 GMT
```

```
PXM1E_SJ.7.PXM.a > dspdate
Dec 06 1999 02:22:53 GMT
```

Configure PXM1E_SJ as a client with a stratum of 5, then display the configuration. Note that both the default and current stratum are 5, and the “sync” status is “no.”

```
PXM1E_SJ.7.PXM.a > cnfsntp -client on
```

```
PXM1E_SJ.7.PXM.a > cnfsntp -stratum 5
```

```
PXM1E_SJ.7.PXM.a > dspntp
```

```
client: yes
server: no
```

```
polling: 64
waiting: 5
rollback: 1024
stratum(default): 5
stratum(current): 5
sync: no
```

On the client PXM1E_SJ, register M8850_NY as the primary server by using the **addsntprmtsvr** command. On the client, display its SNTP status and note that it has synchronized to the server: the “sync” is now “yes.” Note also that the *current* stratum is 2, which is the stratum of the primary server plus 1. The default stratum remains as you configured it.

```
PXM1E_SJ.7.PXM.a > addsntprmtsvr 172.29.52.56 -primary yes
```

```
PXM1E_SJ.7.PXM.a > dspntp
```

```
client: yes
server: no
```

```
polling: 64
waiting: 5
rollback: 1024
stratum(default): 5
stratum(current): 2
sync: yes
```

Display the date and time on the client and note its proximity to those of the primary server.

```
PXM1E_SJ.7.PXM.a > dspdate
Dec 27 2002 02:41:07 GMT
```

```
M8850_NY.7.PXM.a > dspdate
Dec 27 2002 02:41:09 GM
```

cnfsntpmtsvr

Configure Simple Network Time Protocol Remote Server—PXM45, PXM1E

The **cnfsntpmtsvr** command lets you modify certain parameters of a switch that was configured to be an SNTP server through the **addsntpmtsvr** command. The configurable parameters are as follows:

- The version of SNTP
- The role of the server as primary

See the **cnfsntp** and **addsntpmtsvr** descriptions for details of SNTP operation in the current release.

Syntax

```
cnfsntpmtsvr {server IP address} [-version {version}] [-primary {yes | no}]
```

Syntax Description

<i>IP address</i>	The IP address of the switch you want to be an SNTP server.
-version	The SNTP version is either 3 or 4. Default: 3
-primary	This parameter lets you identify the switch as the primary SNTP server. Type -primary yes to make the primary server. To change the remote switch to a secondary server, type -primary no . Default: no

Related Commands

addsntpmtsvr, **cnfsntp**, **delsntpmtsvr**, **dpsntp**, **dpsntpmtsvr**, **dbgsntp**, **clrsntpstats**, **dpsntpstats**, **dpsntp-dbg**

Attributes

Log: yes State: active Privilege: GROUP1

Example

cnfspvcprfx

Configure SPVC Prefix—PXM45, PXM1E

The **cnfspvcprfx** command lets you configure a node-level ATM prefix for SPVCs. For the node to support SPVCs, it must have a 13-byte SPVC prefix for the entire node. No SPVCs can exist on the switch until it has an SPVC prefix. Likewise, to change this prefix, no SPVCs can exist on the node.

Prerequisites

Setting up a node and a network requires advance planning for the PNNI node addressing scheme. For basic guidance on the topic of address planning, see related material in the software configuration guide.

Cisco provides a default SPVC prefix that is the same as the Cisco-supplied ATM address prefix. Each of these default prefixes contains an International Code Designator (ICD) that is unique to Cisco Systems. Therefore, Cisco Systems recommends that you change the ICD identifier for both the ATM address prefix and the SPVC prefix if the node is planned for operation in a public ATM network. If the node operates in a private ATM network, it can keep the default ATM and SPVC address prefixes.

The following list shows the order of prerequisite commands and the **cnfspvcprfx** command.

1. **addcontroller**
2. **cnfpnni-intf**
3. **cnfspvcprfx**
4. **dspsvcprfx**

Syntax

```
cnfspvcprfx -prfx <prefix | "default">
```

Syntax Description

-prfx	The prefix is either a 13-byte value you enter or the character string "default" (to keep the Cisco factory default).
--------------	---

Related Commands

dspsvcprfx

Attributes

Log: yes State: active Privilege: SUPER_GP

Example

First display the current SPVC prefix. The ICD field shows the prefix is the default from Cisco (0091). Configure the SPVC prefix 47.007778000000aa2109ff214.

```
pop20one.7.PXM.a > dspsvcprfx
SPVC Node Prefix: 47.00918100000000107b65f33c
pop20one.7.PXM.a > cnfspvcprfx 47.007778000000aa2109ff214
```

cnfspvcres

Configure SPVC Reserve—PXM45, PXM1E

The **cnfspvcres** command lets you reserve a number of connections for SPVCs or SPVPs. The maximum number of reserved connections cannot exceed the maximum number of logical connections on the switch. The remainder of logical connection numbers goes to SVCs or SVPs.

Syntax

cnfspvcres *<num-conn-reserved-for-svc>*

Syntax Description

<i>num-conn-reserved-for-svc</i>	The number of connections to reserve for SPVCs has a range of 0 to the maximum number of connections that the switch can support. Default: 0
----------------------------------	---

Related Commands

dspsvcres

Attributes

Log: yes State: active Privilege: GROUP1

Example

Reserve 1000 connections to be SPVCs, then display the number of connections reserved for SPVCs.

```
PhattyEmre.2.PXM.a > cnfspvcres 1000

PhattyEmre.2.PXM.a > dspsvcres
reserved terminating spvc number is 1000
```


cnfsrcmclsrc

Configure Service Resource Module Clock Source—PXM45, PXM1E

The **cnfsrcmclsrc** command lets you configure a clock source for the transmit direction of a line on an SRME. If the SRME is using the “no back card” option, this command is meaningless. The choices are local timing and loop timing.



Note

On an SRM-3T3, the **cnfsrcmclsrc** command appears but has *no effect*. Instead, the SRM-3T3 has an internal oscillator for its synchronization and line transmit clock.

For the local transmit clock source, the SRM has its own clock generation circuitry. For loop timing, the SRM takes the receive clock off the receive line and redirects as the transmit clock signal. You can use either the **dspsrmclsrc** or **dspln** command to see the transmit clock source.

Syntax

```
cnfsrcmclsrc <-lineType> <logicalslot>.<line> <-clk 1 | 2>
```

Syntax Description

<i>-lineType</i>	The line type depends on the type of SRM: <ul style="list-style-type: none"> For SRM-3T3/C, type -ds3. For SRME, type -sonet.
<i>logicalslot</i>	The logical slot applies to either physical slot: for example, if the only SRM resides in slot 16, you type its logical equivalent: 15. Also, the logical slot depends on the chassis: <ul style="list-style-type: none"> In an MGX 8850 chassis, the logical slot is either 15 or 31. In an MGX 8830 chassis, the logical slot is 7.
<i>line</i>	The line depends on the type of SRM, as follows: <ul style="list-style-type: none"> On an SRME <i>line</i> is 1. On an SRM-3T3/C, <i>line</i> has a range of 1–3.
-clk	Type a 1 for loop timing or a 2 for local timing. Default: 2 (local timing)

Related Commands

dspsrmclsrc

Attributes

Log: no State: active Privilege: ANYUSER

Example

Display the SRM clock source for 15.1. Change to loop timing as needed. At first, the line is down.

```
PXM1E_SJ.7.PXM.a > dspsrmclksrc -sonet 15.1
  LineNum:          15.1
  LineXmtClockSource: localTiming

PXM1E_SJ.7.PXM.a > cnfsrcmclksrc -sonet 15.1 -clk 1
ERROR: The line is disabled

PXM1E_SJ.7.PXM.a > upln 15.1

PXM1E_SJ.7.PXM.a > cnfsrcmclksrc -sonet 15.1 -clk 1

PXM1E_SJ.7.PXM.a > dspsrmclksrc -sonet 15.1
  LineNum:          15.1
  LineXmtClockSource: loopTiming
```

cnfsscop

Configure SSCOP—PXM45, PXM1E

The **cnfsscop** command lets you configure service-specific connection-oriented protocol (SSCOP) on a port. You can use this command regardless of the state of the port. You must specify at least one of the optional parameters.

Syntax

```
cnfsscop <portid> [-polltmr {poll-timer | 0}] [-keepalivetmr {keepalive-timer | 0}]
[-idletmr {idle-timer | 0}] [-cctmr {cc-timer | 0}] [-norsptmr {noresponse-timer | 0}]
[-t309tmr {t309-timer | 0}] [-maxcc {retries | 0}] [-sndwnd {send-window-packets | 0}]
[-rcvwnd {recv-window-packets | 0}]
```

Syntax Description

<i>portid</i>	The format of the PNNI physical port identifier can vary, as follows: <ul style="list-style-type: none"> On a PXM45: <i>slot:subslot.port:subport</i> On a PXM1E for UNI/NNI back card: <i>slot:subslot.port:subport</i>. On the UNI/NNI back card, the subslot is always 2, but the <i>slot</i> depends on the chassis, as follows: <ul style="list-style-type: none"> In an MGX 8850 chassis, <i>slot</i> is always the logical slot 7. In an MGX 8830 chassis, <i>slot</i> is always the logical slot 1. On a PXM1E for a narrowband service module (NBSM): <i>slot.port</i>. For more details, see the section, “PNNI Format,” in Chapter 1, “Introduction.”
-polltmr	Number of seconds to send POLL PDUs in the active phase. Range: 1–5 seconds Default: 1 second Reset: A 0 forces a return to the default value.
-keepalivetmr	Number of seconds to send POLL PDUs in the transient phase. Range: 1–10 seconds Default: 5 seconds Reset: A 0 forces a return to the default value.
-idletmr	Number of seconds to send POLL PDUs in the idle phase. Range: 1–20 seconds Default: 10 seconds Reset: A 0 forces a return to the default value.

-cctmr	<p>The <i>connection control timer</i> times the retransmission of a variety of PDUs that apply during the connection control phase. This parameter specifies the number of seconds the port waits for an ACK before it retransmits a particular PDU. The timer value is the same for all these PDUs. The control phase PDUs are as follows:</p> <ul style="list-style-type: none"> • BEGIN: The BEGIN PDU establishes a connection between two peer SSCOP entities. • END: The END PDU releases the control connection between the two communicating SSCOP peer entities. • RS: The RESYNC PDU acts as a conventional connection-oriented reset. This PDU re-synchronizes buffers and receive and transmit state. • ER: The ERROR RECOVERY PDU is responsible for recovery from errors that occurred during the connection control phase. <p>Note The <i>number</i> of retransmissions is controlled by the -maxcc parameter.</p> <p>Range: 1–5 seconds Default: 1 second Reset: A 0 forces a return to the default value.</p>
-norsptmr	<p>Number of seconds after which at least one STAT PDU must be received (for the No Response timer).</p> <p>Range: 1–45 seconds Default: 30 seconds Reset: A 0 forces a return to the default value.</p>
-t309tmr	<p>Number of seconds before SAAL reconnects after disconnection.</p> <p>Range: 1 –15 seconds Default: 10 seconds</p>
-maxcc	<p>The <i>retries</i> is the maximum number of times the port can resend one of the PDUs that goes out during the connection control phase. See the cctmr description for information on these PDUs.</p> <p>Range: 1–15 Default: 10 Reset: A 0 forces a return to the default value.</p>
-sndwnd	<p>The value for <i>send-window-packets</i> is the number of packets the port can send before it must receive an acknowledgment from an ATM switch. In practical terms, this parameter represents the size of the send window.</p> <p>Range: 1–127 Default: 30 Reset: A 0 forces a return to the default value</p>
-rcvwnd	<p>The value for <i>receive-window-packets</i> is the number of packets the port can receive before it must send an acknowledgment to an ATM switch. In practical terms, this parameter represents the size of the receive window</p> <p>Range: 1–127 Default: 30 Reset: A 0 forces a return to the default value</p>

Related Commands

disablesscop, dspsscop, dspsscopstats, cnfnpportloscallrel

Attributes

Log: yes State: active Privilege: GROUP1

Example

Set the no-response timer to 40 seconds, check the results. Change it back to the default by specifying 0.

