NORTEL NETWORKS

Passport 7400, 15000, 20000

Path-Oriented Routing System Guide

241-5701-435

Passport 7400, 15000, 20000

Path-Oriented Routing System Guide

Publication: 241-5701-435
Document status: Standard
Document version: 5.2S1
Document date: November 2003

Printed in Canada

Publication history

November 2003

5.2S1 Standard
General availability. Contains information on Passport 7400, Passport 15000, and Passport 20000 for the PCR5.2 GA release.

Contents

List of figures

List of tables

About this document

This user guide describes how to configure, monitor and troubleshoot the Passport Path-oriented Routing System (PORS). For an overview of PORS, see 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*.

Who should read this document and why

This guide is for anyone who performs the following tasks for PORS on the Passport system:

- planning
- installing and provisioning
- operating and maintaining

What you need to know

This guide assumes that you are familiar with the concepts relating to PORS and Passport trunking. For and over of PORS, refer to 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*.

How this document is organized

The 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide* is organized as follows:

What’s new in this document

There were no new features added to this document.

Text conventions

This document uses the following text conventions:

- **`nonproportional spaced bold type`**

  Nonproportional spaced bold type represents words that you should type or that you should select on the screen.

- *italics*

  Words that appear in italics indicate a software component or attribute name.

- `[optional_parameter]`

  Words in square brackets represent optional parameters. The command can be entered with or without the words in the square brackets.

- `<general_term>`

  Words in angle brackets represent variables which are to be replaced with specific values.

Procedure conventions

This document uses the following procedure conventions:

- You can enter commands using full component and attribute names, or you can abbreviate them. The commands used in the procedures contain the full component and attribute names in the first instance. In the second instance, the component and attribute names are abbreviated. For more

information on abbreviating component and attribute names, see 241-5701-060 *Passport 7400, 15000, 20000 Components*. All component and attribute names are formatted in italics.

- The introduction of every procedure states whether you must perform the procedure in operational mode or provisioning mode. For more information on these modes, see “Operational mode” (page 17) or “Provisioning mode” (page 17).
- When you complete a procedure, you can verify your changes and then activate them as the new node configuration. For more information on completing configuration changes and exiting provisioning mode, see “Activating configuration changes” (page 18).

Operational mode

Procedures contained within this document can either be performed in operational mode or provisioning mode. When you initially log into a Passport node, you are in operational mode. Passport uses the following command prompt when you are in operational mode:

```
#>
```

where:
is the current command number

In operational mode, you work with operational components and attributes. In operational mode, you can

- list operational components and display operational attributes to determine the current operating parameters for the node
- control the state of parts of the node by locking and unlocking components
- set certain operational attributes and enter commands to perform diagnostic tests

Provisioning mode

To change from operational mode to provisioning mode, type the following command at the operator prompt:

```
start Prov
```

Only one user can be in provisioning mode at a time. Passport uses the following command prompt whenever you are in provisioning mode:

```
PROV #>
```

where:

is the current command number

In provisioning mode, you work with the provisionable components and attributes that contain the current and future configurations of the node. You can add and delete components, and display and set provisionable attributes. For information on completing the configuration changes, exiting provisioning mode, and returning to operational mode see “Activating configuration changes” (page 18).

For information on operational and provisionable attributes, see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Activating configuration changes

Several procedures in this document ask that you complete the configuration changes. When you complete the configuration changes, you are activating the configuration changes, confirming that you want to activate them, and saving the changes. You are instructed to complete the configuration changes only at the end of procedures that you perform in provisioning mode.

CAUTION
Activating a provisioning view can affect service
Activating a provisioning view can result in a CP reload or restart, causing all services on the Passport node to fail. See 241-5701-050 *Passport 7400, 15000, 20000 Commands*, for more information.

1 Verify that the provisioning changes you have made are acceptable:

```
check Prov
```

Correct any errors and then verify the provisioning changes again.

2 If you want to store the provisioning changes in a file, save the provisioning view:

```
save Prov
```

3 If you want these changes as well as other changes made in the edit view to take effect immediately, activate, confirm, and commit the provisioning changes:

```
activate Prov
```

```
confirm Prov
```

```
commit Prov
```

4 End the provisioning session:

```
end Prov
```

Related documents

For the complete list of documents in the Passport documentation library, see 241-5701-001 *Passport 7400, 15000, 20000 Documentation Guide*.

The following sections contain documents related to the information in this guide:

• "Passport documents" (page 19)

Passport documents

The following documents containing information related to PORS and the Passport system are available from Nortel Networks:

- 241-5701-001 *Passport 7400, 15000, 20000 Documentation Guide*
- 241-5701-005 *Passport 7400, 15000, 20000 List of Terms*
- 241-5701-030 *Passport 7400, 15000, 20000 Overview*
- 241-5701-050 *Passport 7400, 15000, 20000 Commands*
- 241-5701-060 *Passport 7400, 15000, 20000 Components*
- 241-5701-420 *Passport 7400, 15000, 20000 Trunking Guide*
- 241-5701-500 *Passport 6400, 7400, 15000, 20000 Alarms*
- 241-5701-600 *Passport 7400, 15000, 20000 Configuration Guide*
- 241-5701-615 *Passport 7400, 15000, 20000 FP Configuration Reference*
- 241-5701-700 *Passport 7400, 15000, 20000 ATM Overview*

- 241-5701-702 *Passport 7400, 15000, 20000 ATM Routing and Signaling Fundamentals*
- 241-7401-200 *Passport 7400 Hardware Description*
- 241-7401-240 *Passport 7400 Hardware Installation, Maintenance and Upgrade*

How to get more help

For information on training, problem reporting, and technical support, see the "Nortel Networks works support services" section in the product overview document.

Chapter 1
PORS task flow overview

PORS configuration is dependent on its supporting Passport nodes being properly installed and configured. This book deals only with the procedures used to configure PORS once the supporting network infrastructure is in place. See “PORS in relation to overall Passport configuration work flow” (page 22) for a view of how PORS configuration fits into overall Passport configuration.

For a detailed view of the sequence of tasks you perform to configure PORS on Passport, see “PORS configuration work flow” (page 23). Each box in the task flow represents a task that comprises one or more procedures. Each task has a corresponding section in this guide that contains the relevant procedures. To link to any task, go to the list that follows the task flow.

If you are unfamiliar with PORS concepts and procedures, conceptual information about PORS can be found in “Introduction to Path-Oriented Routing System” (page 83).

Figure 1
PORS in relation to overall Passport configuration work flow

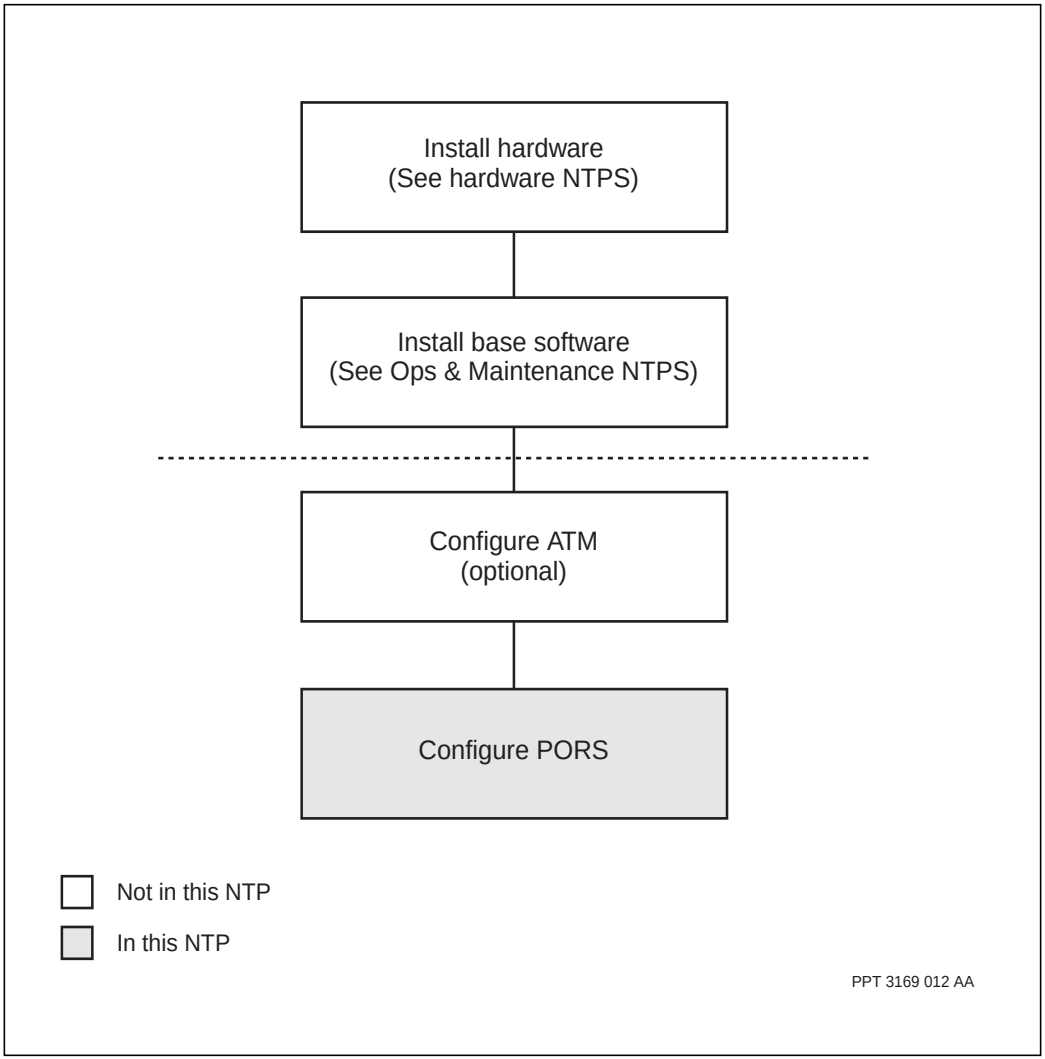

Figure 2
PORS configuration work flow

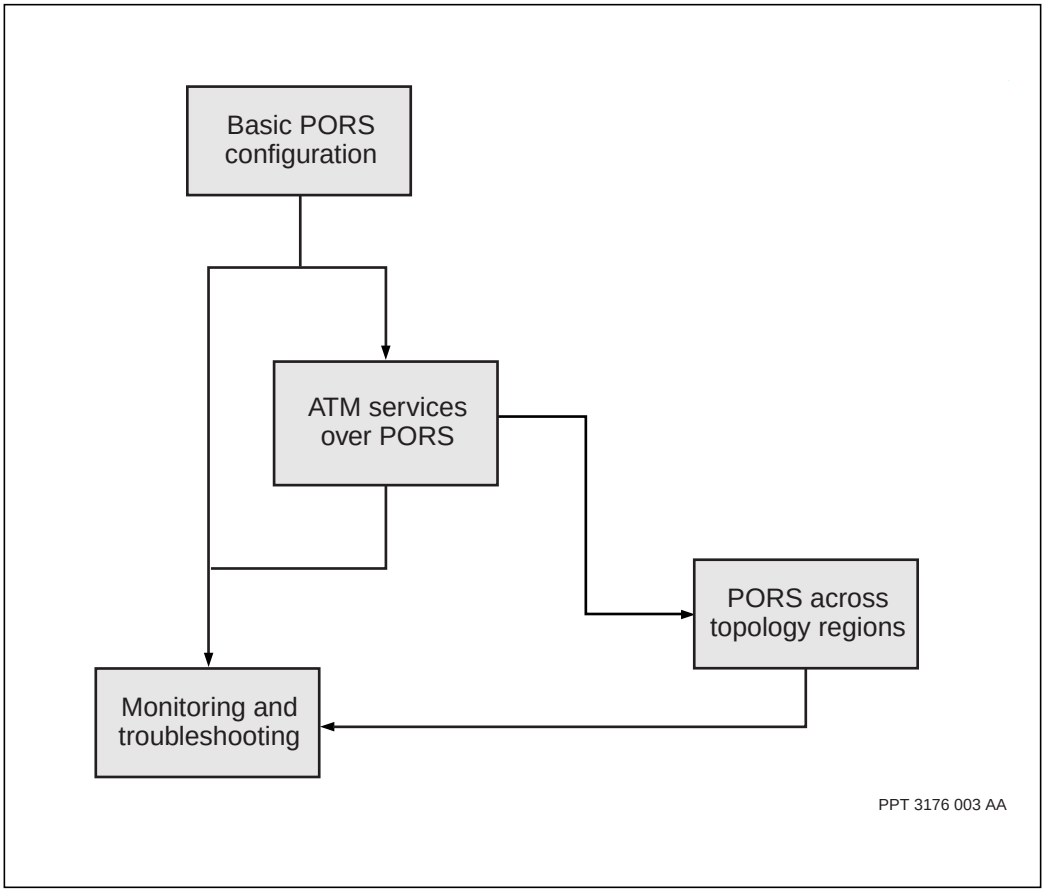

- “Configuring PORS” (page 25)
- “Configuring ATM service over PORS” (page 43)
- “Configuring PORS across topology regions” (page 49)
- “Monitoring and troubleshooting PORS” (page 59)

Chapter 2
Configuring PORS

Configure the Path-Oriented Routing System on Passport trunks

Prerequisite to PORS configuration

- The required processor cards have been provisioned.
- For supporting information about overall PORS configuration, see “Introduction to Path-Oriented Routing System” (page 83)
- For supporting information about PORS routing, see “Routing fundamentals” (page 89)
- For supporting information about PORS applications, see “Routing applications” (page 121)
- For supporting information about traffic management with PORS, see “Traffic management” (page 143).
- For procedure to monitor and troubleshoot PORS connections, see “Monitoring and troubleshooting PORS” (page 59).

PORS configuration task flow

This task flow shows you the sequence of procedures you perform to configure PORS. To link to any procedure, go to the list that follows the task flow.