```
M8850_LA.8.PXM.a > cnfsscop 12:1.1:1 -norsptmr 40
M8850_LA.8.PXM.a > dspsscop 12:1.1:1

SSCOP details for interface: 12:1.1:1

    Current State = enabled, Current Link State = Established State,
SSCOP version = 3.1
    Send Sequence Number: Current = 69893, Maximum = 69923
    Send Sequence Number Acked = 69893
    Rcv Sequence Number: Lower Edge = 69893, Upper Edge = 69893, Max = 69923
    Poll Sequence Number = 329199, Poll Ack Sequence Number = 329199
    Vt(Pd) = 0    Vt(Sq) = 1
    Timer_IDLE = 10 - Active
    Timer_CC = 1 - Inactive
    Timer_POLL = 1 - Inactive
    Timer_KEEPA L I V E = 5 - Inactive
    Timer_NO-RESPONSE = 40 - Inactive
    Timer_T309 = 10 - Inactive
    Max CC = 10
    Send Window = 30
    Recv Window = 30
    Current Retry Count = 0, Maximum Retry Count = 10
    AckQ count = 0, RcvQ count = 0, TxQ count = 0
    AckQ HWM = 2, RcvQ HWM = 0, TxQ HWM = 1

Type <CR> to continue, Q<CR> to stop: q

M8850_LA.8.PXM.a > cnfsscop 12:1.1:1 -norsptmr 0

M8850_LA.8.PXM.a > dspsscop 12:1.1:1

SSCOP details for interface: 12:1.1:1

    Current State = enabled, Current Link State = Established State ,
SSCOP version = 3.1
    Send Sequence Number: Current = 69902, Maximum = 69932
    Send Sequence Number Acked = 69902
    Rcv Sequence Number: Lower Edge = 69902, Upper Edge = 69902, Max = 69932
    Poll Sequence Number = 329225, Poll Ack Sequence Number = 329225
    Vt(Pd) = 0    Vt(Sq) = 1
    Timer_IDLE = 10 - Active
    Timer_CC = 1 - Inactive
    Timer_POLL = 1 - Inactive
    Timer_KEEPA L I V E = 5 - Inactive
    Timer_NO-RESPONSE = 30 - Inactive
    Timer_T309 = 10 - Inactive
    Max CC = 10
    Send Window = 30
    Recv Window = 30
    Current Retry Count = 0, Maximum Retry Count = 10
    AckQ count = 0, RcvQ count = 0, TxQ count = 0
    AckQ HWM = 2, RcvQ HWM = 0, TxQ HWM = 1

Type <CR> to continue, Q<CR> to stop: q
```

cnfstatsmgr

Configure Statistics Manager—PXM45, PXM1E

The **cnfstatsmgr** command lets you configure multiple statistics managers. The configuration consists of the IP address of the workstation and an index that assigns a place in the hierarchy of statistics managers.

Syntax

cnfstatsmgr <index> <IP Address>

Syntax Description



Note

In the current release, index options 1 through 3 are not applicable to the present switch architecture. Use index option 4 to specify the IP address of the statistics master for the switch, which defines the IP address of the workstation authorized to enable or disable statistics on the switch.

<i>index</i>	The <i>index</i> is a number that points to the place of a statistics manager in the hierarchy of managers. <ul style="list-style-type: none"> • 1: Primary statistics manager • 2: Secondary statistics manager • 3: Tertiary statistics manager • 4: Master statistics manager
<i>IP Address</i>	The IP address of a workstation that serves as a statistics manager.

Related Commands

dspstatsmgr

Attributes

Log: no

State: active

Privilege: SERVICE_GP

Example

cnfsvcoverride

Configure SVC Override—PXM45, PXM1E

The **cnfsvcoverride** command lets you specify three node-level settings that relate to overriding an SVC's or SVP's ownership of a VPI and VCI (in the case of an SVC). The possibilities are as follows:

- An SPVC can override an SVCC.
- An SPVC can override an SVP.
- An SPVP can override an SVP.

This override capability exists to support single-ended SPVCs and SPVPs. In turn, the support for single-ended connections exists to support interoperability between the *current series* of PNNI routing switches and other types of ATM WAN switches. The current series of PNNI routing switches consists of Release 4 Cisco MGX 8950, MGX 8850, and MGX 8830 switches as well as the Service Expansion Shelf (SES). The SPVC/SPVP interoperability pertains to other Cisco ATM WAN switches as well as ATM WAN switches from other manufacturers. Refer to the section, "[Interoperability With Other Switches](#)," in the **addcon** description for details on SPVC interoperability and a definition of a single-ended connection.

Syntax

```
cnfsvcoverride [-spvcoverridesvc {enable | disable}] [-spvcoverridesvp {enable | disable}]
[-spvpoverridesvp {enable | disable}]
```

Syntax Description

-spvcoverridesvc	If you enable this option, the SPVC overrides the SVC on the same port and uses the SVC's VPI and VCI. Type the entire word "enable" or "disable." Default: disable
-spvcoverridesvp	If you enable this option, the SPVC overrides the SVP on the same port and uses the SVP's VPI. Type the entire word "enable" or "disable." Default: disable
-spvpoverridesvp	If you enable this option, the SPVP overrides the SVP on the same port and uses the SVP's VPI. Type the entire word "enable" or "disable." Default: disable

Related Commands

dspsvcoverride

Attributes

Log: yes State: active Privilege: SUPER_GP

Example

Enable SPVC override of SVCs.

```
PXM1E_SJ.7.PXM.a > cnfsvcoverride -spvcoverridesvc enable
```

```
PXM1E_SJ.7.PXM.a > dspsvcoverride  
spvcoverridesvc:           Enabled  
spvcoverridesvp:           Disabled  
spvpoverridesvp:           Disabled
```

```
PXM1E_SJ.7.PXM.a >
```


cnftime

Configure Time—PXM45, PXM1E

Configures the time for the node. To see the time after you execute **cnftime**, use **dsptime**. The system displays the time in 24-hour format.



Note

Configure a time zone through **cnftmzn** and optional GMT offset through **cnftmzngmt** before you configure the time through **cnftime**.

Syntax

cnftime <hh:mm:ss>

Syntax Description

hh:mm:ss The format for time specification is:

- *hh* is the hour in the range 01–24.
- *mm* is the minute in the range 01–60.
- *ss* is the second in the range 01–60.

Related Commands

cnfdate, **cnftmzn**, **dsptime**

Attributes

Log: yes State: active Privilege: SUPER_GP

Example

Set time for 2 PM, plus 11 minutes and 22 seconds.

```
excel.1.1.3.PXM.a > cnftime 14:11:22
```

cnftmzn

Configure Time Zone—PXM45, PXM1E

Configures the time zone in the Western Hemisphere for the switch. To configure a time zone outside the four standard time zones of the Western Hemisphere, enter the GMT argument, then execute **cnftmzngmt** to specify an offset in hours from Greenwich Mean Time.

The system returns no messages unless an error occurs. To see the time zone, execute **dsptime**.

Syntax

```
cnftmzn <timezone>
```

Syntax Description

<i>timezone</i>	The possible time zones requires all uppercase characters. GMT, Greenwich Mean Time EST, Eastern Standard Time CST, Central Standard Time MST, Mountain Standard Time PST, Pacific Standard Time EDT, Eastern Daylight Time CDT, Central Daylight Time MDT, Mountain Daylight Time PDT, Pacific Daylight Time
-----------------	--

Related Commands

cnftime, **cnfdate**, **cnftmzngmt**, **dsptime**

Attributes

Log: yes State: active Privilege: SUPER_GP

Example

Configure the time zone in the node to U.S. Pacific Standard Time.

```
excel.1.1.3.PXM.a > cnftmzn PST
```

cnftmzngmt

Configure Time Zone Relative to GMT—PXM45, PXM1E

Configures the time zone for the node relative to GMT. Typically, this command applies to nodes outside the four standard time zones of the Western Hemisphere. Use **cnftmzngmt** according to the following sequence:

- First use **cnftmzn** to specify the time zone as GMT.
- Then specify an offset in hours relative to Greenwich Mean Time by executing **cnftmzngmt**. The values are GMT plus or minus an integer in the range 1–12.

Use **dsptime** to see the time.

Syntax

```
cnftmzngmt <timeoffsetGMT>
```

Syntax Description

<i>timeoffsetGMT</i>	Number of hours offset from GMT in the range -12 through 12.
----------------------	--

Related Commands

cnftmzn, **cnftime**, **cnfdate**, **dsptime**

Attributes

Log: yes State: active Privilege: SUPER_GP

Example

Set time zone in the shelf to GMT plus 4 hours.

```
excel.1.3.PXM.a > cnftmzngmt 4
```

cnftopogw

Configure Topology Gateway—PXM45, PXM1E

The **cnftopogw** command lets you configure the current switch to be a gateway node for the persistent topology feature. In the context of the Persistent Topology feature, a gateway node is a switch that can provide Cisco WAN Manager (CWM) with information about the persistent topology within a single peer group or in a multiple peer group. CWM communicates with only the gateway nodes regarding the persistent topology database. You can also use CWM to identify a gateway node.



Note

Although the **cnftopogw** command runs on a PXM1E, Cisco recommends that you do not use the PXM1E as a gateway node unless the only PXM type in the network is PXM1E.

Beginning with Release 4.0, a PXM45/A cannot be configured as the gateway node. Note that a Service Expansion Shelf (SES)/BPX node cannot be a gateway node.

The *functional* difference between a gateway node and a non-gateway node is that a gateway node sends traps to CWM when an add, delete, or a modification is done on the persistent topology database on any switch. (For a list of the information elements that cause such a trap, see the section, “[Feature Details for Persistent Topology Database](#).”) In all other respects, a non-gateway node maintains its topology database in the same way as a gateway node. The gateway functionality begins or ends as soon as you run the **cnftopogw** command with the **on** or **off** parameter, and the switch immediately sends a trap to CWM to notify it of the change. For important information about the **cnftopogw** command and the Persistent Topology Database feature, see the forthcoming sections, “[Usage Guidelines for the cnftopogw Command](#)” and “[Feature Details for Persistent Topology Database](#).”

Syntax

```
cnftopogw <yes | no>
```

Syntax Description

yes | no Type **yes** to enable or **no** to disable the current node as a gateway node.
 Default: no

Related Commands

dsptopogw, **dsptopondlist**, **dsptopogwndlist**, **deltopond**

Attributes

Log: yes

State: active

Privilege: SUPER_GP

Usage Guidelines for the cnftopogw Command

Before you use the **cnftopogw** command, consider the behaviors of the feature, as follows:

- The Persistent Topology Database feature does not work for a non PNNI switch.
- This feature does not work across IISP or AINI links because it relies on the PNNI PTSE database for acquiring nodal information.
- For an SPG network or the lowest level of an MPG, you must configure at least one gateway node. Cisco recommends that each SPG or MPG have at least two gateway nodes. If one gateway node becomes disabled or unreachable, CWM can use the other gateway node.
- To disable the role of a switch as a gateway, peer group node, you must either use the **cnftopogw** command on the gateway node or disable the gateway functionality through CWM.
- At the moment you configure a node to be a gateway, it has a snapshot of the network. This snapshot does not include nodes that are disabled or unreachable at that moment.

Feature Details for Persistent Topology Database

Gateway and non-gateway nodes in a network maintain a persistent topology database. Upon boot-up, each node populates its topology database with the information about the other nodes in the same peer group. From that time onwards, each node continues to add to its topology database whenever it discovers a new neighbor in the same peer group.

A persistent topology database entry consists of the following:

- Node identifier
- Node name
- LAN IP address
- ATM IP address
- System object identifier
- Gateway node flag to indicate whether the node indicated by the node ID is a gateway node

Propagation of persistent topology information relies on nodal PTSEs. If any of the preceding information on any node changes, nodal PTSEs are generated and flooded. When the nodal PTSE reaches a gateway node, the gateway node examines the PTSE to see if any of the topology information has changed. If it detects a change, the gateway nodes updates the topology database with the new information and sends a trap to CWM to notify it of the change.

If a PTSE with a new node ID reaches the gateway node, the gateway node responds by adding a *new* node in the peer group. (A gateway node needs only the node ID to identify a node.) Therefore, if the node ID of an existing node in the peer group is changed, the gateway node considers it to be a new node in the group.

In a network that consists of nodes running the current software and nodes with Release 2 software, the gateway node gathers nodal PTSEs regardless of the release number, but the database entries for the Release 2 nodes have only the node name and node ID.

The topology databases on the gateway nodes in the same peer group can become out of sync. A loss of synchronization can happen for a variety of reasons. Consider the following sequence, for example:

1. A gateway node is configured for a peer group.
2. A (non-gateway) node is deleted from the same peer group.
3. Another gateway node is later enabled in this peer group.

The information for the deleted node may not reach the second gateway node in a timely manner. The database on each gateway node would become synchronized upon the eventual flooding of nodal PTSEs.

**Note**

When you delete a node (not the gateway) from the persistent node database by using either CWM or the **deltopond** command, all the links that are associated with the deleted node are also deleted from the persistent link database. Also, if an *inside* link becomes an *outside* link in the common outside state, the persistent link entry belonging to the outside node is removed from the persistent topology database.

Examples

Enable the current switch to be a gateway node for the persistent topology database. Check the administrative and operational states by using the **dsptopogw** command. Follow up by using **dsptopogwndlist** to display all nodes in the persistent topology database. Only the current node appears in the list. Display details of this gateway node by using the **dsptopond** command.

```
M8850_LA.8.PXM.a > cnftopogw on

M8850_LA.8.PXM.a > dsptopogw

Admin State:      ENABLED Operational State      ENABLED

M8850_LA.8.PXM.a > dsptopogwndlist

table index:    1  node name: M8850_LA
node id:56:160:47.00918100000100001a531c2a.00001a531c2a.01

M8850_LA.8.PXM.a > dsptopond

Number of Entries = 1

Table Index: 1 Node Name: M8850_LA
Node ID: 56:160:47.00918100000100001a531c2a.00001a531c2a.01
Primary IP: 0.0.0.0
Primary IP Type: None
Secondary IP: 0.0.0.0
Secondary IP Type: None
SysObjId: 1.3.6.1.4.1.9.1.228
Gateway Mode ENABLED
PTSE in DB: YES
```

Disable the current node as a gateway for the persistent topology database then check its status.

```
M8850_LA.8.PXM.a > cnftopogw off

M8850_LA.8.PXM.a > dsptopogw

Admin State:      DISABLED Operational State      DISABLED
M8850_LA.8.PXM.a >
```

cnftrapip

Configure Trap IP—PXM45, PXM1E

The **cnftrapip** command lets you configure the trap IP for Cisco WAN Manager. You can then use the command **dspttrapip** to confirm the value.

Before you use **cnftrapip**:

1. The SNMP agent must be installed.
2. The switch's interface must have an IP address. (To assign an IP address to a switch's interface, use **ipifconfig**.)

For information about installing the SNMP agent for CWM, see *Cisco WAN Manager Installation for Solaris, Release 10* or a higher version.

Syntax

```
cnftrapip <interface> | <ip address>
```

Syntax Description

<i>interface</i>	Specify the interface type as either "atm0" or "lnPci0."
<i>ip address</i>	The switch's Ethernet IP address on which the traps are configured.

Related Commands

dspttrapip, **dspttrapmgr**, **addtrapmgr**, **deltrapmgr**

Attributes

Log: yes State: active Privilege: GROUP1

Example

Assign IP address 172.27.27.184 to the switch, then use the **dspttrapip** command to check it.

```
SanJose.7.PXM.a > cnftrapip 172.27.27.184
```

```
SanJose.7.PXM.a > dspttrapip  
Trap IP Address : 172.27.27.184
```

```
SanJose.7.PXM.a >
```

cnftrftolerance

Configure Traffic Conformance Tolerance—PXM45, PXM1E

The **cnftrftolerance** command specifies a node-level setting to permit the addition of connections whose traffic parameters *may* have a slight, unintentional mismatch between the master and slave endpoints. This feature affects only dual-ended connections.

The tolerance is a percent. The default is 5%, and any deviation from this default tightens the compliance requirement.

Syntax

cnftrftolerance %

Syntax Description

%	The percent of tolerance has a range of 0–5. Default: 5
---	--

Related Commands

dsptrftolerance

Attributes

Log: yes State: active Privilege: SUPER_GP

Example

Configure a tolerance of 2% and check the configuration by using the **dsptrftolerance** command.

```
PXM1E_SJ.7.PXM.a > cnftrftolerance 2
```

```
PXM1E_SJ.7.PXM.a > dsptrftolerance
Trf Tolerance for SPVCs: 2
```


cnfuser

Configure User—PXM45, PXM1E

The **cnfuser** command lets you configure a new password or privilege level for a user. If the user does not already exist, use the **cnfuser command** with a new user-name creates that user.

If you do not specify a user-name (*userID*) but include one or more of the other parameters, the command applies to the current user.

Syntax

```
cnfuser -u <userID> [-p <password>] [-l <accessLevel>] [-i <expiration interval>]
```

Syntax Description

-u	Keyword that specifies a string of 1–12 characters that identifies a user. The maximum number of users a switch can support is 100.
-p	(Optional) Keyword that specifies a new password with 5–15 characters for <i>userid</i> .
-l	(Optional) Keyword that specifies a new access level for the user. The <i>accessLevel</i> can be SERVICE_GP, SUPER_GP, GROUP1, or ANYUSER. The new level you type must be lower than the privilege of the current user. See adduser description for an explanation of privilege levels.
-i	(Optional) Keyword that specifies the password expiration interval. The password expiration interval is the number of days that a password is valid. At the end of those days, the user must change the password. The range for the expiration interval is 1–1000 days.

Related Commands

adduser, deluser, dspusers

Attributes

Log: no State: active Privilege: GROUP1

Example

Change the password and privilege lever of user “rocky.” New password is “nevermind,” and the privilege level is GROUP1. Note that the you must be logged in at a higher than GROUP1 privilege level to specify GROUP1 for “rocky.” If the “-u” and userID (rocky) were not entered, this command would change the password and privilege of the current user.