Figure 3
PORS configuration task flow

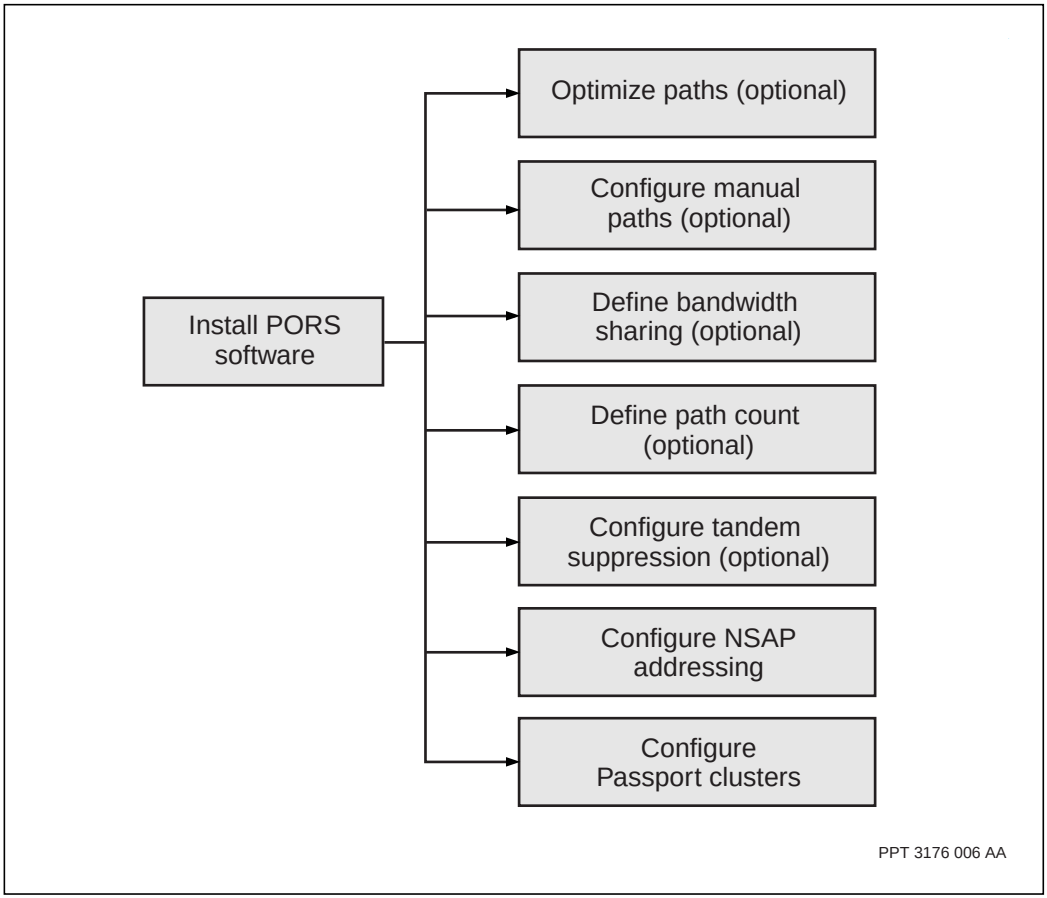

- “Installing PORS software” (page 28)
- “Optimizing paths” (page 29)
- “Configuring optional manual paths for managed cut through switching” (page 31) and “Configuring optional manual paths for a PLC” (page 33)
- “Defining bandwidth sharing on Passport trunks” (page 35)
- “Configuring the path count for Passport trunks” (page 36)
- “Configuring tandem suppression” (page 37)

- “Configuring NSAP addressing” (page 38)
- “Configuring Passport clusters” (page 39)

Installing PORS software

Use the procedures in 241-5701-270 *Passport 7400, 15000, 20000 Software Installation Guide* to install PORS software.

Optimizing paths

Specify whether the connection allows PORS optimization. If this attribute is set to disabled, the PLC will not optimize the connection regardless of the setting for the *optimizationInterval* attribute. If set to enabled, the connection is optimized as specified through the *optimizationInterval* attribute.

> ***Note:*** A connection must be in service for at least 1 minute prior to an optimization request otherwise the optimization process does not proceed for that connection.

Prerequisites

- For supporting information about optimizing paths, see "Routing fundamentals" (page 89)

Procedure

1. On each switch that supports PORS, start the optimize process:

   ```
   optimize Routing Pors
   ```

2. Display the attributes for the *Routing Pors* component to determine the effects of the optimization.

   ```
   display -o Routing Pors
   ```

3. Optionally, set automated optimization.

   ```
   set <service>/<instance> PermanentLogicalConnection
   optimization <opt_setting>
   ```

   ```
   set Routing Pors optimizationInterval <opt_value>
   ```

Variable definitions

Variable	Value
<service>	This value is either HdlcTransparentDataService or BitTransparentDataService. There is no default value.

(Sheet 1 of 2)

Variable	Value
<instance>	The instance value assigned is a decimal from 1 to 65 535, which identifies the instance of the service on the switch. There is no default value.
<opt_setting>	This value is set to either disabled or enabled. The default value is enabled.

(Sheet 2 of 2)

Configuring optional manual paths for managed cut through switching

Configure a manual path, which consists of a sequence of hops, identified by node name and egress trunk for managed cut through switching. Services linked to a manual path can also override the manual path setting through the *manualPath* or *porsManualPath* attribute associated with connections for that service.

Note: If you change the configuration of an existing manual path to specify a different route, the services that use that *ManualPath* component automatically switch to the new route. You do not need to lock or disable the service.

Prerequisites

- For supporting information about manual paths, see “Routing applications” (page 121) and “Routing fundamentals” (page 89)

Procedure

1. Add the *ManualPath* component with an instance number

   `add Routing PathOrientedRoutingSystem ManualPath/<n>`

2. Define the first hop in the route by specifying the originating node and the trunk used to leave that node for the next node in the path.

   `set Rtg Pors MPath/<n> route <index> “<switch> <trunk>”`

   For example, to define the first hop in a route that originates at node EM/a, leaving on trunk 10, use this command:

   `set Rtg Pors MPath/1 route 0 “EM/a trk/10”`

3. Define the next hop in the route by specifying the trunk used to leave the second node in the path.

   `set Rtg Pors MPath/<n> route ,index. “<switch> <trunk>”`

   For example, to define the secondhop in a route that leaves node EM/b, on trunk 20, use this command:

   `set Rtg Pors MPath/1 route 1“EM/b trk/20”`

4. Repeat step 3 for each hop in the route, specifying the trunk that leaves the node at each hop. The last entry in the route identifies the trunk

leaving the second-last node in the path. The destination node in not explicitly defined.

5 Link the service to the *ManualPath* component.

```
set Rtg Pors MPath/<n> users <comp_type>/instance>
<comp_type>/instance ... <comp_type>/instance>
```

Variable definitions

Variable	Value
<index>	This value is a decimal from 0 to 19 that indicates the position of the hop in the route. The first hop is assigned index 0, the second is assigned index 1, and so on.
<n>	This value is a decimal of 1 to 255 that identifies the PORS manual path. There is no default value.
“<switch> <trunk>”	An ASCII string of up to 40 characters that defines the originating node in the path and the trunk used to leave the node. The string must be enclosed in double quotation marks (“ “).

Additional Supporting information

You can link multiple services to the same manual path on a Passport series switch. See the following service guides:

- 241-5701-720 *Passport 7400, 15000, 20000 AAL1 Circuit Emulation Guide*
- 241-7401-440 *Passport 7400 Frame Relay Managed Cut-through Switching Guide*

Configuring optional manual paths for a PLC

Configure a manual path, which consists of a sequence of hops, identified by node name and egress trunk for transparent data services (TDS). Services linked to a manual path can also override the manual path setting through the *manualPath* or *porsManualPath* attribute associated with connections for that service.

Note: If you change the configuration of an existing manual path to specify a different route, the services that use that *ManualPath* component automatically switch to the new route. You do not need to lock or disable the service.

Prerequisites

- For supporting information about manual paths, see “Routing applications” (page 121) and “Routing fundamentals” (page 89)

Procedure

1. At the source node, add the btds instance

   `add btds/<instance>`

2. Set the interface name

   `set btds/<instance> framer intefaceName Lp/<instance> x21/<instance>`

3. Set the path type to manual.

   `set btds/<instance> plc pathType manual`

4. Set the remote name

   `set btds/<instance> plc remoteName “<switch>”`

5. Define the first hop in the route by specifying the originating node and the path to be taken in order to exit the node.

   `set btds/<instance> plc manual path 0 “<switch> <trunk>”`

   Note: 0 is the decimal that indicates the position of the hop. the first hop is assigned index 0, the seconds assigned index 1 and so on.

6. At the destination node, repeat steps 1 and 2.

7. At the destination node, choose one of the following options:

a. Set the PathType to manual, optionally with a different path to accommodate any failure on the primary path

b. Leave the PathType as default normal, but do not provision any remote name at this end

Variable definitions

Variable	Value
<index>	This value is a decimal from 0 to 19 that indicates the position of the hop in the route. The first hop is assigned index 0, the second is assigned index 1, and so on.
<n>	This value is a decimal of 1 to 255 that identifies the PORS manual path. There is no default value.
“<switch> <trunk>”	An ASCII string of up to 40 characters that defines the originating node in the path and the trunk used to leave the node. The string must be enclosed in double quotation marks (“ “).

Additional Supporting information

You can link multiple services to the same manual path on a Passport series switch. See the following service guides:

- 241-5701-720 *Passport 7400, 15000, 20000 AAL1 Circuit Emulation Guide*
- 241-7401-440 *Passport 7400 Frame Relay Managed Cut-through Switching Guide*

Defining bandwidth sharing on Passport trunks

Define a percentage of trunk bandwidth can be reserved for path-oriented traffic. This setting also prevents path-oriented traffic from using more than the percentage specified.

Prerequisites

- For supporting information, see “Passport trunk bandwidth sharing” (page 150)

Procedure

1 Set the reserved bandwidth as a percentage:

```
set Trunk/<trk_n> PathAdministrator maxReservedBwOut
<percentage>
```

Note: This attribute controls the allocation limit for path-oriented traffic. This limit is not enforced in real time on the Passport trunk.

Variable definitions

Variable	Value
<percentage>	This value is a decimal of 0 to 100, which defines the percentage of trunk bandwidth allocated to path-oriented traffic. The default value is 50.
<trk_n>	This value is a decimal instance value of the trunk.

Configuring the path count for Passport trunks

Configure the maximum number of logical channels permitted on a trunk. The value for the *maxLc* attribute can be as large as FP memory allows. The memory requirements for a given value in maxLc is defined as:

```
memory = 10,000 bytes + (maxLc * 36 bytes)
```

Note: If the *PathAdmin* components at each end of a Passport trunk are configured with different values for the *maxLc* attribute, the switch generates an alarm. PORS selects the value that is smallest for the trunk instance. If the value for the *maxLc* attribute is too large for available memory, the switch generates an alarm and the connection request is ignored.

Procedure

1 Set the number of logical channels for the trunk.

```
set Trunk/<trk_n> PathAdministrator maxLc <value>
```

Note: The value set for the *maxLc* attribute should be added to the *connectionPoolCapacity* attribute under the *Lp Eng Arc Ov* component. This ensures that PORS connections will establish on map mode trunks. For details, see 241-5701-710 *Passport 7400, 15000, 20000 ATM Configuration Guide*.

Variable definitions

Variable	Value
<trk_n>	This value is a decimal representing the instance value of the trunk.
<value>	This value is a decimal between 0 and 65 435, and indicates the maximum number of logical channels that the trunk can support. The default value is 512.

Configuring tandem suppression

Configure tandem suppression to prevent tandem traffic from traversing selected switches and to give the network operator control over routing behavior. This feature only applies to Passport 7400 series switches.

Prerequisites

- For supporting information, see “Tandem suppression” (page 127)

Procedure

1. Set the *tandemTraffic* attribute under the *Routing* component.

   ```
   set Routing tandemTraffic <permissions>
   ```

Variable definitions

Variable	Value
<permission>	This value is either allowed or denied. The default value is denied.

Configuring NSAP addressing

Configure an NSAP address for the node prefix and the alternate PORS prefixes. For routing purposes, PORS converts these address to Passport switch identifiers.

Note 1: Changing a node prefix is a critical configuration change that results in the node rebooting (unless the prefix was initially blank).

Note 2: Changing a prefix is a critical configuration change, causing the switch to reboot (unless the prefix was initially blank). Once the prefix is changed, it can never be set back to a blank.

Prerequisites

- for supporting information, see “NSAP addressing” (page 132)
- for additional information about NSAP addressing, see 241-5701-702 *Passport 7400, 15000, 20000 ATM Routing and Signaling Fundamentals*

Procedure

1 Set the node prefix for the switch

```
set ModuleData nodePrefix “<prefix>”
```

2 Optionally, set the alternate PORS prefixes for the switch.

```
set ModuleData alternatePorsPrefix 1:”<alt_pref>”
[2:”<alt_pref>” [3:”alt_pref>” [4:”<alt_pref>”]]]
```

Variable definitions

Variable	Value
<prefix>	This value is a string of 0 to 26 characters, and is a valid NSAP address prefix. There is no default value.
<alt_pref>	This value is a string of 1 to 40 characters. The default is “ “.

Configuring Passport clusters

Configure Passport clusters to allow for a larger number of Passports to be deployed in a topology region or RID subnet. Passport clusters also allow for improved network scaling in terms of CPU, memory, and control traffic bandwidth

Note: This procedure assumes that Passport clusters are being deployed in a network with existing PORS services present on planned cluster links.

Prerequisites

- Identify the backbone-cluster boundary that divides the backbone nodes and the cluster nodes. The backbone nodes and cluster nodes are connected by cluster links.
- Provision node NSAP addresses on all nodes running PORS services. Refer to “Configuring NSAP addressing” (page 38) for details.
- Reconfigure VTDS services for addressing based on NSAP addressing. Refer to “Reconfiguring VTDS services to use NSAP” (page 51) for details.
- Verify that all function processors (FPs) being used for border nodes are supported. Refer to 241-5701-615 *Passport 7400, 15000, 20000 FP Configuration Reference* for a list of supported FPs.
- Ensure that the *overrideTrunkDelay* attribute under the *Trunk Pa* component in properly configured for all future cluster links. Refer to “Bandwidth load spreading” (page 138) for more details.
- For DPN considerations, ensure that the call server resource module (CSRM) is in the backbone, and that all local call routers on cluster nodes are removed.
- For supporting information, see “Passport clusters” (page 140)

Procedure

1 In order to route PORS services across the planned cluster link, configure the *RoutingGateway* component on the backbone border nodes and cluster border nodes. To share load and provide some redundancy in the event of an LP failure, more than one *RoutingGateway* instance can be configured on each cluster border node and backbone border node. For

details on configuring the *RoutingGateway* component, refer to "Configuring a routing gateway" (page 53).

2 Optionally, if legacy PORS traffic exists on the planned cluster, set the *fullTopExchgOnClusterLinks* attribute to enabled on all planned cluster nodes:

```
set rtg fullTopExchgOnClusterLinks enabled
```

3 To provision a cluster node, set the *clusterNode* attribute to yes. This attribute should be configured on a node-by-node basis, with Passport cluster nodes being deployed starting from the edge of the topology region or RID subnet.

```
set rtg clusterNode yes
```

4 Configure reachable addresses on all cluster links. Refer to "Configuring inter-region Passport trunks" (page 55) for details.

Note: Ensure that you also unlock all *RoutingGateway* components.

5 Optionally, if the NSAP address plan allows for summarization, configure summary addresses on cluster border nodes or backbone nodes.

```
add Routing SummaryAddress/<sum_add>
```

```
set Routing SummaryAddress/<sum_add> cost <cost_value>
```

For more information on summary addresses, refer to "Displaying NSAP address information" (page 72).