```
raviraj.7.PXM.a > cnfuser -u rocky -p nevermind -l GROUP1 -i 20
```

Example screens of a user changing the password when password expiration check is On:

```
MGX8850.8.PXM.a >
```

```
Login: superuser  
password:
```

```
Password has expired.  
New password:  
Re-enter new password:  
ERR: Password too short (minimum: 5 characters)  
New password:  
Re-enter new password:  
ERR: Twice-entered passwords mis-match  
New password:  
Re-enter new password:  
ERR: Password too short (minimum: 5 characters)  
You have exceeded the number of attempts allowed.
```

```
MGX8850.8.PXM.a >
```

```
Login: superuser  
password:
```

```
Password has expired.  
New password:  
New password:  
New password:  
You have exceeded the number of attempts allowed.
```

cnfuplinkbert

Configure Uplink Bit Error Rate Test—PXM1E

The **cnfuplinkbert** command sets up a bit error rate test (BERT) for a line on the PXM1E UNI/NNI back card. To start and stop BERT on the line, use the **startuplinkbert** and **stopuplinkbert** commands, respectively.

Syntax

```
cnfuplinkbert -ln <bay.line> -tp <test pattern> -tpi <transmit pattern inverse>  
-rpi <receive pattern inverse> -eir <error insertion rate>
```

Syntax Description

bay.line The required parameter identifies the bay and line that is under test. The possible values are as follows:

- *bay*: always 2 on the PXM1E
 - *line*: 1–16 maximum but depends on the number of lines supported by the back card
-

-tp	<p>The value of <i>test pattern</i> is a number in the range 1–32 and has the following significance:</p> <p>1 = all-zeros: All 0s (Continuous spaces).</p> <p>2 = all-ones: All 1s (Continuous Marks).</p> <p>3 = alt-One-Zero: Alternate 1/0 pattern (..1010..).</p> <p>4 = doubleAltOnesZeros: Double alternate 1/0 (..1100..).</p> <p>5 = oneIn4: Standard loop up remote code.</p> <p>6 = oneIn8: An eight-bit pattern which contains single 1.</p> <p>7 = oneIn16: N repetitive pattern, 1 in 16.</p> <p>8 = threeIn24: A 24-bit pattern that contains three 1s.</p> <p>9 = inbandLoopup: D4/SF Loopback activate.</p> <p>10 = inbandLoopdown: D4/SF Loopback deactivate.</p> <p>11 = twoE3MinusOne: This is a 2^3-1 pattern (7 bits in length).</p> <p>12 = twoE4MinusOne: This is a 2^4-1 pattern (15 bits in length).</p> <p>13 = twoE5MinusOne: This is a 2^5-1 pattern (31 bits in length).</p> <p>14 = twoE6MinusOne: This is a 2^6-1 pattern (63 bits in length).</p> <p>15 = twoE7MinusOne: This is a 2^7-1 pattern (127 bits in length).</p> <p>16 = twoE7MinusOneFT1Loopup: 2^7-1 Fractional T1 Loop Back Activate.</p> <p>18 = twoE9MinusOne: 2^9-1 (511 bits in length). It has the maximum of 8 (non-inverted) sequential 0s and 9 sequential 1s.</p> <p>19 = twoE10MinusOne: $2^{10}-1$ (1023 bits in length).</p> <p>20 = twoE11MinusOne: $2^{11}-1$ (2047 bits). Maximum of 15 (inverted) sequential 0s.</p> <p>21 = twoE15MinusOne: $2^{15}-1$ (32767 bits long). Max. of 15 (inverted) sequential 0s.</p> <p>22 = twoE17MinusOne: $2^{17}-1$ (131071 bits in length).</p> <p>23 = twoE18MinusOne: $2^{18}-1$ (262144 bits in length).</p> <p>24 = twoE20MinusOne: $2^{20}-1$ (1048575 bits long). Max. 19 (non-inverted), sequential 0s.</p> <p>25 = twoE20MinusOneQRSS: $2^{20}-1$ (1048575 bits). This pattern has zero-suppression (Quasi-random Signal Source).</p> <p>26 = twoE21MinusOne: $2^{21}-1$ (2097151-bit length).</p> <p>27 = twoE22MinusOne: $2^{22}-1$ (4194303-bit length).</p> <p>28 = twoE23MinusOne: $2^{23}-1$ (8388607-bit length). Highest stress, pseudo-random pattern, with a maximum of 23 (inverted) sequential 0s and 23 sequential 1s.</p> <p>29 = twoE25MinusOne: $2^{25}-1$ (33554431 bits long).</p> <p>30 = twoE28MinusOne: $2^{28}-1$ (268435455 bits long).</p> <p>31 = twoE29MinusOne: Highest stress pseudo-random pattern, with a maximum of 29 (inverted) sequential 0s.</p> <p>32 = twoE31MinusOne: Maximum 31 sequential 0s.</p>
------------	--

-tpi	You can specify that the test pattern in the transmit direction is inverted. Type a “1” for non-inverted or a “2” for inverted. Default: non-inverted
-rpi	You can specify that the test pattern in the receive direction is inverted. Type a “1” for non-inverted or a “2” for inverted. Default: non-inverted
-eir	The error insertion rate has a range of 1–7 and has the following significance: <ul style="list-style-type: none">• 1: No error insertion• 2: 1-in-100• 3: 1-in-1,000• 4: 1-in-10,000• 5: 1-in-100,000• 6: 1-in-1,000,000• 7: 1-in-10,000,000

Related Commands

dspuplinkbert, dspuplinkbertstats, startuplinkbert, stopuplinkbert

Attributes

Log: yes

State: active

Privilege: GROUP1

cnfxbaradmin

Configure Crossbar Administration—PXM45

The **cnfxbaradmin** command lets you turn off or turn on the switch planes on a switch card. The purpose for turning off switch planes is a graceful removal of the standby PXM in an MGX 8850 (PXM45) switch or an XM60 in an MGX 8950 switch. The circumstance for removing the card is a bad switch plane. If you change your mind about removing a card after turning off switch planes, you can turn them back on with the appropriate command parameter.

Graceful removal reduces lost traffic due to card removal. Removing a switching card normally causes very little cell loss, and using the **cnfxbaradmin** command further reduces the loss. When you plug in a card whose switch planes have been turned off before the card was removed, all the switch planes are automatically on.



Note

You should never disable all planes 0-3 on all XM-60s, which would be the equivalent of turning off power.

Syntax

```
cnfxbaradmin <xbarSlot> <adminStatus>
```

Syntax Description

<i>xbarSlot</i>	Identifies the slot where the switch planes reside. In an MGX 8850 switch, the range for <i>xbarSlot</i> is 7–8. In an MGX 8950 switch, the range for <i>xbarSlot</i> is 9–10 and 25–26.
<i>adminStatus</i>	Specifies whether the switch planes are on or off. Type the entire word “on” or “off.” Default: on

Related Commands

cnfxbarerrthresh, **cnfxbarmgmt**, **dspxbar**, **dspdevalms**, **dspdeverrhist**

Attributes

Log: yes

State: active

Privilege: SUPER_GP

Examples

Do the following:

1. Turn off the switch fabric for the PXM45 in slot 8—the standby card.
2. Turn on the switch fabric for slot 8.
3. Use the **dspxbar** command to check the operational state.

```
M8850_NY.7.PXM.a > cnfxbaradmin 8 off
Shut-down of switch could be service affecting!
cnfxbaradmin: Do you want to proceed (Yes/No)? y
```

```
M8850_NY.7.PXM.a > cnfxbaradmin 8 on
```

Do the following on an MGX 8950 switch:

1. Turn off the switch fabric in slot 9.
2. Use the **dspxbar** command to check the operational state.
3. Turn on the switch fabric for slot 9.

```
M8950_DC.7.PXM.a > cnfxbaradmin 9 off
Shut-down of switch could be service affecting!
cnfxbaradmin: Do you want to proceed (Yes/No)? y
```

```
M8950_DC.7.PXM.a > cnfxbaradmin 9 on
```

cnfxbarerrthresh

Configure Crossbar Error Threshold—PXM45

The **cnfxbarerrthresh** command lets you set the thresholds for the setting and clearing of crossbar error thresholds for each alarm severity. When the error rate crosses one of the thresholds, the controller shuts down the affected switch plane.



Note

Cisco Systems recommends that you leave all settings at the default values during normal operation.

The Syntax Description contains a list of possible errors. A crossbar error *threshold* consists of:

- A period for counting errors
- Severity of the resulting alarm (minor, major, and critical)
- Setting and clearing thresholds for each alarm severity

Usage Guidelines for cnfxbarerrthresh

Note the following details about this command:

- You can specify only one error type and one severity for each execution of the **cnfxbarerrthresh** command. For example, to change the thresholds for minor, major, and critical alarms for one error type, you must run **cnfxbarerrthresh** three times.
- You must enter all parameters of the threshold whether or not you change them. For example, if you want to change only the number of milliseconds (*threshtime*), you must include the existing error type, severity, alarm count, and clear count.



Note

The default settings for crossbar error thresholds are optimal for most applications. The **dspbarerrthresh** command shows the existing thresholds. Before changing thresholds, consider using **dspbarerrthresh** to check current thresholds.

Syntax

cnfxbarerrthresh *<errtype>* *<threshtime>* *<severity>* *<clrcount>* *<almcount>*

Syntax Description

<i>errtype</i>	<p>A number that identifies the type of error, as follows:</p> <ol style="list-style-type: none"> 1. Loss of synchronization (LossOfSync). 2. Transceiver error (TransceiverErr). 3. DisparityErr—an accumulation of five ASIC-level errors. 4. ParityErr—a parity error in the switch frame as a whole. 5. HeaderCRCErr—a CRC error for the switch frame header. 6. PayloadCRCErr—a CRC error for the switch frame payload. 7. RemapTwiceErr—in a redundant configuration, if the primary card that has gone to standby state, requests still go to the primary card but are redirected to the active-secondary card. If <i>more</i> than one slot is redirecting requests to the same card, that is a <i>remap twice</i> error. A single remap error raises an alarm. 8. RemapRecurrErr—if one slot redirects traffic to another slot, and that slot redirects to yet a different slot, that is a “remap recurrence” error. A single remap error raises an alarm. 9. Backpressure parity error—a parity error in the signaling for backpressure.
<i>threshtime</i>	<p>This value is a duration of error-counting. The switch continuously checks for crossbar errors, but it still uses a specific duration of time to determine whether a threshold has been crossed. This window is specified in milliseconds. As the window gets smaller, the chances for the switch to raise an alarm become smaller.</p> <p>Range: 1–20000 milliseconds</p> <p>Default: 20000</p>
<i>severity</i>	<p>The severity of the alarm resulting from the error-count per threshold-time. The range is 0–2, as follows</p> <ul style="list-style-type: none"> • 0: minor alarm • 1: major alarm • 2: critical alarm
<i>clrnt</i>	<p>The clear count is the number of errors below which the alarm changes to the next lowest severity. For example, the system clears a minor alarm for a particular type of error when the number of errors goes to 0. Similarly, if <i>clrnt</i> for a major alarm is 30, the alarm goes to minor when the count drops below 30.</p>
<i>almnt</i>	<p>The number of errors for an alarm severity above which the alarm goes to the next higher severity.</p>

Related Commands

dspxbarerrthresh

Attributes

Log: yes

State: active

Privilege: SERVICE_GP

Example

For Loss of Sync, set the clear count for critical alarms to 200.

```
pop20two.7.PXM.a > cnfxbarerrthresh 1 2000 2 200 301
```

In the sequence of command and arguments, the only value that differs from the existing threshold is the clear count of 200. If the operation is successful, the system displays the error threshold for the type of error you specified—Loss of Sync in this example.