6 After one hour, verify the backbone topology has aged out of the topology database on cluster nodes:

```
display rtg top node/*
```

7 Segment all PORS calls crossing the cluster-backbone boundary. This is achieved by either locking and subsequently unlocking each service instance, causing the call to reestablish on the source end point (SEP), or by clearing all the logical channels on all cluster links.

8 If set previously in step 2, set the *fullTopExchgOnClusterLinks* attribute to disabled on all cluster nodes that have PORS traffic.

```
set rtg fullTopExchgOnClusterLinks disabled
```

Variable definitions

Variable	Value
<cost_value>	This is the value of the *cost* attribute.
<sum_add>	This is the instance value of the *SummaryAddress* subcomponent.

Chapter 3
Configuring ATM service over PORS

Configure ATM service over PORS.

Prerequisites to configuring ATM services over PORS

- The PORS service must be initially installed and configured. See "Configuring PORS" (page 25).
- To understand how to set up a PORS network for ATM services, see "Configuring ATM services over PORS on a Passport switch" (page 129)
- For supporting conceptual and reference information about ATM configuration, see 241-5701-710 *Passport 7400, 15000, 20000 ATM Configuration Guide*.

ATM service over PORS configuration task flow

This task flow show you the sequence of procedures you perform to configure ATM service over PORS. To link to any procedure, go to the list that follows the task flow.

Figure 4
ATM service over PORS configuration task flow

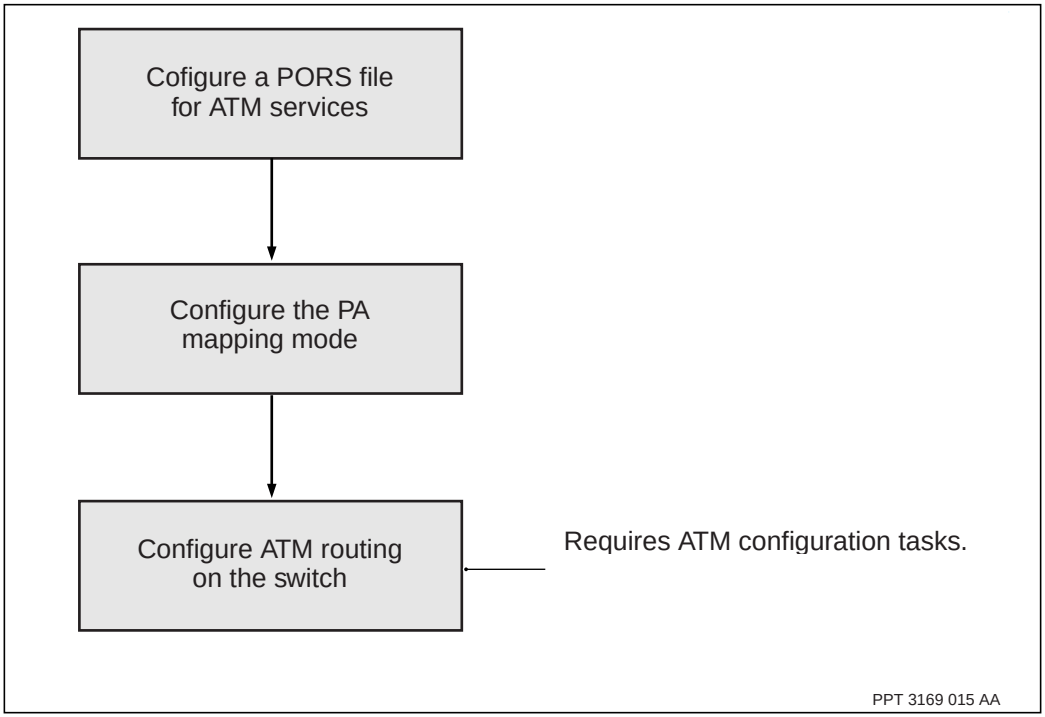

- “Configuring a PORS profile for ATM services” (page 45)
- “Configuring the path administrator mapping mode” (page 46)
- “Configuring the switch for ATM routing” (page 47)

Configuring a PORS profile for ATM services

Create a PORS profile that supports ATM services

Prerequisites

- For supporting information, see “PORS profiles for ATM services” (page 129)

Procedure

1. Add a routing profile for PORS

   add Routing PathOrientedRoutingSystem Profile/ <instance>

2. Display the default PORS attributes for this profile.

   display Rtg Pors Prof/<instance>

3. Override the discard priority for the ATM service category that uses this profile.

   set Rtg Pors Prof/<instance> discardPriority <dp_value>

4. Override the ATM bandwidth values for both the receive and transmit directions.

   set Rtg Pors Prof/<instance> requiredTx 18100, requiredRx 18100

Variable definitions

Variable	Value
<instance>	This value is a decimal from 1 to 255. There is no default value.
<dp_value>	This value is one of sameAsApplication, dp1, dp2, or dp3. The default value is sameAsApplication.

Configuring the path administrator mapping mode

Configure the ATM subcomponent of the connection path administrator in mapping mode to support AAL1 traffic on PORS. In this configuration, the path administrator creates new virtual channel connections (VCC) for each new logical channel that it sets up.

Note: Mapping can only be used in a directly connected Passport configuration or in a network configuration where a virtual path is dedicated to the PA. Logical trunks cannot be configured in mapped mode.

Prerequisites

- For supporting information about the path administrator, see “Routing fundamentals” (page 89) and “Path administrator mode for the trunk” (page 132).

Procedure

1. Set the connection path administrator for ATM access.

   ```
   set Trunk/<trk_n> PathAdmin AtmAccess
   ```

2. Set map mode

   ```
   set Trunk/<trk_n> PathAdmin AtmAccess mapping
   ```

Variable definitions

Variable	Value
<trk_n>	This value is a decimal from 1 to 65 535. This procedure assumes that the trunk is already configured. There is no default value.

Configuring the switch for ATM routing

Configure the destination module of the service for ATM routing. ATM routing protocols handles the address resolution.

Prerequisites

- for supporting information, see “Routing fundamentals” (page 89)
- for additional information, see 241-5701-702 *Passport 7400, 15000, 20000 ATM Routing and Signaling Fundamentals*

Procedure

1. Add the atmlisp, atmUni, and atmPnni features to the feature list for at least one logical processor.

   ```
   set Software LogicalProcessorType/ATM featureList
   atmIisp atmUni atmPnni
   ```

2. Add ATM routing functionality.

   ```
   add AtmRouting
   ```

Chapter 4
Configuring PORS across topology regions

Configure PORS across topology regions to allow Passport to select and establish a call path from a source end point in one topology region to a destination end point in another topology region.

Prerequisites to PORS across topology regions configuration

- The PORS service must be initially installed and configured. See "Configuring PORS" (page 25).
- PORS must be configured for ATM routing and NSAP switch addresses must be defined. See "Configuring ATM service over PORS" (page 43).
- For supporting conceptual and reference information about PORS across topology regions, see "Selecting paths based on cost and delay" (page 121)

PORS across topology regions configuration task flow

This task flow shows you the sequence of procedures you perform to configure PORS across topology regions. The link to any procedure, go to the list that follows the task flow.

Figure 5
PORS across topology regions configuration task flow

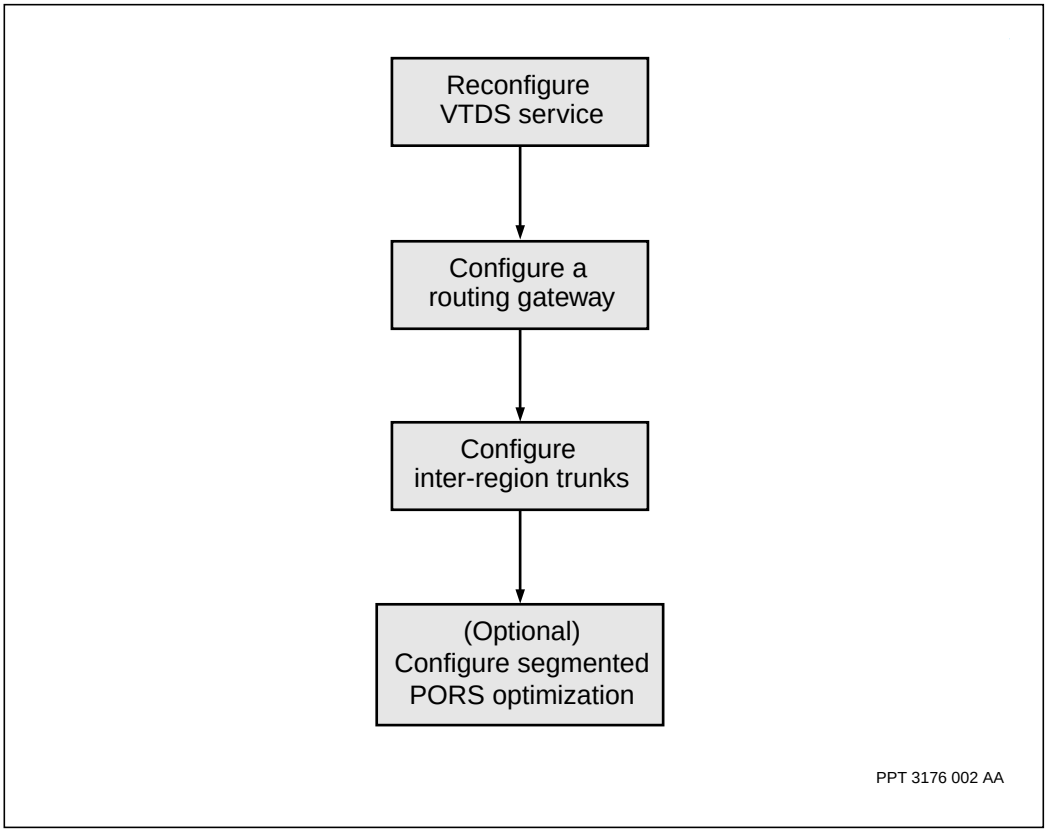

- “Reconfiguring VTDS services to use NSAP” (page 51)
- “Configuring a routing gateway” (page 53)
- “Configuring inter-region Passport trunks” (page 55)
- “Configure segmented PORS optimization” (page 56)

Reconfiguring VTDS services to use NSAP

Reconfigure voice transparent data services (VTDS) to make calls across topology regions using NSAP addressing. You can configure VTDS services for calling from source to destination and from destination to source; or only from source to destination.

Prerequisites

- “PORS across topology regions” (page 133)

Procedure

Note: If you are configuring calls only from source to destination, you only need to define a remote name setting on the destination switch.

1. On the switch at the near end point, define the local address.

   set <service>/<instance> PermanentLogicalConnection localAddress “<loc_addr>”

2. Define the address of the far end point.

   set <service>/<instance> Plc addressToCall “<rem_addr>”

3. Set the upper bound on the gateway cost metric for the route.

   set <service>/<instance> Plc maximumAcceptableGatewayCost “<value>”

4. On the switch at the remote end point, repeat step 1 through step 3.
5. Activate the configuration on both switches.
6. When the call has been verified, set the name of the remote endpoint address for both the near and far end switches to an empty string.

   set <service>/<instance> Plc remoteName “ “

Variable definitions

Variable	Value
<instance>	This value is a decimal from 1 to 65 535, which identifies the instance of the service on the switch. There is no default value.
<loc_addr>	This value is a string of 1 to 40 characters. The default is " ". This string is a valid NSAP address and represents the address for the switch.
<rem_addr>	This value is a string of 1 to 40 characters. The default is " ". This string is a valid NSAP address and represents the address for the switch at the remote end. This address is the destination relative to the near end switch.
<service>	This value is either HdlcTransparentDataService or BitTransparentDataService.
<value>	This value is a decimal from 1 to 65 535, which identifies the instance of the service on the switch. The default value is 2048.

Configuring a routing gateway

Configure a routing gateway to bridge calls from one topology region to another.

Note: Because the *way* component is memory and CPU intensive, allocate the component on a service function processor if one is available. When a *RoutingGateway* component is configured on a function processor that is dedicated to other tasks such as trunking, the *RoutingGateway* component is configured with a limit on the maximum number of calls supported. This call limit prevents the *RoutingGateway* component from exhausting all memory and CPU resources for the function processor.

Prerequisites

- “PORS across topology regions” (page 133)
- “Passport clusters” (page 140)

Procedure

1. Add a routing gateway

   ```
   add RoutingGateway/<instance>
   ```

2. Specify the logical processor on which the routing gateway resides.

   ```
   set RGty/<instance> logicalProcessor LogicalProcessor/
   <lp_n>
   ```

3. Specify the maximum number of calls for the routing gateway.

   ```
   set RGty/<instance> maxcalls <max>
   ```

Variable definitions

Variable	Value
<instance>	This value is a decimal from 1 to 65 535. There is no default value.

(Sheet 1 of 2)

Variable	Value
<lp_n>	This value is a decimal between 0 to 15, and identifies the logical processor number. There is no default value.
<max>	This value is a decimal from 0 to 1 000 000. The default value is 1 000 000.
(Sheet 2 of 2)	

Configuring inter-region Passport trunks

Configure inter-region Passport trunks with link-reachable addresses. In a multi-topology region network, switches at the edge of each topology region are configured with an inter-region Passport trunk to another topology region; and the addresses that can be reached through that inter-region Passport trunk.

Prerequisites

- “PORS across topology regions” (page 133)

Procedure

1. Add a trunk

   `add TRunk/<trk_n>`

2. Define an address that can be reached through that trunk

   `add Trunk/<trk_n> Address <address>`

3. Set the cost for the address.

   `set Trunk/<trk_n> Address <address> cost <cost>`

4. Repeat step 2 and step 3 for each address that is reachable through this trunk.

Variable definitions

Variable	Value
<address>	This value is a string of 1 to 40 characters, and is a valid NSAP address. This is no default value.
<cost>	This value is a decimal from 1 to 65 435. The default value is 200.
<trk_n>	This value is a decimal from 1 to 65 535. There is no default value.

Configure segmented PORS optimization

Specify whether each segment of a segmented PORS connection allows PORS optimization. You can configure segmented optimization at the source node, the outbound gateway node, and the inbound gateway node of a segmented PORS connection.

Prerequisites

- The *Routing Pors* component must be added. For supporting information, see "Path optimization" (page 115).

Procedure

1. Configure the *optimizationScope* attribute at the source node to enable optimization of segmented calls originated on the node:

   ```
   set Routing Pors optimizationScope applicationService
   ```

2. Invoke segmented PORS optimization on all segmented PORS connections on the source node:

   ```
   optimize -segmented rtg pors
   ```

3. Configure the *optimizationScope* attribute at a gateway node to enable optimization of outbound call segments:

   ```
   set Routing Pors optimizationScope outboundRGty
   ```

4. Invoke segmented PORS optimization on all segmented PORS connections on the outbound gateway node:

   ```
   optimize -segmented rtg pors
   ```

5. Configure the *optimizationScope* attribute at a gateway node to enable optimization of inbound call segments:

   ```
   set Routing Pors optimizationScope inboundRGty
   ```

6. Invoke segmented PORS optimization on all segmented PORS connections on the inbound gateway node:

   ```
   optimize -segmented rtg pors
   ```

7. To disable segmented optimization on a node using pling:

   ```
   set Routing Pors optimizationScope !
   ```

8. Optionally, configure the *optimizationScope* attribute to enable multiple options (for example, inboundRGty and outboundRGty options on a gateway node):

```
set Routing Pors optimizationScope ! outboundRGty
inboundRGty
```

Chapter 5
Monitoring and troubleshooting PORS

Monitor and troubleshoot the Path-Oriented Routing System using the procedures in this section.

Prerequisites to monitoring and troubleshooting PORS

- The PORS service must be initially installed and configured. See "Configuring PORS" (page 25), "Configuring ATM service over PORS" (page 43) and "Configuring PORS across topology regions" (page 49).

Monitoring and troubleshooting PORS task overview

This task flow show you the procedures you can perform to monitor and troubleshoot PORS. To link to any procedure, go to the list that follows the task overview.

Figure 6
Monitoring and troubleshooting PORS task flow

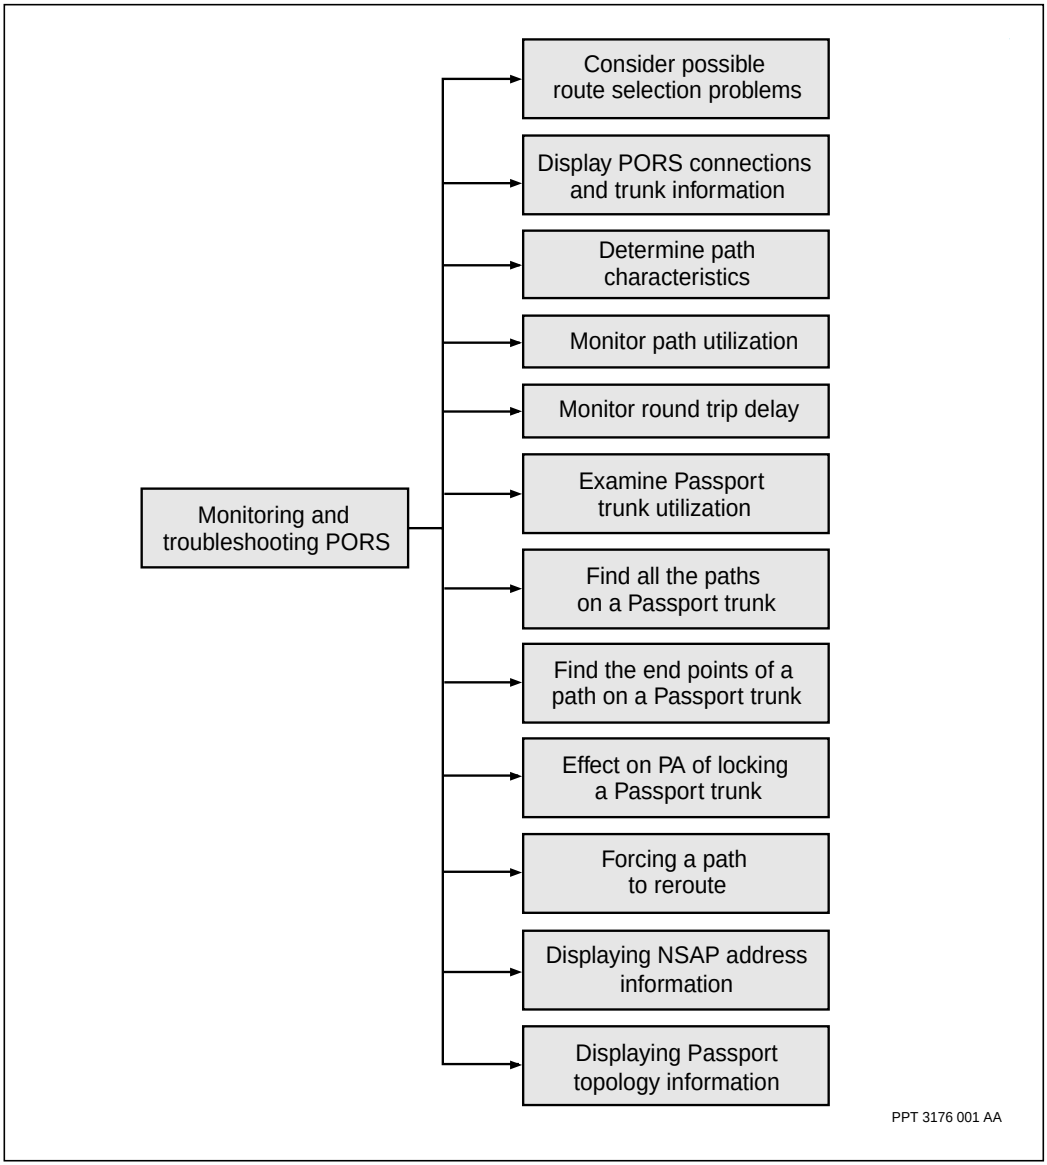

• “Possible route selection problems” (page 62)

Possible route selection problems

The route selection process can fail for the following reasons:

- There is no route that satisfies the transport characteristics.
- All available routes have a cost metric higher than the maximum cost metric and bumping is impossible.
- All available routes have a delay metric higher than the maximum delay metric and bumping is impossible.
- All available routes have number-of-hops larger than the maximum of ten and bumping is impossible.

Failure to select a route causes the logical connection and the call to retry periodically or to terminate if the path is rerouting due to a failure.

Display PORS connections and trunk information

Display a list of operational and configuration attributes for each PORS trunk going in the direction from the first node specified to the second node specified.

Procedure

1. Display information on the number of PORS connections on a switch and the status of optimization for all paths.

   display Routing PathOrientedRoutingSystem

   Information on the number of PORS connection and the optimization status appears. For example:

   ```
   Rtg PORS
   activeConnections = 7
   optimizationState = scheduled
   lastOptimizationTime = 200-04-01 20:06:42.56
   nextOptimizationTime = 2000-04-01 20:36:42.56
   optimizationPasses = 5
   optimizationProgress = 0%
   pathsOptimized = 3
   totalpathsOptimized = 4
   ```

2. Display data about PORS trunks in the network.

   display Routing Topology Node/<node_name1> LinkGroup/ <node_name2> PorsTrunkObject/*

Variable definitions

Variable	Value
<node_name1>	The name of the node from which the PORS trunks are outgoing.
<node_name2>	The name of the node to which the PORS trunks are going

Determining path characteristics

Examine information on path characteristics on a per logical channel basis. Information can be queried from both the trunk level and the service level. If all path characteristics are required, the end point of the path must be queried.

Procedure

1. Display information for a specific logical channel under a specific trunk.

   `display Trunk/<trk_n> LogicalChannnel/lc_n>`

2. From a service these characteristics can be queried by

   `display <servicetype>/st_n> LogicalConnection`

3. From a Passport trunk, all of these characteristics are be queried by displaying the appropriate *LogicalChannel* component. For example

   `display Trunk/<trk_n> LogicalChannel/lc_n>`

4. From a service these characteristics can be queried by

   `display <servicetype>/st_n> LogicalConnection`

Variable definitions

Variable	Value
<lc_n>	This value is a decimal from 0 to 65 535, representing the logical channel trunk instance. There is no default value.
<servicetype>	This value is one of VoiceService, BitTransparentDataService, HdlcTransparentDataService, or Aal1Ces
<st_n>	The instance identifier fro the service type.
<trk_n>	This value is a decimal from 0 to 65 535, representing the trunk instance. There is no default value.

Determining path utilization

Determining path utilization is not possible on a per logical channel basis.

Determining round trip delay

The round trip delay of a path must be queried from one of its end points. If you do not know the end points, see “Finding the end points of a path on a Passport trunk” (page 69).

Determining Passport trunk utilization

Determine how much path-oriented traffic a particular Passport trunk is forwarding. Statistics showing how many interrupting and non-interrupting packets were received as well as discarded are displayed.

Procedure

1 Display PORS statistics for the trunk.

```
display Trunk/<tnk_n> porsStatistics
```

Variable definitions

Variable	Value
<trk_n>	This value is a decimal from 0 to 65 535, representing the trunk instance. There is no default value.

Finding all the paths on a Passport trunk

List all logical channels on a trunk.

Procedure

1. List the logical channel for the *Trunk* component to display a summary listing

   `display Trunk/<trk_n> LogicalChannel/*`

2. Display the logical channels fro the *Trunk* component for a full listing.

   `display Trunk/<trk_n> LogicalChannel/*`

Variable definitions

Variable	Value
<trk_n>	This value is a decimal from 0 to 65 535, representing the trunk instance. There is no default value.

Finding the end points of a path on a Passport trunk

Determine the end points of a path.

Procedure

1. Display the next hop for

   `display Trunk/<trk_n> LogicalChannel/<lch_n> nextHop`

2. Repeat step 1, using the value of the nextHop attribute displayed through step 1 as the logical channel instance until the value of the nextHop attribute is a service component and not another logical channel instance.

Variable definitions

Variable	Value
<lch_n>	This value is a decimal from 1 to 65 535, representing the logical channel instance. There is no default value.
<trk_n>	This value is a decimal from 0 to 65 535, representing the trunk instance. There is no default value.

Effect on PA of locking a Passport trunk

When a *Trunk* component is locked, it enters the shutting down state and the path administrator of that Passport trunk no longer permits new paths to set up on the Passport trunk. All existing paths remain up but when they clear (of their own accord) they are no longer permitted back on this Passport trunk. Eventually there are no more logical channels on the Passport trunk at which time the Passport trunk immediately enters the locked state. Note that since the logical channels are not forced off the Passport trunk and the majority of the paths are permanent it can take a very long time before a shutting down Passport trunk enters the locked state. If you want to take the Passport trunk rapidly down you must issue a lock-force. However, lock-force disrupts service on all existing logical channels.

It is worth noting that while a Passport trunk in the shutting down state is still supporting logical connections the connectionless traffic can still flow unaffected over that Passport trunk. It is only when the last logical connection supported by that Passport trunk clears that the connectionless traffic is pushed off the Passport trunk and the Passport trunk enters the locked state.

Forcing a path to reroute

Trigger the rerouting of an established path which runs a specific service.

Procedure

1. Attempt to reroute the path. This command attempts to reroute the service path to a new path as long are the new path is not worse than the current path in terms of delay or cost, depending on whichever attribute the service is attempting to minimize.

   `reroute <servicetype>/<st_n> LogicalConnection`

2. If the reroute request in step 1 does not reroute the current connection, force the path to reroute using the clear command. This command interrupts in service connections and the reroute attempt occurs immediately.

   `clear Trunk/<trk-n> LogicalConnection/<lc_n>`

Variable definitions

Variable	Value
<lch_n>	This value is a decimal from 1 to 65 535, representing the logical channel instance. There is no default value.
<servicetype>	One of VoiceService, BitTransparentDataService, HdlcTransparentDataService, or Aal1Ces
<st_n>	The instance identifier for the service type.
<trk_n>	This value is a decimal from 0 to 65 535, representing the trunk instance. There is no default value.

Displaying NSAP address information

A component administration system (CAS) interface is available to display NSAP address information. While there is no configuration involved in setting up the CAS interface, an understanding of node addresses and reachable addresses is required.

Node NSAP addresses refer to the attributes *nodePrefix* and *alternatePorsPrefixes* that are provisioned under the *moduleData* component. Reachable NSAP addresses refer to the subcomponents *address* and *summaryAddress*, provisioned under the *Trunk* and *Routing* components respectively. Both of these NSAP addresses are broadcast and stored in the Base Routing Topology database of each node. Refer to the diagram "Node address and reachable address components and attributes" (page 73) for a visual representation of these attributes and subcomponents.

Figure 7
Node address and reachable address components and attributes

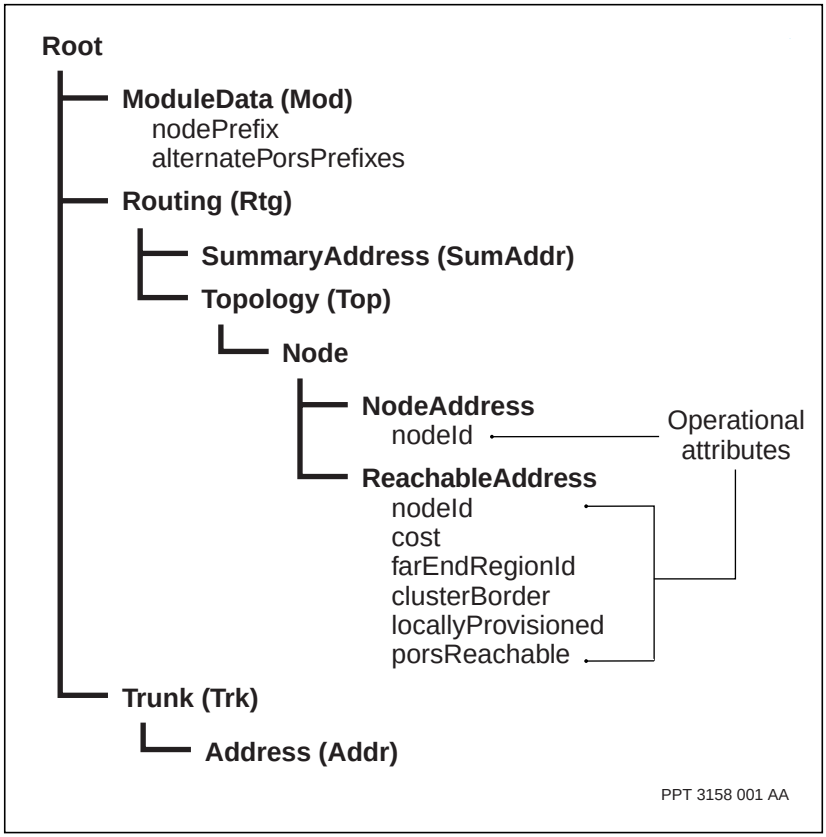

For detailed information on the operational attributes of node addresses and reachable addresses, refer to 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Node NSAP addresses

To display the node addresses stored in the Base Routing Topology database, enter either of the following commands:

```
display rtg top node/<x> nodeAddress/<y>
```

```
list rtg top node/<x> nodeAddress/<y>
```

In this case, <x> refers to the node stored in the Base Routing Topology database. If <x> is the local node, the displayed addresses <y> are provisioned locally. If <x> is a remote node, the displayed addresses <y> have been learned by topology broadcasts from its neighboring nodes.

The node address <y> is displayed on a local node if it is present in the local Base Routing Topology database. To be present in the local Base Routing Topology database:

- node <x> must be reachable from the local node
- on node <x>, at least one of the attributes *nodePrefix* or *alternatePorsPrefixes* must be provisioned as equal to <y>
- the local Base Routing Topology database must be up-to-date

Note: These commands can also be used on a network divided into topology regions; however, any command issued within a specific region will only display the node address contained within that region.