```
pop20two                System Rev: 02.01   Dec. 05, 2000 02:29:17 GMT
MGX8850                 Node Alarm: MAJOR

                        CROSSBAR ERROR CONFIGURATION
Device Error           Thresh -- MINOR --      -- MAJOR --      -- CRITICAL --
Type                   Time   Clear Alarm    Clear Alarm    Clear Alarm
                       (msec) Count Count      Count Count      Count Count
-----
LossOfSync             2000    0    3          4    15          200    301
```

cnfxbarmgmt

Configure Crossbar Management—PXM45

The **cnfxbarmgmt** command lets you configure the following operational parameters:

- Load sharing
- Automatic shutdown
- Plane alarm threshold

The purpose of load sharing is to provide backup for the operational switch planes on the active PXM45. The backup is provided by the switch planes on the standby PXM45. The hardware on which **cnfxbarmgmt** focuses is a redundant pair of PXM45s in an MGX 8850 switch. (In an MGX 8950 switch, load sharing is always enabled.) If a switch plane fails on the active PXM45, a switch plane on the standby PXM45 takes over the switching tasks of the failed ASIC. If the active PXM45 itself fails and the standby takes over, load sharing is not available.

With automatic shutdown, the controller turns off the switch plane if the error threshold is reached.

The alarm threshold is the number of errors that can occur before the controller shuts down a link. When the threshold is exceeded, a major alarm is declared. When the number of errors falls below the threshold, the controller resets the alarm.



Note

The switch blocks the **cnfxbarmgmt** command if a bad switch ASIC exists on the active PXM45.

Regardless of whether the node has redundant PXM45s or a load-sharing configuration, you can still investigate alarms and errors through a hierarchy of shelf-management and crossbar-related commands:

1. **dspndalms**
2. **dspswalms**
3. **dspdevalms**
4. **dspdeverrhist**

Syntax

```
cnfxbarmgmt <loadSharing> <autoShutdown> <planeAlarmThresh>
```

Syntax Description

<i>loadSharing</i>	<p>Cisco Systems recommends that you leave load sharing enabled on the PXM45. On the XM60, load sharing is always enabled.</p> <ul style="list-style-type: none"> 0: Disable load sharing. 1: Enable load sharing. -1: Force load sharing to be disabled when one or more bad switch ASICs exist on the active PXM45. <p>Default: enabled</p>
<i>autoShutdown</i>	<p>If auto shutdown is enabled, a crossbar link is shut down when errors cross a major-alarm threshold for the crossbar core or crossbar port. Type 1 or 0.</p> <ul style="list-style-type: none"> 0: disables automatic shut-down 1: enable automatic shut-down <p>Default: enabled</p>
<i>planeAlarmThresh</i>	<p>Note In the current release, this parameter is not used. Nevertheless, this value is fixed at 1.</p> <p>The <i>planeAlarmThresh</i> value is an alarm threshold for declaring that a switch plane is bad. The determination of a bad link depends on the crossbar error threshold. Each unit of the threshold represents a link between the switch ASIC and the card.</p> <p>If the number of bad links reaches the threshold, the active PXM45 shuts down the ASIC and shifts the load to other switch planes.</p> <p>Range: 1–32 links</p> <p>Default: 1</p>

Related Commands

dspxbarmgmt, dspxbarerrthresh, dspdeverrhist

Attributes

Log: yes State: active Privilege: SERVICE_GP

Example

Disable load sharing, but leave auto-shutdown and plane alarm threshold at the defaults. Use the **dspxbarmgmt** command to check the configuration.

```
pop20two.7.PXM.a > cnfxbarmgmt 1 1 1

pop20two.7.PXM.a > dspxbarmgmt
pop20two                               System Rev: 03.01   Dec. 06, 2002 00:44:20 GMT
MGX8850                                 Node Alarm: MAJOR
Load Sharing: Disable
Auto Shutdown: Enable
Plane Alarm Threshold: 1
```

cnfxbarpathenable

Configure Crossbar Path Enable—PXM45

The **cnfxbarpathenable** command lets you enable a particular crossbar path. To identify the path, provide the following:

- The number of the slot in which the switch plane resides
- The number (or index) of the switch plane itself
- The number of the slot in which the pertinent service module resides



Note

Due to the momentary traffic loss, the system prompts you to confirm the task after you identify a path.

Syntax

```
cnfxbarpathenable <xbarSlot> <planeIndex> <slot>
```

Syntax Description

<i>xbarSlot</i>	The possible crossbar slot numbers depend on the chassis, as follows: <ul style="list-style-type: none"> • MGX 8850 chassis: 7, 8 • MGX 8950 chassis: 9, 10, 25, 26
<i>planeIndex</i>	The possible plane index numbers depend on the chassis, as follows: <ul style="list-style-type: none"> • MGX 8850 chassis: 0–2 • MGX 8950 chassis: 0–3
<i>slot</i>	The possible service module slot numbers depend on the chassis, as follows: <ul style="list-style-type: none"> • MGX 8850 chassis: 1–6 and 9–14 • MGX 8950 chassis: 1–6 and 11–16

Related Commands

cnfxbaradmin, **cnfxbarmgmt**, **dspxbarmgmt**, **dspxbarerrthresh**, **dspdeverrhist**

Attributes

Log: yes State: active Privilege: SERVICE_GP

Example

This example begins with a display of the crossbars (**dspxbar** command). It shows that a disable request was granted for the source and destination—the crossbar and the service module—in slot 1. Use the command to enable this path, then check the result by again displaying the crossbars.

```
GoldenXbar.7.PXM.a > dspxbar 0
GoldenXbar                System Rev: 02.01   May. 08, 2002 19:44:28 GMT
MGX8850                   Node Alarm: CRITICAL

                          PXM45 CROSSBAR CONFIGURATION
Crossbar Slot No: 7       Switch Asic No: 0   Status: OK
Cell Grant Mode: Multicast Pref      Resync Sframe Tic: Rising-Edge Detect
Asic Revision: 1

Slot  BACK PRESSURE  DISABLE REQUEST  DISABLE DATA  REDUNDANCY CONFIG
No    Grant  Mode      Dest    Src      Dest    Src      Mode    Slot
----  -
 1   Valid  InBand    Yes    Yes      No     No     Remap   1
 2   Valid  InBand    No     No       No     No     Remap   2
 3   Valid  InBand    No     No       No     No     Remap   3
 4   Valid  InBand    No     No       No     No     Remap   4
 5   Valid  InBand    No     No       No     No     Remap   5
 6   Valid  InBand    No     No       No     No     Remap   6
 7   Valid  InBand    Yes    Yes      No     No     Remap   7
 8   Valid  InBand    Yes    Yes      No     No     Remap   8
 9   Valid  InBand    No     No       No     No     Remap   9
10   Valid  InBand    No     No       No     No     Remap  10
11   Valid  InBand    No     No       No     No     Remap  11
12   Valid  InBand    No     No       No     No     Remap  12
13   Valid  InBand    No     No       No     No     Remap  13
14   Valid  InBand    No     No       No     No     Remap  14
```

```
GoldenXbar.7.PXM.a >
```

```
GoldenXbar.7.PXM.a > cnfxbarpathenable 7 0 1
There may be momentary traffic loss towards the Slot
the crossbar path is being enabled.
cnfxbarpathenable: Do you want to proceed (Yes/No)? Yes
Crossbar path on Xbar Slot: 7 Plane: 0 for slot 1 re-enabled
```

```
GoldenXbar.7.PXM.a > dspxbar 0
GoldenXbar                System Rev: 02.01   May. 08, 2002 19:44:58 GMT
MGX8850                   Node Alarm: MINOR

                          PXM45 CROSSBAR CONFIGURATION
Crossbar Slot No: 7       Switch Asic No: 0   Status: OK
Cell Grant Mode: Multicast Pref      Resync Sframe Tic: Rising-Edge Detect
Asic Revision: 1

Slot  BACK PRESSURE  DISABLE REQUEST  DISABLE DATA  REDUNDANCY CONFIG
No    Grant  Mode      Dest    Src      Dest    Src      Mode    Slot
----  -
 1   Valid  InBand    No     No       No     No     Remap   1
 2   Valid  InBand    No     No       No     No     Remap   2
 3   Valid  InBand    No     No       No     No     Remap   3
 4   Valid  InBand    No     No       No     No     Remap   4
 5   Valid  InBand    No     No       No     No     Remap   5
 6   Valid  InBand    No     No       No     No     Remap   6
 7   Valid  InBand    Yes    Yes      No     No     Remap   7
 8   Valid  InBand    Yes    Yes      No     No     Remap   8
 9   Valid  InBand    No     No       No     No     Remap   9
10   Valid  InBand    No     No       No     No     Remap  10
11   Valid  InBand    No     No       No     No     Remap  11
12   Valid  InBand    No     No       No     No     Remap  12
13   Valid  InBand    No     No       No     No     Remap  13
14   Valid  InBand    No     No       No     No     Remap  14
```

commithw

Commit Hardware—PXM1E

The **commithw** command is part of a graceful upgrade for the OC3 back card on a PXM1E. This command informs the system that an 8-line OC3 UNI/NNI back card now resides in the active and standby card slots. Thus, the switch changes the reserved card type for the slots. The upgrade replaces the 4-line OC-3 front and back card with an eight-line front and back card set.

The tasks for the graceful upgrade are as follows:

1. Upgrade the firmware on the 4-OC3 nodes to Release 4.
2. On the *standby* card slot, remove the 4 OC3 card set (front and back cards) and replace with an 8-OC3 card-set.
3. The PXM1E-8-OC3 card set comes up and appears as a 4-OC3 set because the *reserved* card type for the slot is 4-OC3.
4. Do a **switchcc** and make the new hardware active.
5. Replace the remaining 4-OC3 card set with an 8-OC3 card set.
6. Run the **commithw** command on the active card. Thereafter, the card set is a PXM1E-8-155, and lines 2.5–2.8 are available for configuration.

Syntax

commithw <Slot Number> <Revision Type>

Syntax Description

<i>Slot Number</i>	The slot number is the logical slot number of the PXM1E and depends on the chassis: <ul style="list-style-type: none"> • MGX 8850 chassis: 7 • MGX 8830: 1
<i>Revision Type</i>	Currently, the only revision type is 1: from PXM1E-4-155 to PXM1E-8-155.

Related Commands

Attributes

Log: yes State: active Privilege: SERVICE_GP

Example

Do the following tasks:

1. With the 8-OC3 card set in place, run **dspcd** and note the inserted and reserved card discrepancy.
2. Run the **commithw** command.
3. Run **dspcd** command and note the change. (Both **dspcd** outputs are truncated here.)

```
HwUpgrade.7.PXM.a > dspcd
HwUpgrade                               System Rev: 04.0 Dec. 05, 2002 19:46:15 GMT
MGX8850                                 Node Alarm: MAJOR
Slot Number    7    Redundant Slot: 8

          Front Card          Upper Card          Lower Card
          -----          -
Inserted Card:    PXM1E-8-155    UI Stratum3/B    SMF_8_OC3_SFP
Reserved Card:    PXM1E-4-155    UI Stratum3      SMFIR_4_OC3
State:            Active          Active          Active
Serial Number:    SAG06260W4D    SAG06270XVH
Prim SW Rev:      4.0(2.0)D    ---            ---
Sec SW Rev:       4.0(2.0)D    ---            ---
Cur SW Rev:      4.0(2.0)D    ---            ---
Boot FW Rev:     4.0(2.5)P1    ---            ---
800-level Rev:   03           03             03
800-level Part#: 800-12345-03    800-21557-01    000-00000-00
CLEI Code:       0             0               0
Reset Reason:    On Power up
Card Alarm:      NONE
Failed Reason:   None
Miscellaneous Information:

Type <CR> to continue, Q<CR> to stop: q
```

```
HwUpgrade.7.PXM.a > commithw 7 1

Hardware Commit Succeeded.

HwUpgrade.7.PXM.a > dspcd
HwUpgrade                               System Rev: 04.0 Dec. 05, 2002 19:48:05 GMT
MGX8850                                 Node Alarm: MAJOR
Slot Number    7    Redundant Slot: 8

          Front Card          Upper Card          Lower Card
          -----          -
Inserted Card:    PXM1E-8-155    UI Stratum3/B    SMF_8_OC3_SFP
Reserved Card:    PXM1E-8-155    UI Stratum3      SMF_8_OC3_SFP
State:            Active          Active          Active
Serial Number:    SAG06260W4D    SAG06270XVH
Prim SW Rev:      4.0(2.0)D    ---            ---
Sec SW Rev:       4.0(2.0)D    ---            ---
Cur SW Rev:      4.0(2.0)D    ---            ---
Boot FW Rev:     4.0(2.5)P1    ---            ---
800-level Rev:   03           03             03
800-level Part#: 800-12345-03    800-21557-01    000-00000-00
CLEI Code:       0             0               0
Reset Reason:    On Power up
Card Alarm:      NONE
Failed Reason:   None
Miscellaneous Information:

Type <CR> to continue, Q<CR> to stop: q
```


commitrev

Commit Revision—PXM45, PXM1E

Completes a graceful upgrade by committing to the operating firmware image as the primary version. The **commitrev** command is the necessary conclusion to a graceful upgrade. See the **loadrev** description for more details about graceful firmware changes.

The impact of **commitrev** is:

- It signifies that the primary firmware image activated through the **runrev** command is accepted.
- The previous image is removed from the card's main memory (but continues to reside on disk).
- Starting another graceful revision change becomes possible. If you attempt **loadrev** on the same card before you execute **commitrev**, the system blocks **loadrev** and states that a revision change is in progress.
- You cannot use **abortrev** to revert to the previous image. To bring a previous image into memory and run it, you must use **setrev** to force-load the image (a non-graceful revision change) or execute **restoreallcnf**.

The order of commands in a graceful upgrade, including the option of aborting the revision change, appears in the following list. For clarification of where firmware resides after each stage of the upgrade, refer to [Table 2-13](#) for a single card and [Table 2-14](#) for a redundant card pair.

1. **loadrev** loads a firmware version from the hard disk to a card's memory. In a non-redundant card setup, **loadrev** does not cause the system to reset the card.
2. **runrev** causes the primary card to start running the new version. For a redundant pair of cards, the standby becomes the active card then starts running the new version.
3. If an unacceptable problem occurs, the optional **abortrev** command restores the previous version of firmware as well as the previous database contents.
4. **commitrev** declares the new primary version to be acceptable and removes the old primary from main memory (but not the hard disk).



Note

You must terminate a current graceful upgrade by using **commitrev** or **abortrev** before starting another upgrade. For example, if you attempt to run **loadrev** on a card before using **commitrev** to complete an upgrade, the system blocks the attempt and returns an error message.

The stages of a graceful upgrade and the reset actions appear in [Table 2-13](#) and [Table 2-14](#). For a single-card upgrade, see [Table 2-13](#). For a redundant-pair upgrade, see [Table 2-14](#). The tables start by showing that, initially, the primary and secondary versions of firmware are 2.x, so the only possible operational version is 2.x. The **loadrev** command loads a generic version called 2.y, and the upgrade sequence progressively changes the primary and secondary firmware versions. If you execute **abortrev** before **commitrev**, one or two cards (redundant pair only) are reset, as the tables show.

Table 2-13 Single-Card Upgrade From 2.x to 2.y

Firmware Status	Initial Version	After loadrev	After runrev	After commitrev
Primary	2.x	2.x	2.y	2.y
Secondary	2.x	2.y	2.x	2.y
Current	2.x	2.x	2.y	2.y
			After runrev , the card resets.	



Note

In Table 2-14, **runrev** causes the standby card to become the active card and run the new version of firmware. The reversed location of the “Active” and “Standby” columns shows these changed states.

Table 2-14 Redundant Pair Upgrade From 2.x to 2.y

Firmware status	Before upgrade		After loadrev		After runrev		After commitrev	
	Active	Standby	Active	Standby	Standby	Active	Standby	Active
Primary	2.x	2.x	2.x	2.x	2.y	2.y	2.y	2.y
Secondary	2.x	2.x	2.y	2.y	2.x	2.x	2.y	2.y
Current	2.x	2.x	2.x	2.y	2.y	2.y	2.y	2.y
			abortrev resets only standby card.		abortrev resets both cards.			

Syntax

```
commitrev <slot> <revision>
```

Syntax Description

<i>slot</i>	Number of the slot where firmware must revert to previous version. In an MGX 8950 switch, the slot number cannot be 9, 10, 25, or 26.
<i>revision</i>	Revision number derived from the name of the firmware file. For an explanation, see the section, “ Version Numbering Conventions ,” in the loadrev description.

Related Commands

loadrev, **abortrev**, **runrev**, **setrev**, **dspversion**

Attributes

Log: yes

State: active

Privilege: SERVICE_GP

Example

Commit the PXM45 in slot 8 to the graceful upgrade to file pxm45_002.001.000.000_mgx.fw (so 2.1(0) is the *version*). The system prompts you to confirm that you want the command to execute.

```
pinnacle.8.PXM.a > commitrev 8 2.1(0)
```

conntrace

Connection Trace—PXM45, PXM1E

The **conntrace** command lets you trace the path of an existing SPVC or SVC. (A path trace captures the path taken through the network during the initial setup of a call.)



Note

Connection tracing does not require tracing to be enabled at the port level. However, path tracing requires you to enable it through the **pathtraceport** command.

You can view the results of a connection trace by using the **dspconntracebuffer** command. To list all connection traces, use the **dspconntracebuffers** command. The switch keeps records for the 100 most recent connection traces. (If the node reroutes a very large number of connections, older trace records are lost.)

Up to five connection traces can run on the node at one time. You can run the **conntrace** command in relation to the following:

- Slave endpoint or master of an SPVC
- Either direction of an SVC
- Either direction of any connection at a via (or transit) node

To clear one or all trace buffers, use the **clrconntracebuffer** or **clrconntracebuffers** command.



Note

The **conntrace** command for Release 4 software cannot inter-operate with Release 2.1 software.

Syntax

Two ways of identifying the connection are possible.

```
conntrace <portid> {-callRef <callRef> [-endptRef <endptRef>] [-callref-flag <flag>]}
```

or

```
conntrace <portid> {-vpi <vpi> -vci <vci>}
```

Syntax Description

<i>portid</i>	The format of the PNNI physical port identifier can vary, as follows: <ul style="list-style-type: none"> • On a PXM45: <i>slot:subslot.port:subport</i> • On a PXM1E for UNI/NNI back card: <i>slot:subslot.port:subport</i>. On the UNI/NNI back card, the subslot is always 2, but the <i>slot</i> depends on the chassis, as follows: <ul style="list-style-type: none"> – In an MGX 8850 chassis, <i>slot</i> is always the logical slot 7. – In an MGX 8830 chassis, <i>slot</i> is always the logical slot 1. • On a PXM1E for a narrowband service module (NBSM): <i>slot.port</i>. For more details, see the section, “PNNI Format,” in Chapter 1, “Introduction.”
-callRef	You can provide a call reference by itself to trace the connection. A call reference by itself identifies a point-to-point call. Range for <i>callRef</i> : 1–16777215

-endptRef	If you provide a call reference, you can also provide an endpoint reference for a point-to-multipoint (P2MP) call. Range for <i>endptRef</i> : 1–32767
-callref-flag	The flag can be either a 0 or a 1. Default: 0
-vpi	The VPI has a range that depends on whether the connection is a virtual path connection or a virtual channel connection, as follows: <ul style="list-style-type: none"> • 1–4095 for a VPC • 0–4095 for a VCC
-vci	The VCI has a range of 0–65355.

Related Commands

dspcontracebuffer, dspcontracebuffers, clrcontracebuffer, clrcontracebuffers

Attributes

Log: yes State: active Privilege: SUPER_GP

Example

Do a connection trace on the connection with a VPI/VCI of 1/100 on port 5:1.1:1, then display the trace buffer for this connection.

```
r18pop157.7.PXM.a > contrace 5:1.1:1 -vpi 1 -vci 100
```

```
r18pop157.7.PXM.a > dspcontracebuffer 5:1.1:1 1 100
```

```
Last update time:Jul 9 2002 23:42:46
Result:SUCCESS    Reason:N/A
```

```
Incoming Port:17111041    Physical PortId:5:1.1:1
VPI    :1    VCI:100    CallRef:147019
Node Name:r18pop157 NodeId:
56:160:47.00918100000000016444456a.00016444456a.01
Outgoing Port:16914433    Physical PortId:2:1.1:1
```

```
VPI    :1    VCI:47161    CallRef:72760
Node Name:r18pop161 NodeId:
56:160:47.009181000000000164444b71.000164444b71.01
Outgoing Port:16848897    Physical PortId:1:1.1:1
```

```
VPI    :1    VCI:55482    CallRef:93315
Node Name: NodeId:56:160:47.0091810000000004c113b985.0004c113b985.01
Outgoing Port:17045505    Physical PortId:4:1.1:1
```

```
VPI    :1    VCI:47074    CallRef:36277
Node Name: NodeId:56:160:47.009181000000000164444b35.000164444b35.01
Outgoing Port:16979969    VPI    :1    VCI:100    CallRef:142329    Physical
PortId:3:1.1:1
```

copy

Copy—PXM45, PXM1E

Use **copy** to copy a file to a new file on the hard disk on the PXM. The **copy** command is the same as the **cp** command.

Syntax

copy *<source file name>* *<destination file name>*

Syntax Description

<i>source file name</i>	The name of the file you intend to copy.
<i>destination file name</i>	The name of the new file resulting from copy or the name of the existing file that is over-written as a result of copy .

Related Commands

cp, cd, ls, rm, pwd, rename

Attributes

Log: yes State: active, standby, init Privilege: GROUP1

Example

Create a new firmware file without the image's suffix by copying the file named pxm_1.0.00Ef.fw to pxm_1.0.00.fw.