Reachable NSAP addresses

In instances where a network is divided into various topology regions, reachable addresses are used to establish connections from one region to another. Each node broadcasts the *address* and *summaryAddress* subcomponents to all nodes in the same region, storing this information in the local Base Routing Topology database of each node within the same region.

To display the reachable addresses stored in the Base Routing Topology database, enter one of the following commands:

```
display rtg top node/<x> reachableAddress/<y,n>
```

```
list rtg top node/<x> reachableAddress/<y,n>
```

In this case, <x> refers to the node stored in the Base Routing Topology database. If <x> is the local node, the displayed addresses <y> are provisioned locally. If <x> is a remote node, the displayed addresses <y> have been learned by topology broadcasts from its neighboring nodes.

If <x> is the local node, the value of <n> is the *linkId* of the trunk to which the address <y> belongs. If <x> is a remote node, the value of <n> is always 0. Though a reachable address <y> can be displayed several times on a local node with different non-zero values of <n>, it appears only once on a remote node, with the value of <n> always equal to 0.

The reachable address is displayed on a local node if it is present in the local Base Routing Topology database. To be present in the local Base Routing Topology database:

- node <x> must be reachable from the local node
- on node <x>, at least one of the subcomponents *address* or *summaryAddress* must be provisioned as equal to <y>
- the local Base Routing Topology database must be up to date

Note 1: These commands can also be used on a network divided into topology regions; however, any command issued within a specific region will only display the reachable addresses contained within that region.

Note 2: Reachable addresses are generally not used in flat networks.

Displaying Passport topology information

Determine the topology view of a backbone node and a cluster node.

Procedure

1. Refer to the table “Operator commands at a backbone node and cluster node” (page 76) for a summary of display and list operator commands to obtain information about a cluster node or a backbone node:

Table 1
Operator commands at a backbone node and cluster node

Operator command	Backbone node	Cluster node
l Rtg Top *	Lists all dynamic subcomponents of the *Topology* component. It shows all nodes in the backbone and cluster link groups staged on the local node.	Lists all dynamic subcomponents of the *Topology* component. It shows all nodes in the cluster and cluster link groups staged on the local node.
l Rtg Top Node/*	Shows the node names of all backbone nodes in the topology region.	Shows the node names of all cluster nodes in the cluster.
(Sheet 1 of 5)		

Table 1 (continued)
Operator commands at a backbone node and cluster node

Operator command	Backbone node	Cluster node
l Rtg Top Node/<x> *	Lists all dynamic subcomponents of the *Node* component (this includes the *LinkGroup*, *ClusterNode*, *Address*, and *ReachableAddress* components). The *LinkGroup* component exists if node x has at least one active neighbor node. A *ClusterNode* component exists if there is at least one cluster connected to node x. The *Address* component exists if there is a node prefix or an alternate PORS prefix provisioned on node x. The *ReachableAddress* component exists if there are trunk addresses or summary addresses provisioned on node x.	Lists all dynamic subcomponents of the *Node* component (this includes the *LinkGroup*, *BackboneNode*, Address, and ReachableAddress components). The *LinkGroup* component exists if node x has at least one active cluster node. A *BackboneNode* component exists when node x is connected to a backbone node. The *Address* component exists if there is a node prefix or an alternate PORS prefix provisioned on node x. The *ReachableAddress* component exists if there are trunk addresses or summary addresses provisioned on node x.
l Rtg Top Node/<x> ClusterNode/*	Lists the node names of all cluster nodes that can be reached through node x.	Shows an “invalid syntax” response.
l Rtg Top Node/<x> PeerBorderNode/*	Shows an “invalid syntax” response.	Lists the node names of all backbone border nodes that can be reached through node x.
d Rtg Top Node/<x> ClusterNode/<y>	Shows the delay and throughput metrics of the best path from node x to the cluster node.	Shows an “invalid syntax” response.
d Rtg Top Node/<x> PeerBorderNode/<z>	Shows an “invalid syntax” response.	Shows the delay and throughput metrics of the best path from node x to the backbone border node.

(Sheet 2 of 5)

Table 1 (continued)
Operator commands at a backbone node and cluster node

Operator command	Backbone node	Cluster node
l Rtg Top Node/* Addr/*	Lists all dynamic *Address* subcomponents. Shows all broadcast node prefixes and alternate PORS prefixes for all backbone nodes in the topology region.	Lists all dynamic *Address* subcomponents. Shows all broadcast node prefixes and alternate PORS prefixes for all cluster nodes in the cluster.
d Rtg Top Node/* Addr/*	Shows the attributes of the dynamic *Address* subcomponent, for all backbone nodes in the topology region.	Shows the attributes of the dynamic *Address* subcomponent, for all cluster nodes in the cluster.
l Rtg Top Node/* Raddr/*	Lists all dynamic *ReachableAddress* subcomponents. The subcomponents represent non-summarized *Address* and *SummaryAddress* components provisioned for all backbone nodes in the topology region.	Lists all dynamic *ReachableAddress* subcomponents. The subcomponents represent non-summarized *Address* and *SummaryAddress* components provisioned for all cluster nodes in the cluster.
d Rtg Top Node/* Raddr/*	Shows the attributes of the dynamic *ReachableAddress* subcomponent for all backbone nodes in the topology region.	Shows the attributes of the dynamic *ReachableAddress* subcomponent for all cluster nodes in the cluster.
l Rtg Top Node/<x> Addr/*	Lists all dynamic *Address* subcomponents on node x. The subcomponents represent the node prefixes and alternate PORS prefixes provisioned on node x.	Lists all dynamic *Address* subcomponents on node x. The subcomponents represent the node prefixes and alternate PORS prefixes provisioned on node x.
d Rtg Top Node/<x> Addr/*	Shows the attributes of the dynamic *Address* subcomponents on node x.	Shows the attributes of the dynamic *Address* subcomponents on node x.

(Sheet 3 of 5)

Table 1 (continued)
Operator commands at a backbone node and cluster node

Operator command	Backbone node	Cluster node
l Rtg Top Node/<x> Raddr/*	Lists all dynamic *ReachableAddress* subcomponents on node x. The subcomponents represent non-summarized *Address* and *SummaryAddress* components provisioned on node x.	Lists all dynamic *ReachableAddress* subcomponents on node x. The subcomponents represent non-summarized *Address* and *SummaryAddress* components provisioned on node x.
d Rtg Top Node/<x> Raddr/*	Shows the attributes of the dynamic *ReachableAddress* subcomponent on node x.	Shows the attributes of the dynamic *ReachableAddress* subcomponent on node x.
l Rtg Top Node/* Lg/*	Lists all dynamic *LinkGroup* subcomponents. The subcomponent represents all the dynamic link groups that are present in the topology region backbone.	Lists all dynamic *LinkGroup* subcomponents. The subcomponent represents all the dynamic link groups that are present in the cluster.
l Rtg Top Node/* Lg/* PorTrunkObject/*	Lists all dynamic *PorTrunkObject* subcomponents. The subcomponent represents PORS trunks provisioned for all backbone nodes in the topology region.	Lists all dynamic *PorTrunkObject* subcomponents. The subcomponent represents PORS trunks provisioned for all cluster nodes in the cluster.
d Rtg Top Node/* Lg/* PorTrunkObject/*	Shows the characteristics and operational attributes of all PORS trunks in the topology region backbone.	Shows the characteristics and operational attributes of all PORS trunks in the cluster.
l Rtg Top Blg/*	Shows all active cluster link groups staged on the local node.	Shows all active cluster link groups staged on the local node.
d Rtg Top Blg/<node_name>	Shows the delay and throughput metrics as well as the LNNs supported by the cluster link group.	Shows the delay and throughput metrics as well as the LNNs supported by the cluster link group.

(Sheet 4 of 5)

Table 1 (continued)
Operator commands at a backbone node and cluster node

Operator command	Backbone node	Cluster node
l Rtg Top Blg/<node_name> *	Shows all dynamic subcomponents of the cluster link group (this includes the *PorTrunkObject* and *ClusterNode* components). The *PorTrunkObject* component exists if there is any PA configured in the cluster link group. The *ClusterNode* component exists if there is at least one cluster node reachable through the cluster link group.	Shows all dynamic subcomponents of the cluster link group (this includes the *PorTrunkObject* component). The *PorTrunkObject* component exists if there is any PA configured in the cluster link group.
l Rtg Top Blg/<node_name> PorTrunkObject/*	Shows all PORS trunks provisioned in the cluster link group.	Shows all PORS trunks provisioned in the cluster link group.
d Rtg Top Blg/<node_name> PorTrunkObject/<link_id>	Shows the characteristics and operational attributes of the specified PORS trunk.	Shows the characteristics and operational attributes of the specified PORS trunk.
l Rtg Top Blg/<node_name> ClusterNode/*	Shows all cluster nodes that are reachable through the cluster link group.	Shows an “invalid syntax” response.
d Rtg Top Blg/<node_name> ClusterNode/<y>	Shows the delay and throughput metrics from the border node to reach the cluster node x through the specified cluster link group.	Shows an “invalid syntax” response.

(Sheet 5 of 5)

Variable definitions

Variable	Value
<link_id>	The instance value of the *PorTrunkObject* component.
<node_name>	The instance value of the *Blg* component. This value refers to a specified cluster link group.
<x>	The instance value of the *Node* component.
<y>	The instance value of the *ClusterNode* component.
<z>	The instance value of the *PeerBorderNode* component.

Procedure job aid

Figure 8
Displaying Passport cluster information component hierarchy

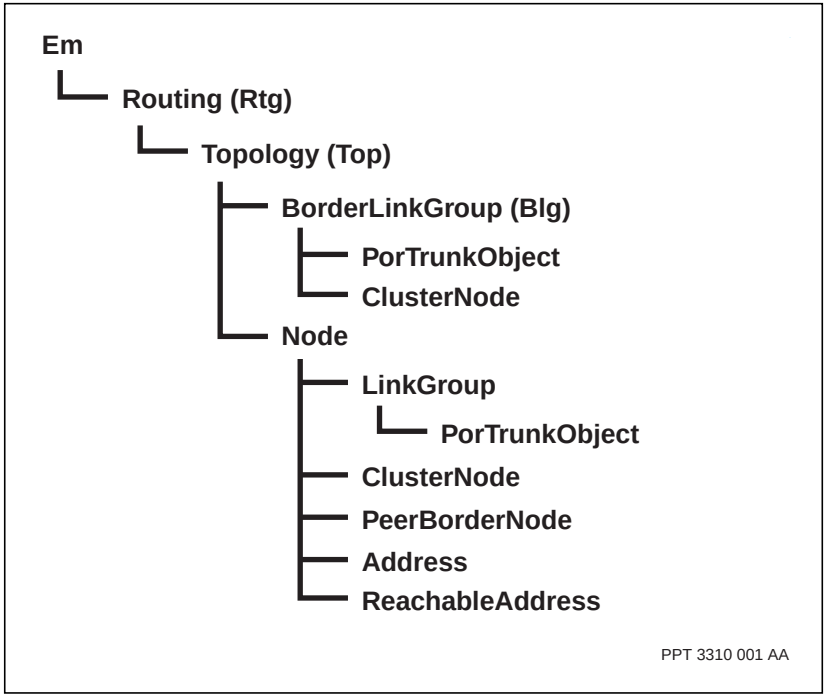

Chapter 6
Introduction to Path-Oriented Routing System

This section introduces conceptual and supporting information about the Path-Oriented Routing System. For an overview of the procedures relating to this service, see “PORS task flow overview” (page 21).

The Passport Path-Oriented Routing System (PORS) is a connection-oriented routing system that runs on Passport 7400, Passport 15000 and 20000 switches. PORS establishes and maintains an end-to-end digital transmission path for both voice and digitized data on a trunk-by-trunk basis. PORS provides ordered transmission of data with minimal delay variance. Examples of non-voice data include HDLC and any bit-stream data.

PORS software is required on all Passport switches in the network that support transport services.

This section provides an overview of the Path-Oriented Routing System (PORS), its features, and how it works. Overview information is organized under the following headings:

Overview

PORS supports both switched logical connections (SLC) and permanent logical connections (PLC). The use of an SLC or a PLC depends on the application that uses PORS. When a call is placed, a PLC or an SLC, and a route are created between the two end points (source and destination) of the

call. Either end point can initially try to establish a path. The first end point to establish a path becomes the calling end (source), the remote end point becomes the called end (destination).

The path is created by:

- selecting a route that satisfies the transport characteristics
- instantiating the route by:
  - — verifying that the required bandwidth is available on all Passport trunks
  - — reserving the bandwidth on all Passport trunks
  - — setting up the *LogicalChannel* components required on all Passport trunks

Upon completion of these steps, a logical connection and path are established. It is up to the access service, for example Voice Transport, to use the path.

Services that use PORS

PORS supports the following services:

- Voice Transport (see 241-7401-750 *Passport 7400 Voice Transport Guide*)
- HDLC transparent data service (see 241-7401-770 *Passport 7400 HDLC Transparent Data Service Guide*)
- bit transparent data service (see 241-7401-775 *Passport 7400 Bit Transparent Data Service Guide*)
- AAL1 circuit emulation service (see 241-5701-720 *Passport 7400, 15000, 20000 AAL1 Circuit Emulation Guide*)

PORS can also be used to transport frame relay data using multiservice cut-through switching (MCS). MCS provides many-to-one multiplexing of connections for data services and uses PORS facilities to route the traffic. For more information on MCS, see 241-7401-440 *Passport 7400 Frame Relay Managed Cut-through Switching Guide*.

Features and characteristics

PORS features and characteristics are presented under the following headings:

- "Summary of PORS transport characteristics" (page 85)
- "Summary of packet management features" (page 86)
- "Summary of bandwidth and congestion management features" (page 87)

Summary of PORS transport characteristics

The selection of a connection path is based on criteria that is set by a network operator. There is no need to manually select a route for every connection, although manual configuration is available. PORS preferentially selects paths that minimize cost or delay, within the bounds of a maximum cost or delay metric. The routing mechanism at each hop consists of a table lookup using the connection identifier as an index. A connection is re-routed if a component on the path fails.

The following transport characteristics of a route are configurable:

- the bandwidth required in each direction of the route
- the priority at which the route is set up and the priority it holds after being set up
- the security level of the route
- the type of Passport trunks that can be used on the route (such as terrestrial, satellite)
- the type of traffic carried on the path (such as voice, data, video)
- the optional customer defined parameter supported on the Passport trunks of the path
- the primary route metric to minimize: cost or delay
- the maximum cost metric that a path can tolerate; calculated from a value assigned to Passport trunking facilities by a network administration
- the maximum delay metric that a path can tolerate; calculated from the transmission delay of the Passport trunking facilities

- NSAP link-reachable addresses and their costs (for routing across topology regions)
- maximum region cost

An operator can manually select a route where all of the configured PORS characteristics must still be met by the manually selected Passport trunks before path is instantiated.

An operator can also force a path to follow the manually selected route, regardless of path characteristics. A forced manual path still reserves the bandwidth on the Passport trunk; if not enough bandwidth is available, it reserves only up to the maximum available limit on that Passport trunk.

For information on the PORS addressing system, see 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*.

Summary of packet management features

The key features of PORS packet management are:

- Identical forward and backward data path facilitates the exchange of network management information between all points on the route.
- Periodic optimization of a PORS path provides network efficiency.
- Minimal delay variance exists between packets. All packets follow the same route.
- Preservation of packet ordering prevents a packet from overtaking another.
- Duplication (retransmission) of packets is avoided.
- Packet emission priority allows packets that are sensitive to delays to be emitted from a Passport trunk urgently.
- Packet discard priority allows less important packets to be discarded during congestion.
- Distribution of congestion information to the service end points is done in the packets.
- Best effort delivery is provided (but packet loss can occur during periods of congestion).

- Loadspreading allows traffic to be spread across link groups.
- Network partitioning allows autonomous topology regions.
- Tandem suppression allows the network operator to control routing behavior by preventing tandem traffic from traversing the selected nodes.

For Passport 15000 and 20000, PORS can be configured as a warm standby feature. A warm standby application or feature can operate together with a hot standby application or feature on the same function processor without affecting the ability of the hot standby application or feature to provide hitless services during an equipment switchover.

See 241-5701-600 *Passport 7400, 15000, 20000 Configuration Guide* for a description of hitless services and hot, warm and cold standby applications and features.

Summary of bandwidth and congestion management features

PORS allows a network administration to efficiently manage the bandwidth on Passport trunking facilities. The key features of PORS bandwidth management are:

- Bandwidth for PORS traffic is reserved on all Passport trunks on the route. The reserved portion can be set to the service peak, average, or any other arbitrary bandwidth usage that the service requires.

  Bandwidth can be allocated in a statistical fashion by reserving only a portion of the bandwidth required for the connection, or in a conservative fashion by reserving the peak bandwidth requirements.

- Bandwidth on a Passport trunk is shared between connectionless and connection-oriented traffic. Bandwidth that one traffic type does not use is unused by one traffic type can be used by the other.

  The network operator can configure the maximum amount of bandwidth allowed for reservation by path-oriented connections.

- Bandwidth on a Passport trunk is statistically shared amongst all of the paths.
- A Passport trunk is not used as part of a path if it exceeds the Passport trunks total reserved bandwidth.

- Bandwidth overhead is low. The packet headers contain a route identifier rather than a full destination address.
- PORS uses both Passport frame-cell trunks and Passport trunks over ATM. There are three modes of operation for PORS traffic on Passport trunks over ATM: AAL5-mux mode, SPO-mux mode and map mode (see 241-5701-420 *Passport 7400, 15000, 20000 Trunking Guide*).

For Passport 15000 and 20000, PORS trunking can be configured as a warm standby feature. However, PORS routing is interrupted during an equipment switchover and becomes available after:

- the switchover to the standby function processor is complete
- Passport trunks over ATM are re-established

Chapter 7
Routing fundamentals

This section describes the fundamentals of routing for the Passport Path-Oriented Routing System (PORS). Information is organized under the following headings:

- “Packet transport characteristics” (page 89)
- “Introduction to route selection” (page 98)
- “Route selection attributes” (page 99)
- “Instantiating the route” (page 109)
- “Rerouting paths” (page 112)
- “Transporting data on the path” (page 118)
- “Terminating a route” (page 119)

For information on specific applications of PORS outside of the default implementation for services, see “Routing applications” (page 121).

Packet transport characteristics

Information on packet transport characteristics is organized under the following headings:

- “A path as a sequence of logical channels” (page 90)
- “Path-oriented packet forwarding (inbound trunk)” (page 91)
- “Path-oriented packet header” (page 93)
- “Path-oriented data mode” (page 95)

- “Path-oriented inband control mode” (page 95)
- “Path-oriented outband control mode” (page 95)
- “Importance and urgency of packets” (page 96)
- “Guidelines for setting emission and discard priorities” (page 96)

A path as a sequence of logical channels

The figure “Example of a logical connection and a route: topology perspective” (page 91) shows a connection between nodes A and D from the perspective of network topology. The Passport route selector defines the connection as

```
EM/A Trunk/4 -> EM/B Trunk/6 -> EM/C Trunk/5
```

The figure “Example of a logical connection and a route: logical perspective” (page 91) shows the same connection from a logical perspective. The connection has the following characteristics:

- the connection consists of three hops: EM/A Trunk/4 over LogicalChannel/25, EM/B Trunk/6 over LogicalChannel/7, and EM/C Trunk/5 over LogicalChannel/12
- The route selector chooses the route from the Passport trunks but does not select the logical channel numbers (LCN) at each hop
- LCNs are allocated during path instantiation in a Passport network
- an LCN is local to a Passport trunk and is generally different at each Passport trunk of the route

Figure 9
Example of a logical connection and a route: topology perspective

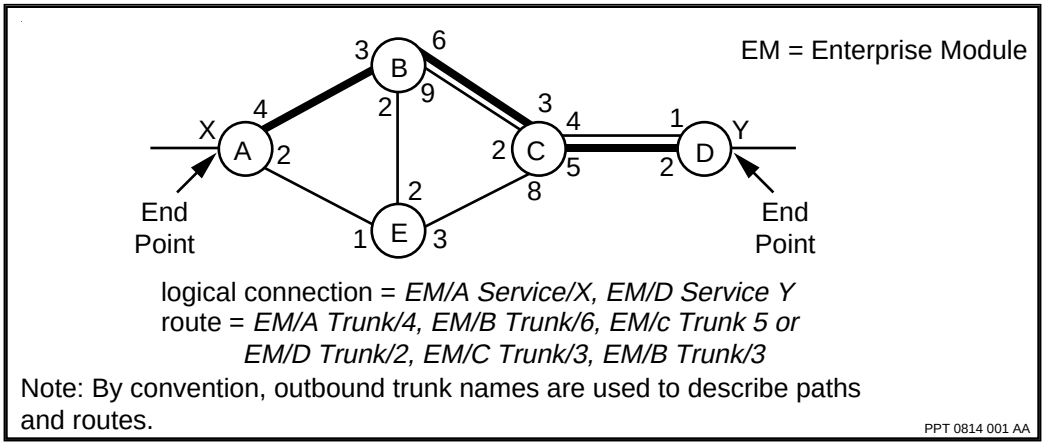

Figure 10
Example of a logical connection and a route: logical perspective

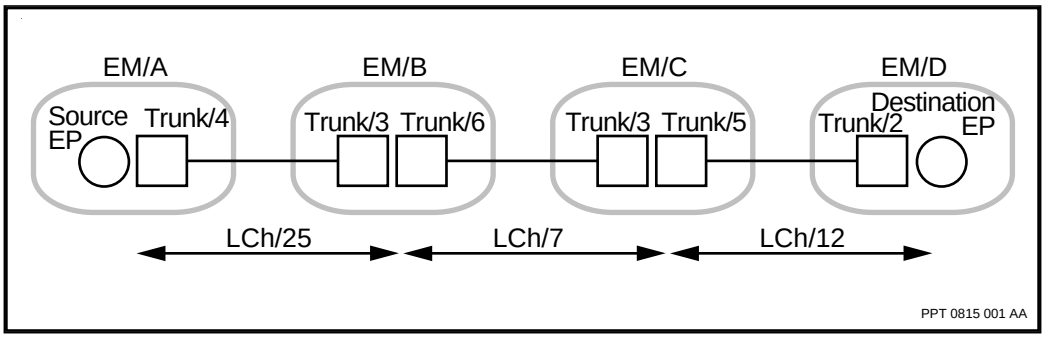

Path-oriented packet forwarding (inbound trunk)

When a path-oriented packet arrives at an inbound Passport trunk for forwarding to the next outbound trunk, PORS packet forwarding uses the inbound logical channel to determine the outbound logical channel. The figure “Example of the path-oriented forwarding (POFWD) tables along a route” (page 92) shows the logical connection from nodes A to node D. The packet forwarding (POFWD) table for the trunk stores the information necessary for the forward and the reverse data flow.

Figure 11
Example of the path-oriented forwarding (POFWD) tables along a route

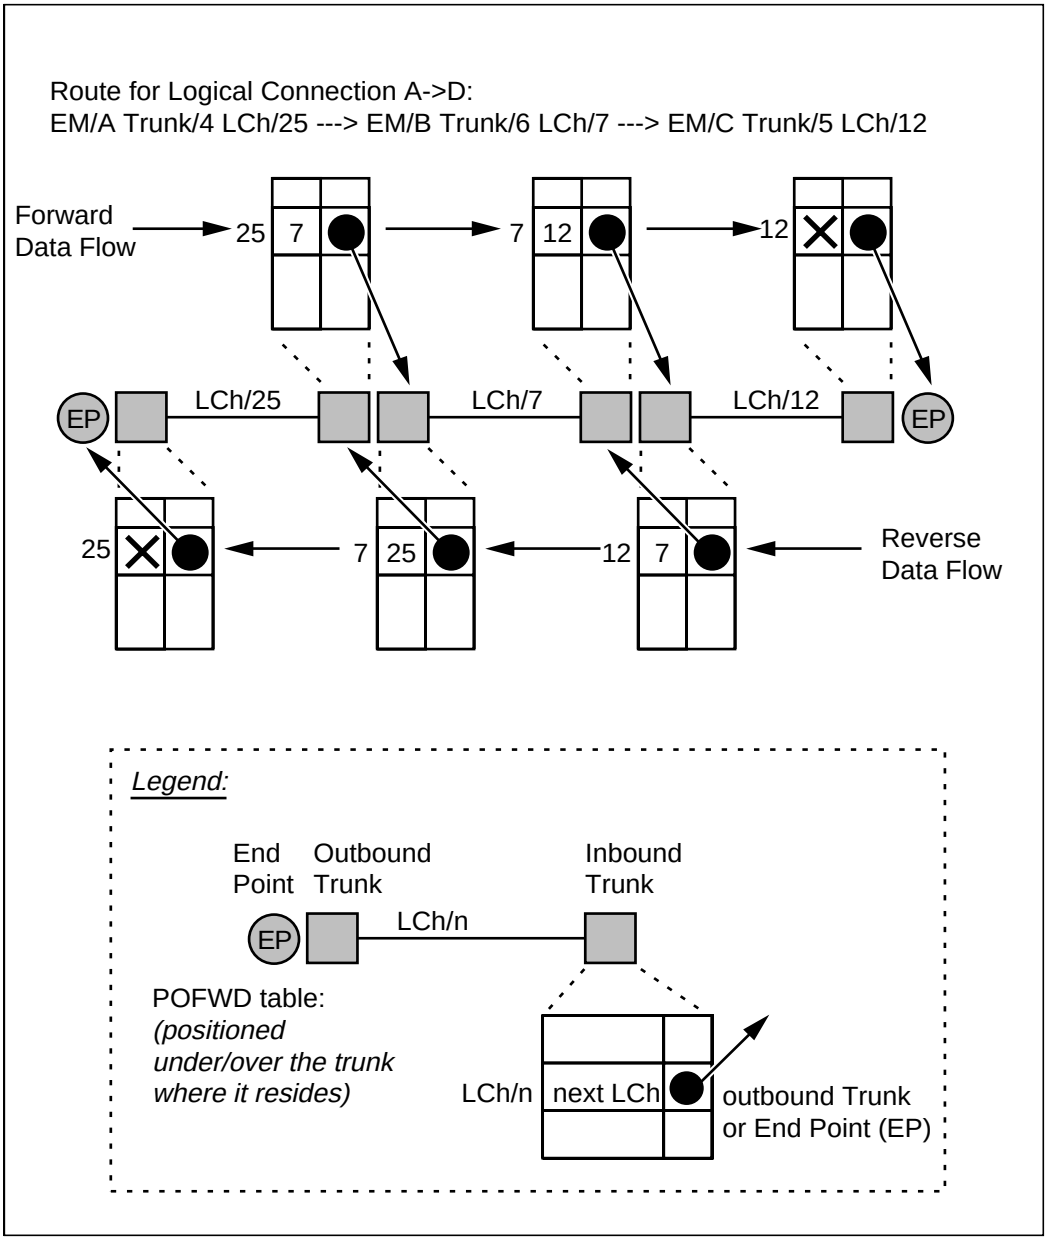

PORS uses the LCN of the packet to index into the POFWD table on the inbound trunk to determine the outbound Passport trunk and next LCN to use. The reverse data flow is much the same. The LCNs used in both directions of a path are identical.

Path-oriented packet header

Path-oriented routing protocol data units (PDU) have a routing common header followed by an LCN followed by the data (packet payload). Two sizes are available for LCN resulting in two general formats for the path-oriented routing PDUs. These headers are shown in the figure "Path-oriented routing protocol data unit" (page 93). All three routing modes of path-oriented routing PDUs share the same format. See

- "Path-oriented data mode" (page 95)
- "Path-oriented inband control mode" (page 95)
- "Path-oriented outband control mode" (page 95)

Note: Traffic from ATM services, such as circuit emulation service (CES), does not need a PORS header. This type of traffic is routed with an ATM header only.

Figure 12
Path-oriented routing protocol data unit

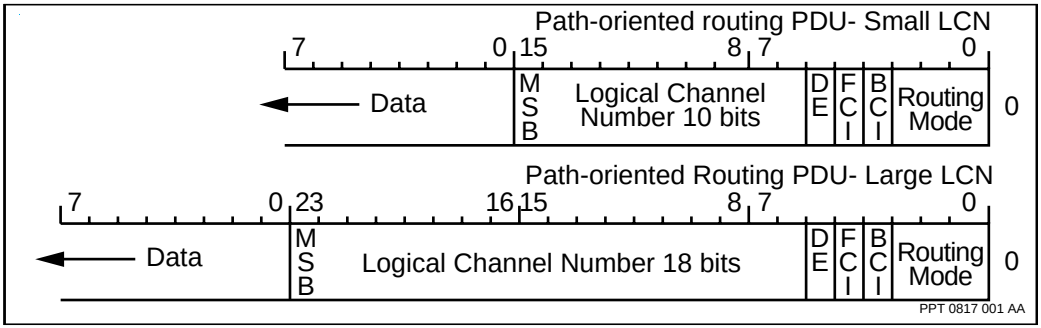

The table "Path-oriented routing PDU protocol data elements" (page 94) describes the contents of the PDUs. The table "Summary of path-oriented routing modes" (page 95) summarizes characteristics of the routing modes.