```
MGX8850.8.PXM.a > copy pxm_1.0.00Ef.fw pxm_1.0.00.fw
```

```
MGX8850.8.PXM.a >
```

copycons

Copy Channels—PXM1E

The **copycons** command lets you copy one or more endpoints from a single endpoint. This command works by incrementing the VCI for a VCC endpoint and the VPI for a VPC endpoint.



Note

The purpose of this command is to facilitate debugging and is not intended to be an easy way to add significant numbers of user connection.

Improper use of this command can result in dangling (unpaired) endpoints in the network.

The following steps are recommended when using this command:

1. Add a slave endpoint then a master endpoint.
2. Copy the slave endpoints by using the **copychans** command.
3. Copy the master endpoints by using the **copychans** command.

Syntax

```
copycons <source> <destn> [-rem <remote Conn Id>] [-num <num. conns to add>] [-verbose <1 | 0>]
```

Syntax Description

<i>source</i>	source ID: The endpoint that serves as a template for copying. The format of this value is: <i>ifNum.vpi.vci</i> . The range for <i>ifNum</i> is 1–31.
<i>destn</i>	destination ID: The endpoint into which the controller pastes the copied connection template. The format of this value is: <i>ifNum.vpi.vci</i> . The range for <i>ifNum</i> is 1–31.
-num	The number of consecutive endpoints to be added, starting from <i>destn</i> endpoint. Default: 1
-rem	The remote connection ID specified in the format: <i>ifNum.vpi.vci</i>
-verbose	Prints the status of cloning process if enabled. 1 = enable verbose 0 = disable verbose Default: disabled

Related Commands

addcon, **delcon**, **dspcons**

Attributes

Log: yes State: active Privilege: SERVICE_GP

Example

```
MGX8850.1.11.AXSME.a > copycons 3.10.50 3.10.60 -num 10
```

core

Core Memory Dump—PXM45, PXM1E, AXSM, AXSM-E, FRSM12

The **core** command applies to core memory dumps that can occur when a card is reset. (Whether a specific reset type leads to a core dump is configurable.) You can copy zipped files to a workstation.

The **core** task has the following functional areas (further described in the Syntax Description sections):

- It *displays*:
 - Whether core files from the processor card exist, the reset reason that triggered the core dump as well as a list of all possible reset reasons, a time stamp, and so on
 - Status of core dumps in progress
 - The current configuration of various parameters
 - A subset of core-related information on the CLI of a service module
- It lets you *configure* a wide variety of applicable functions.
- It can *take an immediate action*, such as aborting an active core dump or acquiring a snapshot of a card's core memory.

Certain functions are complex enough to warrant a detailed description. These functions are noted in the Syntax Description tables and have details in the Usage Guidelines section.

The processor card and AXSM modules support a different number of command parameters. The parameters are described by card type in subsequent sections. Furthermore, the processing of the captured memory contents differs for the processor card and the service modules, as follows:

- For the PXM, the controller writes the RAM contents as a raw data image to an unmarked part of the hard disk on the back card. Before doing so, the processor compares the reset reason to the core mask. For any match, it writes core memory to the unmarked part of the disk drive. The drive holds only two *raw* data images for the PXM, so you must copy that data to a zip file before it is overwritten. Using parameters described in the section, "[Syntax Description \(PXM\)](#)," you can save an image to a zipped file having a name you choose, as follows:

specified_name.zip

- For any model of the AXSM, a core dump can occur during card boot-up after a reset. The processor compares the reset reason to the core mask for that slot. For any match, core memory is written to a file in the root directory of the C drive. The zipped file has the following format:

core_slotslot_num.zip, where *slot_num* is the number of the slot where the AXSM resides

The node logs messages for a service module core dump. The log shows when the core dump started, finished, and aborted as well as any exceptions. To see these logs, use **dsplug -mod CRDMP**.

Ftp files to a work station. You can send files to the Cisco TAC to be unzipped and debugged.

Syntax

core (*optional parameters*—see Syntax Description for AXSM or PXM)

Syntax Description (AXSM or FRSM12)

The first entry in the following list is the **core** command itself because it produces a unique output when no parameters follow it. The list entries after the **core** command are the parameters. Each of these parameters must follow the **core** command even though each parameter description does not show **core**.

core	The core command without parameters indicates whether core dumps are enabled for the current slot and that files reside on the C drive.
?	The core command with a question mark lists the optional parameters.
mask	Enter core mask to display the following: <ul style="list-style-type: none"> • A list of all possible reset reasons • Whether the reset is enabled to trigger a core dump • The associated hexadecimal value of each reason <p>The default mask is 0x262ee. To modify the mask, use mask hex-mask. See also the section, “Usage Guidelines,” for the core mask details.</p>
mask default	Enter mask default to return the mask to the default value (0x262ee).
mask hex-mask	Type core mask followed by a hexadecimal value to modify the mask. You can specify a mask regardless of whether core dumping is enabled for the card. See the section, “ Usage Guidelines ,” for the core mask details. See also Examples.
enable	Enter core enable to enable automatic core dumping for the current slot.
disable	Enter core disable to disable automatic core dumping for the current slot.

Syntax Description (PXM)

The first entry in the following list is the **core** command itself because it produces a unique output when no parameters follow it. The list entries after the **core** command are the parameters. Each of these parameters must follow the **core** command even though each parameter description does not show **core**.

core	The core command without parameters shows the priority of core dumping at the switch level, whether images exist and the number of the image (0 or 1), the reason for the reset that caused the core dump, and a time stamp.
?	The core command followed by a question mark lists the optional parameters.
mask	Enter core mask to display the following: <ul style="list-style-type: none"> • A list of all possible reset reasons • Whether the reset is enabled to trigger a core dump • The associated hexadecimal value of each reason <p>The default mask is 0x262ee. To modify the mask, use mask hex-mask. For details, see the paragraphs on the core mask in the section, “Usage Guidelines.”</p>
mask default	Enter mask default to return the mask to the default value (0x262ee).
mask hex-mask	Type core mask followed by a hex value to modify the mask. You can specify a mask whether or not core dumping is enabled on the card. For details, see the paragraphs on core masks in the section, “ Usage Guidelines .” See also Examples.
enable	Type core enable to enable automatic core dumping for the current slot.
disable	Type core disable to disable automatic core dumping for the current slot.
hot-dump	The hot dump option directs the system to save memory but <i>not</i> to reset the card. The memory-read occurs during normal operation, so RAM is frequently changing, and the accuracy is not high. This parameter targets only the PXM. The hot dump is the only dump during which traffic continues to flow.

dump-and-reset	This parameter causes a core dump then a reset (the reverse of the usual order). It applies to only the PXM. This parameter is infrequently used. It could apply to a situation where you must reset the PXM and want the memory contents before the reset. The same results result from a hot dump followed by a manual reset.
save <i>bin-num</i> <i>file-name</i>	To save a raw data image residing in the unmarked part of the disk to a zip file, type save followed by the bin number (0 or 1) then the file name in the format <i>filename.zip</i> . For example: save 0 new_york010202.zip
clear	Enter core clear to delete the raw core data. These files are the PXM core dumps in the unmarked region of disk. The service module files remain on the C drive.
red-policy	Enter core red-policy to show the core dump policy for non-redundant service modules (not PXMs). See redundancy policy in " Usage Guidelines ."
red-policy enable	At the switch-level, enable core dumps for non-redundant service modules.
red-policy disable	At the switch-level, disable core dumps for non-redundant service modules.
dump-status	This parameter lets you see the progress of core dumps.
abort-dump	To abort a core dump in progress, core abort-dump followed by the slot number. See the abort paragraph in the section, " Usage Guidelines ."
time-out	Display the maximum number of seconds a core dump can take before the system terminates the process.
time-out <i>timeInSecs</i>	You can specify the number of seconds allowed for a core dump to complete. Range: 240–7200 seconds Default: 360 seconds
priority	Enter core priority to see if core dumping is a high-priority task or a low-priority task. The default is high priority. For some purposes, you may not want core dumping to be a high-priority task. Priority is node-level and not configurable by slot. See the section, " Usage Guidelines ," for more information on the priority.
priority high	Enter core priority high to make core-dumping a high-priority task at the node level. High priority is the default for the switch to ensure that no cores are lost, so this command is necessary only if the priority is low and needs to change.
priority low	Enter core priority low to make core-dumping a low-priority at the node level.

Usage Guidelines

A description of usage considerations for the more complex parameters follows.

Disabling Core Dumps, Timeout, and Priority

You may want to disable core dumps for a slot due to the time to write core memory to disk. For example:

- You may have isolated a problem and want to save the time required to write RAM contents to disk.
- The traffic on a card may be of such high priority that you do not want to dump core memory to disk.

As the processor gets busier, core dumps require more time. In addition to disabling core dumps for a slot, you can set the priority of core dumps to **low** at the switch level or specify a timeout period for core dumps. See **priority** and **timeout** explanations in the section, "[Syntax Description \(PXM\)](#)."

Specifying the Core Mask

The core mask is the sum of the hexadecimal numbers associated with reset reasons that are enabled to trigger a core dump. Most reasons for a card reset can be enabled to trigger a core dump. (The reasons that cannot trigger a core dump are indicated as such.) Each reset reason has an associated hexadecimal number—regardless of whether it can trigger a core dump. If the reset reason is ON, the associated hex number is an element of the mask.

To create a core mask, add the hexadecimal values for the reset reasons that you want to be in the mask. The list that follows shows the reset reasons and the default enables. For a simplified example, enter **core mask c** to specify that only a resource overflow or watchdog timeout can cause a core dump *for the slot where you enter this command*. The *default* mask as displayed by **core mask** follows:

OFF 00001 not used (cannot be turned ON)

ON 00002 DRAM Parity Error

ON 00004 WatchDog Timeout Reset

ON 00008 Resource Overflow

OFF 00010 Clear All Configuration (cannot be turned ON)

ON 00020 Missing Task

ON 00040 Reset because of PXM Low Voltage

ON 00080 Reset By Event Log Task

OFF 00100 Reset from Shell—a reset issued from a low-level debugging shell used by Cisco engineers

ON 00200 Unknown

OFF 00400 Reset from PXM—of the reasons PXM causes reset, some (e.g., **resetcd**) can cause a dump

OFF 00800 Reset System (cannot be turned ON)—the system reset triggered by the **resetsys** command

OFF 01000 Switch Core Card—the reset caused by the **switchcc** command (core card switch-over)

ON 02000 Secondary Cache Error

ON 04000 Software Error Reset

OFF 08000 S/W reset due to upgrade (cannot be turned ON)

OFF 10000 Restore All Configuration (cannot be turned ON)

ON 20000 Device Driver Error

If you add all the reset reasons that are ON in the default mask, the sum is the hexadecimal number 262ee. A reason that cannot trigger a core dump is indicated in the preceding list with "can't be turned ON." A reset reason that cannot trigger a core dump removes pertinent information from memory.

Redundancy Policy

After a redundant pair of service modules switches over, the former active card is rebooting, so a core dump is possible. Because the activated card is carrying the traffic, the time to write RAM contents from the reset card to disk normally is not an issue. For non-redundant service modules, however, the dump time may be a concern. The parameters for redundancy policy let you determine whether core dumps can occur in non-redundant service modules.

The redundancy policy is a node-level configuration. You can override the configuration on a per slot basis by enabling or disabling core dumps at the CLI of the individual card.

Aborting a Core Dump

In some circumstances, you would want to abort a service module core dump. Example situations follow:

- Two or three core dumps begin, but you do not want the switch to take the time or resources to complete these processes. Additionally, one core dump may be crucial, so to ensure that it does not time out, you could abort one or two of the other core dumps.
- You could have removed redundancy from a pair of card slots but did not disable core dumps on a card where you do not want core dumps. If a core dump begins at such a slot, you can abort the core dump from the PXM then change the configuration on the service module after it comes up.

Related Commands

ftp, ll, cd, dsplog (use the **dsplog** command with the following parameter: **-mod CRDMP**)

Attributes

Log: yes State: active, standby, init Privilege: SERVICE_GP

Examples (PXM)

First, check the current status for PXM files. Next, clear the raw core dumps then re-check the status.