Table 2
Path-oriented routing PDU protocol data elements

Field	Meaning
Routing mode	The following routing modes define path-oriented routing PDUs: • 000 - path-oriented data • 010 - path-oriented inband control • 100 - path-oriented outband control These Routing Modes are described in more detail later.
FCI	On a received packet, the presence of the Forward Congestion Indication bit indicates that this packet has encountered congestion on travels across the network.
BCI	On a received packet, the presence of Backward Congestion Indication bit indicates that packets being sent by this end are encountering congestion as they cross the network.
DE	The Discard Eligibility bit is set on a packet to cause the packet to be treated with lowest possible discard priority whenever discard priority decisions are made. Packets with this bit set are the first candidates for discard.
LCN	The LCN associates this PDU with a particular logical channel. LCN 0 is a special logical channel used for path management on a Passport trunk. All other logical channels are available for data flow. Two ranges are available for LCNs: small, 1-1023 (10 bits); and large, 1-262144 (18 bits) depending on the needs of the facility. (Actually, nothing prevents future faster facilities from having larger LCNs if they want to spend the bits). Facilities operating at DS1, E1 speeds, and lower use the small logical channel range to conserve bandwidth. Facilities operating above these speeds use the large LCNs to permit greater channel handling capacity. For convenience, in the rest of this document these are referred to as small and large Passport trunks, respectively. The size of the LCN is a local issue for the facility.

Table 3
Summary of path-oriented routing modes

Characteristic	Path-oriented data	Path-oriented inband control	Path-oriented outband Control
Discard priority	From connection table	0	0
Emission priority	From connection table	Same as path-oriented data	1
Packet size	Selected by service	Less or equal path-oriented data	Any

Path-oriented data mode

Packets with this mode follow the path using the emission and discard priority associated with the path during path establishment. The LCN field contains the LCN of the logical channel for this hop of the path and is changed for each hop. The packet is delivered to the service at the far end. This mode is used exclusively for VC data packets which are being transported on behalf of the service.

Path-oriented inband control mode

Packets with this mode are handled along the path in exactly the same fashion as path-oriented data except that at the destination they are extracted from the data stream and delivered to the PEP. Round-trip delay measuring packets use this routing mode so that they follow exactly the same path as the data and therefore return an accurate measure. These packets must meet the same size criteria as the data so that they do not upset queuing and congestion management, especially when they go through the interrupting queue.

Path-oriented outband control mode

These packets come in two types: those with an LCN of 0 and those without. LCN 0 packets go hop-by-hop and are delivered to the TPA at each hop. Path setup packets are an example of this kind of packet. LCN 0 is also used for all TPA management packets which cross a Passport trunk. Non-LCN 0 packets follow the same path as the data packets except that they are always treated as Emission Priority 1 (high, not interrupting) and Discard Priority 0 (high). This allows the VC to send larger control packets on this mode since these packets do not end up on the interrupting queue.

Importance and urgency of packets

The next logical channel and outbound Passport trunk fields in the forwarding table are sufficient to forward a packet along a fixed route. Other transport properties include the importance of delivery and the urgency of emission. Urgency and importance define the quality of service (QOS) that PORS provides on the path. The path-oriented header does not carry fields to indicate packet importance and urgency since these attributes are defined at the path level.

Packet importance is defined as discard priority, which implies the likelihood of the packet reaching its destination (the probability of loss or discard). Paths can contain important data which, if lost, can cause problems for a service or customer. Important packets are given a high guarantee of delivery and are therefore less likely to be discarded in the event of network problems or congestion. The discard priority defined for a path indicates the importance of all the packets on the path. Independently of the path discard priority, individual packets can be marked as the least important packets by setting the discard eligible (DE) bit in the packet header.

Packet urgency defines emission priority, which implies how quickly a packet is forwarded to its destination. Urgent traffic can consist of packets that must be delivered within a specific time frame for the application to function correctly. Packets of paths with a high emission priority are queued on a Passport trunk queue that is serviced before other Passport trunk queues.

Guidelines for setting emission and discard priorities

The purpose of having four different emission and discard priorities is to allow the network traffic to be partitioned into groups requiring similar quality of service. In this manner, traffic requiring better service can be dealt with more promptly and/or with less loss than a service requiring a lesser QOS. Normally, an access service is aware of the QOS required for its path and automatically sets the attribute to the desired values. However, some applications allow one or both attributes to be adjusted.

For each type of priority there are two extremes, either all traffic has the same value, or traffic is distributes among the priorities. See the table “Four extremes of traffic assignment to emission and discard priority” (page 97).

Table 4
Four extremes of traffic assignment to emission and discard priority

Attribute	All traffic at same priority	Traffic distributed throughout all levels
discard priority	All traffic is considered equally important. Under congestion, the decision of which packet to discard is randomly distributed amongst all packets.	Traffic has four distinct levels of importance. Under congestion, DE priority packets are discarded first, then priority 3, then 2, then 1 and finally 0.
emission priority	All traffic is considered equally urgent and is emitted on a first come first serve basis. When congested, the Passport trunk emission delay increases for all packets.	Traffic has three distinct levels of urgency. When congested, the first packets emitted are those at priority 0 followed by priority 1 and finally priority 2.

Taking into consideration the fact that high urgency voice and bit-sliced transparent data must run at the highest emission priority, use the following guidelines:

- If all traffic is equally important, use a middle discard priority (that is, 2) for all traffic.

  This approach permits raising or lowering the priority of individual traffic sources later to increase or decrease their relative importance.

- If there are clearly different levels of importance, the priority levels must be spread evenly among the different kinds of traffic.

  This approach permits the network to make the most precise choices about packets to throw away at congested points. Consider the tolerance to loss of the individual service when setting the discard priority. If the service can tolerate small amounts of loss, do not use discard priority 0. Use discard priority 0 for services that are seriously degraded by a lost packet.

- If all traffic is equally urgent, use a middle emission priority (that is, 1) for all traffic except voice and bit transparent data.

  This approach permits raising or lowering the priority of individual traffic sources later to increase or decrease their relative urgency.

- If all traffic is equally urgent, use emission priority 0 for voice and bit-sliced transparent data only.
- If there are clearly different levels of urgency, use emission priority 0 for voice and bit-sliced transparent data. Evenly apply the remaining priority levels to the remaining services.

Introduction to route selection

In a well-engineered network, there can be several routes between two nodes. A route selector chooses the route to the destination end point based on the transport characteristics configured for the permanent logical connection.

Route selection is made using a modified Dijkstra shortest path routing algorithm and the same topology database as is used by Passport connectionless routing.

The route selector prunes from the topology the Passport trunks that do not conform to the transport characteristics and solves the shortest route problem on the remaining Passport trunks, using cost or delay characteristics as the criteria.

The route selector requires an accurate view of the bandwidth available on each Passport trunk in the network to select the optimum route. It learns this information from

- regular topology updates
- reports from path setups that have failed or cannot be instantiated
- reports from local Passport trunk path administrators (PAs) about unreserved PORS bandwidth on outgoing Passport trunks
- reports from distant PAs about significant bandwidth releases on distant Passport trunks

The description of route selected by the route selector is returned to the end point.

Route selection attributes

A Passport trunk is a candidate for route selection if its configuration agrees with the values configured for the logical connection.

The following transport characteristics are defined for each logical connection. Most characteristics apply to both switched and permanent connections except where noted in the following descriptions. The transport characteristics marked with an asterisk (*) are not used during route selection.

- “Destination application name” (page 100)
- “Setup and holding priorities” (page 100) (holding priorities only*)
- “Emission and discard priority” (page 101)*
- “Required bandwidth” (page 102)
- “Security” (page 102)
- “Traffic type” (page 103)
- “Passport trunk type” (page 103)
- “Customer defined parameter” (page 104)
- “Minimization criteria” (page 105)
- “Maximum acceptable cost” (page 105)
- “Maximum acceptable delay” (page 106)
- “Manual path” (page 106)*
- “Manual path with forced bandwidth” (page 107)*
- “Terminate when rerouting” (page 107)*
- “Optimization” (page 107)*
- “Bump preference” (page 108)

For ATM services, these attributes (except destination application name) can be specified in a PORS routing profile instead for the PLC.

Destination application name

The application name is the PLC remote end point address in the format of a Passport component name. This name is specified by the *remoteName* attribute located under the *PermanentLogicalConnection* component. End points are identified using a node name and a service name. If the service uses data network address (DNA) or network service access point (NSAP) addressing, PORS converts the address to a node ID. If the *remoteName* attribute is configured, it must match the remote name of the other end or the connection is not be established even if the remote end is configured correctly.

The *overrideRemoteName* attribute located under the *LogicalConnection* component can also be used to change the remote end points address. If this attribute is set, and the *remoteName* attribute is blank, then the remote end points address specified by the *overrideRemoteName* attribute overrides the blank PLC *remoteName* attribute. The advantage of using the *overrideRemoteName* attribute is that the user does not use configuration to change the remote end point address. The disadvantage of this attribute is that the *overrideRemoteName* attribute is not permanent so in the event of a resetting of the function processor or node, this override setting is lost.

Note 1: To use the *overrideRemoteName* attribute, the user configures the new remote end points address using the *overrideRemoteName* attribute.

Note 2: The *override* attribute cannot be used with non-PLC based PORS services, such as AAL1 CES.

Setup and holding priorities

PORS reserves the bandwidth required by a path on each Passport trunk of the route. If a route with sufficient bandwidth cannot be found, existing paths can be rerouted to reallocate the bandwidth to the new path. This is the process of bumping paths. Setup and holding priorities are used to rank existing paths (holding priority) and the new path (setup priority) to determine if the new path can bump an existing path.

The setupPriority of a new *PermanentLogicalConnection* component and the *holdingPriority* attributes of the *PermanentLogicalConnection* components for existing paths are used to specify these priorities. The higher the holding

priority, the less likely it is for PORS to reallocate its bandwidth to a new path. Similarly, the higher the setup priority, the more paths PORS can bump to set up the path. See “Bumping paths” (page 112) for more details.

The setup and holding priority values range from zero (0) to four (4). The value zero (0) is the priority assigned to the most important path. It is referred to as the highest priority. Four (4) is the priority for the least important path.

An existing path can be bumped if and only if the value for the *setupPriority* attribute of the new path is numerically less than the value for the *holdingPriority* attribute of the existing path.

To illustrate the use of the setup and holding priority, consider a network which supports both voice and data services in its corporate network. The voice traffic is given a low setup priority because new voice paths can use the public network if the corporate network cannot accommodate the new path. However, the voice traffic is given a high holding priority since it is undesirable for the path to be rerouted during an active voice call. For data traffic, high setup and holding priorities are desirable since data paths cannot be established on an alternate network.

Emission and discard priority

The emission priority is a measure of how urgently a packet is emitted on the Passport trunk. There are three possible values, from zero (0) to two (2). Zero (0) is the highest priority, or the fastest emission, and two (2) is the lowest priority, or the slowest emission. Zero uses the interrupting queue on Passport trunks.

The discard priority is a measure of the importance of a packet. It is used to decide on which packets are discarded first during congestion. There are four discard levels: zero (0) is assigned for paths whose packets are the last ones to be discarded, and three (3) is assigned to paths whose packets can be discarded first. Discard level 0 is reserved for control packets.

Emission and discard priorities are set independently. Taking a voice path as an example, it has a high emission priority because a low delay must be experienced on the path, but it has a low discard priority since a small number of discarded packets do not significantly impact the quality of a voice call.

The emission and discard priorities are transport characteristics of a path. These are not used for route selection. “Importance and urgency of packets” (page 96) discusses these transport characteristics in detail.

Required bandwidth

PORS reserves the required bandwidth for a path on every Passport trunk on the route and in each direction of transmission. The required bandwidth is an indication of the potential bandwidth utilization of the path. It is not enforced by PORS, and is only used as an indication of the number and size of paths that can be set up on a Passport trunk.

The *PermanentLogicalConnection* component has the *requiredTxBandwidth* and *requiredRxBandwidth* attributes to indicate the bandwidth to reserve in transmit and receive directions. It is an ordered pair that specifies the bandwidth required in bits/second in the outgoing and incoming direction (that is, to and from the far end point respectively). This is the minimum bandwidth that PORS must provide in order to set up the path. The required bandwidth attributes must take into account both the payload and the packet overhead.

Bandwidth can be determined automatically by the application and signalled to PORS.

If enough free bandwidth is not available, path bumping is required. In this case, the bandwidths of enough existing paths with a holding priority less than the setup priority of the new path is made available for the new path. Bumping is possible if there is enough bandwidth currently reserved for lower holding priority paths.

Security

The security option specifies the minimum security level required for transporting the packets for the service. PORS supports eight possible security levels, between zero (0) and seven (7). A value of zero (0) indicates the highest security, while seven (7) indicates the lowest.

A Passport trunk is a candidate for route selection if it provides at least the required path security. PORS can use a more secure Passport trunk than the one requested.

The security level is specified in the *requiredSecurity* attribute of the *PermanentLogicalConnection* component. A corresponding attribute exists on Passport trunks, called *trunkSecurity*. A Passport trunk is chosen for route selection, if the numeric value of the *trunkSecurity* attribute is equal or less than the numeric value of the *requiredSecurity* attribute.

Traffic type

The traffic type is used to allow Passport trunks on a route that support a specific traffic type. During route selection, potential candidate Passport trunks for a route are those that support the required traffic type. An administration can use this attribute to prevent a traffic type from using certain Passport trunks, or to reserve Passport trunks for one traffic type.

The *PermanentLogicalConnection* component has an attribute called the *requiredTrafficType*, which specifies the traffic type transported on the path. There is a corresponding attribute on Passport trunks, called the supportedTrafficType, which is a set of traffic types that can be carried on the Passport trunk.

A Passport trunk can be used for route selection, if the *requiredTrafficType* attribute value for the path is a member of the *supportedTrafficTypes* attribute on the Passport trunk.

PORS allows eight possible traffic types for both the *requiredTrafficType* and *supportedTrafficTypes* attributes.

As an example, voice, data and video are possible traffic types supported by PORS. If voice or data are specified in the *requiredTrafficType* attribute, only those Passport trunks which include voice or data in their *supportedTrafficTypes* attribute are considered for a route.

Passport trunk type

The *trunkType* attribute on a Passport trunk indicates the type of facilities that a Passport trunk uses. Examples of Passport trunking facilities are terrestrial or satellite links. A Passport trunk can be one of eight possible Passport trunk types.

The *PermanentLogicalConnection* component has an attribute called *permittedTrunkTypes*, which specifies a set of possible Passport trunk types that can be used on the route. There can be up to eight Passport trunk types specified.

A Passport trunk is considered a potential candidate for route selection, if an only if its *trunkType* is one specified in the *permittedTrunkTypes* attribute in the *PermanentLogicalConnection* component.

As an example, if the *permittedTrunkTypes* attribute is the set to terrestrial or satellite, then only those Passport trunks with terrestrial or satellite *trunkType* attribute are considered potential candidates for route selection. To prevent a route from using satellite links, the *trunkType* representing satellite links must be omitted from the *permittedTrunkTypes* set.

Customer defined parameter

PORS provides the network administration (our customer) with the flexibility to define a parameter for route selection to tailor the route to their special needs that are not explicitly addressed by PORS.