```
M8850_LA.8.PXM.a > core
```

```
Saved PXM core images:
```

Bin	Reset Reason	Type	Slot	Dump Time
0	Manual dump	POPEYE 2	8	WED JAN 30 19:28:13 2002
1	Software Error Reset	POPEYE 2	8	TUE FEB 27 17:48:24 2001

```
The current core bin is 1.
Automatic core dumping is ** disabled ** for this slot.
Current Task Priority for CoreDump is [High].
```

```
Saved SM core images can be found in C:/.
```

```
M8850_LA.8.PXM.a > core clear
```

```
All the raw core dump images will be erased.
core: Do you want to proceed (Yes/No)? y
```

Check the mask and observe it has returned to the default.

```
M8850_LA.8.PXM.a > core
```

Saved PXM core images:

Bin	Reset Reason	Type	Slot	Dump Time
-----	--------------	------	------	-----------

```
-----
The current core bin is 0.
Automatic core dumping is ** disabled ** for this slot.
Current Task Priority for CoreDump is [High].
```

Saved SM core images can be found in C:/.

```
M8850_LA.8.PXM.a >
```

Get a snapshot of core memory on the PXM45. The gradual display of dots across the screen shows the progress of the dump to disk. Check the status by using the **core** command without parameters. The display shows that a snapshot occurred (“Manual dump”). Save the core dump to a zipped file called test_1.zip and confirm that it resides on the C drive.

```
M8850_LA.8.PXM.a > core hot-dump
core: Do you want to proceed (Yes/No)? y
Dumping PXM Core Image[0]:
```

```
.....
Done.
```

```
M8850_LA.8.PXM.a > core
```

Saved PXM core images:

Bin	Reset Reason	Type	Slot	Dump Time
0	Manual dump	POPEYE 2	8	WED JAN 30 19:28:13 2002
1	Software Error Reset	POPEYE 2	8	TUE FEB 27 17:48:24 2001

```
-----
The current core bin is 1.
Automatic core dumping is ** disabled ** for this slot.
Current Task Priority for CoreDump is [High].
```

Saved SM core images can be found in C:/.

```
M8850_LA.8.PXM.a >
```

Enter **ll C:** to check for the test_1.zip file. The example display is truncated to show just the zip file.

```
M8850_LA.8.PXM.a > ll C:
  size      date       time       name
-----
19517299   JAN-31-2002  22:42:22  test_1.zip
```

```
In the file system :
total space : 819200 K bytes
free space  : 689221 K bytes
```

This example shows three actions in relation to the mask: showing the current mask, specifying a mask, and restoring the default mask. Show the current core mask for the PXM45 by entering **core mask**.

```
M8850_LA.8.PXM.a > core mask
Automatic core dumping is ** disabled ** for this slot.
The current core mask is 0x64e0.

OFF 00001 not used (can't be turned ON)
OFF 00002 DRAM Parity Error
OFF 00004 WatchDog Timeout Reset
OFF 00008 Resource Overflow
OFF 00010 Clear All Configuration (can't be turned ON)
ON 00020 Missing Task
ON 00040 Reset because of PXM Low Voltage
ON 00080 Reset By Event Log Task
OFF 00100 Reset from Shell
OFF 00200 Unknown
ON 00400 Reset from PXM
OFF 00800 Reset System (can't be turned ON)
OFF 01000 Switch Core Card
ON 02000 Secondary Cache Error
ON 04000 Software Error Reset
OFF 08000 S/W reset due to upgrade (can't be turned ON)
OFF 10000 Restore All Configuration (can't be turned ON)
OFF 20000 Device Driver Error
```

Set the mask to 0xc and observe the results in the display. Most reasons are off (some are always off).

```
M8850_LA.8.PXM.a > core mask c
Automatic core dumping is ** disabled ** for this slot.
The current core mask is 0xc.

OFF 00001 not used (can't be turned ON)
OFF 00002 DRAM Parity Error
ON 00004 WatchDog Timeout Reset
ON 00008 Resource Overflow
OFF 00010 Clear All Configuration (can't be turned ON)
OFF 00020 Missing Task
OFF 00040 Reset because of PXM Low Voltage
OFF 00080 Reset By Event Log Task
OFF 00100 Reset from Shell
OFF 00200 Unknown
OFF 00400 Reset from PXM
OFF 00800 Reset System (can't be turned ON)
OFF 01000 Switch Core Card
OFF 02000 Secondary Cache Error
OFF 04000 Software Error Reset
OFF 08000 S/W reset due to upgrade (can't be turned ON)
OFF 10000 Restore All Configuration (can't be turned ON)
OFF 20000 Device Driver Error
```

Reset the mask to the default by entering **core mask default**. Note that core dumping remains disabled regardless of the mask.

```
M8850_LA.8.PXM.a > core mask default
Automatic core dumping is ** disabled ** for this slot.
The current core mask is 0x262ee.

OFF 00001 not used (can't be turned ON)
ON 00002 DRAM Parity Error
ON 00004 WatchDog Timeout Reset
ON 00008 Resource Overflow
OFF 00010 Clear All Configuration (can't be turned ON)
ON 00020 Missing Task
ON 00040 Reset because of PXM Low Voltage
ON 00080 Reset By Event Log Task
OFF 00100 Reset from Shell
ON 00200 Unknown
OFF 00400 Reset from PXM
OFF 00800 Reset System (can't be turned ON)
OFF 01000 Switch Core Card
ON 02000 Secondary Cache Error
ON 04000 Software Error Reset
OFF 08000 S/W reset due to upgrade (can't be turned ON)
OFF 10000 Restore All Configuration (can't be turned ON)
ON 20000 Device Driver Error
```

```
M8850_LA.8.PXM.a >
```

Display the redundancy policy.

```
M8850_LA.8.PXM.a > core red-policy
Non-redundant SMs are allowed to dump core.
```

Examples (AXSM)

Check the core mask on the current AXSM.

```
M8850_LA.1.AXSM.s > core mask
Automatic core dumping is enabled for this slot.
The current core mask is 0x262ee.

OFF 00001 not used (can't be turned ON)
ON 00002 DRAM Parity Error
ON 00004 WatchDog Timeout Reset
ON 00008 Resource Overflow
OFF 00010 Clear All Configuration (can't be turned ON)
ON 00020 Missing Task
ON 00040 Reset because of PXM Low Voltage
ON 00080 Reset By Event Log Task
OFF 00100 Reset from Shell
ON 00200 Unknown
OFF 00400 Reset from PXM
OFF 00800 Reset System (can't be turned ON)
OFF 01000 Switch Core Card
ON 02000 Secondary Cache Error
ON 04000 Software Error Reset
OFF 08000 S/W reset due to upgrade (can't be turned ON)
OFF 10000 Restore All Configuration (can't be turned ON)
ON 20000 Device Driver Error
```

Set the core mask to 0xc and note the result in the display.

```
M8850_LA.1.AXSM.s > core mask c
Automatic core dumping is enabled for this slot.
The current core mask is 0xc.

OFF 00001 not used (can't be turned ON)
OFF 00002 DRAM Parity Error
ON 00004 WatchDog Timeout Reset
ON 00008 Resource Overflow
OFF 00010 Clear All Configuration (can't be turned ON)
OFF 00020 Missing Task
OFF 00040 Reset because of PXM Low Voltage
OFF 00080 Reset By Event Log Task
OFF 00100 Reset from Shell
OFF 00200 Unknown
OFF 00400 Reset from PXM
OFF 00800 Reset System (can't be turned ON)
OFF 01000 Switch Core Card
OFF 02000 Secondary Cache Error
OFF 04000 Software Error Reset
OFF 08000 S/W reset due to upgrade (can't be turned ON)
OFF 10000 Restore All Configuration (can't be turned ON)
OFF 20000 Device Driver Error
```

Determine if core dump is enabled for the current slot.

```
M8850_NY.1.AXSM.s > core
Automatic core dumping is enabled for this slot.

Saved core images are on PXM's hard disk (C:/).

M8850_NY.1.AXSM.s >
```


cp

Copy—PXM45, PXM1E

Use **cp** to copy a file to a new file on the disk on the PXM45 or the PXM1E. This command is the same as the **copy** command.

Syntax

```
cp <source file name> <destination file name>
```

Syntax Description

<i>source file name</i>	The name of the file you intend to copy.
<i>destination file name</i>	The name of the new file resulting from cp or the name of the existing file that is over-written as a result of cp .

Related Commands

cd, ls, rm, pwd, rename

Attributes

Log: yes State: active, standby, init Privilege: GROUP1

Example

Create a new firmware file without the image's suffix by copying the file named pxm_1.0.00Ef.fw to pxm_1.0.00.fw.

```
MGX8850.8.PXM.a > cp pxm_1.0.00Ef.fw pxm_1.0.00.fw
```