The customer can specify an attribute in the *PermanentLogicalConnection* component, called *requiredCustomerParameter*, which defines a level which must be satisfied by the Passport on the chosen route. It has possible values from zero (0) to seven (7). A corresponding Passport trunk attribute, the *trunkCustomerParameter*, specifies the level supported on the Passport trunk.

The Route Selector only considers Passport trunks with a *trunkCustomerParameter* attribute which is numerically less than or equal to the *requiredCustomerParameter* attribute of the path.

Figure 13, “Passport trunk candidates based on customer defined parameters,” (page 105) shows an example of candidate Passport trunks based on the customer defined parameter. The *requiredCustomerParameter* attribute is 4, and the *trunkCustomerParameter* attribute is indicated on each Passport trunk. The solid line shows the Passport trunks that are candidates for establishing a route.

Figure 13
Passport trunk candidates based on customer defined parameters

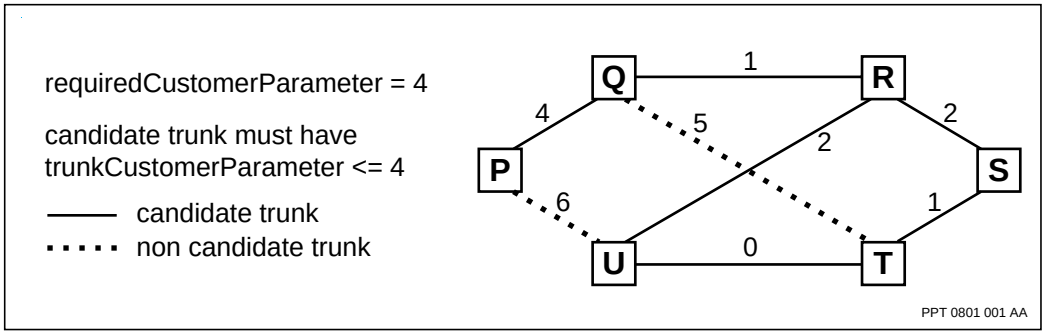

Minimization criteria

The Route Selector algorithm computes a minimum path with respect to only one metric: cost or delay for the Passport switch. It uses the second metric to select between Passport trunks of equal (first) metric within a link group. Therefore, the *trunkAttributeToMinimize* attribute of the *PermanentLogicalConnection* specifies one of two possible minimization criteria:

- minimization of the route cost. For Passport trunks of equal cost, the Passport trunk with the least delay is selected.
- minimization of the route delay. For Passport trunks of equal delay, the Passport trunk with least cost is selected.

Maximum acceptable cost

The cost is a metric assigned to a Passport trunk in the *trunkCost* attribute, which is representative of the amount of network resources consumed when the Passport trunk is used. For example, cost can indicate monetary instrument such as dollar value that an administration charges for the use of a Passport trunk, a geographic distance, a hop count, or the speed of the Passport trunk or a composite function of the above. The cost metric can be used by an administration to favor or discourage the use of some Passport trunks over others.

The *maximumAcceptableCost* attribute specifies the highest value of the total cost of all Passport trunks of a route. If a potential route has a cost greater than the value specified for the *maximumAcceptableCost* attribute, the route is not selected.

The following situation can happen when route delay is minimized. Route Selector for the Passport switch can find a non-optimal route within the *maximumAcceptableDelay* and *maximumAcceptableCost* attributes, which does not cause bumping. Otherwise bumping at the next priority level (starting at the lowest) is attempted until a route is found or all choices are exhausted.

Maximum acceptable delay

The *maximumAcceptableDelay* attribute of a *PermanentLogicalConnection* specifies the upper bound of the total delay experienced on any selected routes. Each Passport trunk has a delay metric associated with it. The metric is measured by the Passport trunks. Currently the metric is a delay value that approximates the delay for packets across a Passport trunk.

During selecting a route, the Route Selector for the Passport switch adds all the delay metrics of all the Passport trunks on the route. The route is acceptable only if this total is less than or equal to the value set for the *maximumAcceptableDelay* attribute.

Manual path

Although PORS selects the best route between two end points, it is desirable to have a network operator manually select the route for a PLC. Typical applications are restrictions on the route that cannot be expressed in the path characteristics (such as policy based routing), the need for a test PLC, or when routing between two or more topology regions. Manual paths apply to PLCs only.

Using the manual path capability depends on the type of service using PORS routing:

- For transparent data services (TDS), use a manual path to select a route manually for a single PLC.

- For managed cut-through switching (MCS), use a manual path to select a route manually for multiple connections. The *ManualPath* component and attributes contain the sequence of Passport trunk component names to be used on the route.

The bandwidth requirements and the configured PORS characteristics still apply to manually-selected paths.

Manual path with forced bandwidth

When a path is selected by an operator, if the *pathType* attribute of the *PermanentLogicalConnection* is set to forced, the path characteristics are not checked on the specified route. (The *pathType* route selection attribute applies to *PermanentLogicalConnection* components only.) The required bandwidth is reserved, even if it exceeds the reservation limit, maxReservedBwOut. The forced manual path can violate the path characteristics and the bandwidth reservation rules and therefore is used with caution for testing purposes only.

Note: If a forced bandwidth pushed the allocated bandwidth over 100%, the Path Administrator will not account for this bandwidth properly when the bandwidth falls below 100%.

Terminate when rerouting

Some services cannot tolerate the delays imposed when a path is rerouted. The *pathFailureAction* attribute of the *PermanentLogicalConnection* is used to cause a logical connection to terminate instead of having its path be rerouted. The logical connection is terminated whenever the path fails because of bumping, a component failure on the route, or load shedding due to congestion.

This requirement is not used during route selection.

Optimization

The *optimization* attribute is used to determine if the connection can attempt to follow through with the optimization process when requested by PORS Connection Control. If this attribute is set to disabled, the connection cannot be optimized even though it has been requested to do so by PORS Connection Control. If set to enabled, then the connection is optimized as requested by PORS Connection Control. For more information on optimizing paths, see “Path optimization” (page 115).

Bump preference

The *bumpPreference* attribute allows bumping to occur in the route selection process. A connection with a higher (numerically less) setup priority can bump a connection with a lower (numerically higher) holding priority in order to acquire its bandwidth and become established.

When set to bumpWhenNecessary, an optimizing connection or any new connection attempts to establish a path where bandwidth is available. At this point, it cannot bump current connections. Where bandwidth is available, the path is chosen along this new route. If bandwidth is not available and the connection is optimizing then the optimization is aborted. This is because optimization is a non-disruptive process by default. That is, optimization of one connection must not cause rerouting of another connection. For new connections attempting to establish a path where bandwidth is not available, the new connection establishes a path by bumping current connections which have a lower (numerically higher) holding priority than the setup priority for this connection. When set to bumpToObtainBestRoute, an optimizing connection or any new connection always bumps current connections in order to establish an absolutely optimal path when bandwidth along that path is reserved by connections with holding priority lower (numerically higher) than the setup priority for this connection. For more information on optimizing paths, see “Path optimization” (page 115).

PORS routing profile

For ATM services transmitted over PORS, the transport characteristics can be defined in a PORS routing profile, instead of the *PermanentLogicalConnection* component. The *Profile* component is associated with the calling, or source, end point of the ATM service.

The PORS profile contains most of the transport characteristic attributes that are defined in the *PermanentLogicalConnection* component for other services. These attributes include the type of routing to be done, cost minimization of delay minimization, and the types of Passport trunks that are acceptable for the call.

A profile can be shared by many services of different types. For example, you can design one profile to represent setup and holding priority 1 CES circuits, another for lower-priority CES circuits, and another for MCS calls.

Alternatively, you can use one profile for each call, or each type of call, or each service. The only limit is the maximum number of profiles for a module, which is 255.

Two areas of the PORS profile that differ from the PORS PLC are bandwidth reservation (*requiredTxBandwidth* and *requiredRxBandwidth* attributes), and emission and discard priorities, known as quality of service (QOS). The QOS and bandwidth reservation values are signalled to PORS by the application. PORS uses the signalled values unless you override them in a profile.

If you do not provide a profile for an ATM service configured over PORS, the default values for the transport characteristics are assigned automatically to each connection.

Instantiating the route

After the route selector selects a route, the source end point instantiates it. At each Passport trunk of the selected route, and for each direction of transmission, PORS

- verifies that the required bandwidth is available from the pool of unreserved bandwidth configured for path-oriented traffic
- reserves the required bandwidth
- assigns an LCN to the path on the trunk. The LCN is the same for both directions of transmission over the Passport trunk

The path instantiation procedure involves laying down the path using the least cost/delay route chosen by the Route Selector on a hop-by-hop basis. Bandwidth availability is verified and is reserved, and the LCNs are allocated on all Passport trunks on the route, for both directions of traffic.

Path instantiation is done in a progressive fashion, by sending a path setup packet along the selected route, as illustrated in Figure 14, “Path setup packet,” (page 110). The packet travels in the forward direction, from the source end point to the destination end point, stopping at all Passport trunks along the way. The backward direction is from the destination to source end point. The outbound Passport trunk is the one which emits the path setup packet, while the inbound Passport trunk receives the packet.

Figure 14
Path setup packet

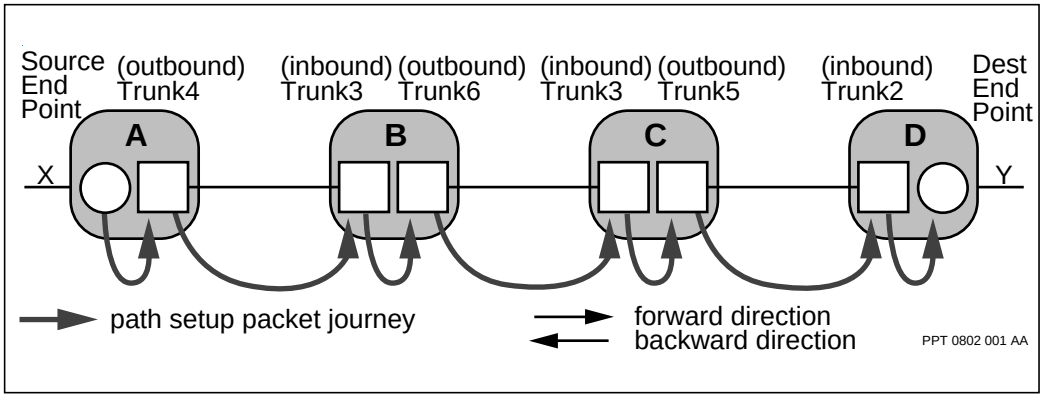

Along the way, the path setup packet triggers the following actions:

- verification of bandwidth availability

  The Passport trunk verifies that the value set for the required bandwidth attribute does not push the current reserved bandwidth beyond the reservation limit (*maxReservedBwOut* attribute). The outbound Passport trunk verifies the bandwidth for the forward direction of the path, while the inbound Passport trunk verifies it for the backward direction.

  Forced manual paths do not need their bandwidth requirements verified.

- reservation of bandwidth

  The Passport trunk adds the path bandwidth to the reserved bandwidth pool.

  Forced manual paths reserve the required bandwidth even if this causes the reserved bandwidth to exceed the reservation limit (*maxReservedBwOut* attribute).

- creation of the *LogicalChannel* components, one per Passport trunk

  On the inbound Passport trunks, the *LogicalChannel* component for the transfer of packets in the forward direction of the path is created. The outbound Passport trunks create the *LogicalChannel* component for the backward direction of the path.

- allocation of the LCN to be used on each Passport trunk

  The outbound Passport trunk chooses the LCN which must be confirmed by the inbound Passport trunk.

When the path setup packet reaches the destination end point, a path setup confirmation packet is returned to the source end point in the backward direction, stopping at all Passport trunks along the way. When this packet arrives at the calling end, path for data transfer is enabled.

The selected path can fail to instantiate under the following conditions:

- There is not enough bandwidth available.
- There is a failure (node, FP or Passport trunk) along the chosen route.
- The Passport trunk has reached the maximum number of paths that it can support (*maxLc* attribute).

In case of failure, a path tear down packet is returned from the point of failure back to the source end point. All the LCNs and the bandwidth reserved prior to the failure point are de-allocated as the packet travels down the backward route. The end point reports to the route selector the failure reason of actual bandwidth available or the component failure, and requests a new route. The route selector updates topological information before selecting another route for the same path, thus learning from the failure, and selects another route. The route selector therefore learns from its mistake and does not repeatedly select unavailable routes. If another route is selected by the route selector, the end point starts the instantiation procedure again. If another route is not available, route selector informs the end point that there is no route.

For manually selected paths, route selector learns from the failure but does not attempt to select another path. Forced manual paths only fail to instantiate if there is a failure on the selected route.

Rerouting paths

A path is expected to follow the same route for the lifetime of the logical connection. However, under certain conditions, a new route is needed for an existing connection.

Manual and forced paths cannot be rerouted or optimized. Also, paths for PLCs that have been configured to disconnect when a momentary interruption occurs cannot reroute to other paths. These paths can, however, be optimized. This capability is possible because rerouting causes the service to go down for the duration of the rerouting while optimization is transparent to the service. Paths that cannot be rerouted are terminated when a rerouting condition is encountered.

Information on rerouting is organized under the following headings:

- “Bumping paths” (page 112)
- “Network rerouting around component failure” (page 113)
- “Rerouting by operator command” (page 114)
- “Path optimization” (page 115)

Bumping paths

Paths being established are assigned a setup priority, while established paths are assigned a holding priority. There are five setup and five holding priorities. A path that is being established can bump an existing path with a holding priority that is lower than the setup priority for the new path. This allows higher priority calls to steal the bandwidth from lower priority calls. The bumped path is rerouted to an alternate route that can satisfy its requirements. If the bumped path cannot be rerouted, it is terminated. The setup and holding priorities can be different to allow setup at one priority and holding at an independent priority. This allows some calls not to invoke bumping and not to be bumped at the same time.

Bumping by default only happens as a last resort when there are no routes available for a given path. That is, *bumpPreference* attribute of the *PermanentLogicalConnection* component is set to bumpWhenNecessary by default. When this attribute is set to bumpToObtainBestRoute then bumping occurs to establish an absolutely optimal path. For more information on this attribute, see “Bump preference” (page 108).

Bumping occurs when there is no single piece of bandwidth to suit the current requirements, or when the cost and delay are such that only certain occupied routes satisfy the current path request.

When bumping is necessary, the RS tries to minimize disruption caused by bumping and for this result it recomputes a route using a topology whose Passport trunks bandwidth does not include as reserved the lowest holding priority paths (numerically highest). If a route is still not available, it considers bandwidth from the next higher holding priority paths. This is repeated until the bandwidth for all paths of lesser holding priority than the setup priority for the new path has been considered.

During the instantiation of a path which must bump other paths, lower holding priority paths are bumped before higher priority paths. The decision as to which of the available paths are bumped at each intermediate node by the new path is arbitrary.

Network rerouting around component failure

When the failure of a component on a route is detected, the network attempts to reroute all affected paths to bypass the failed component.

Detection of a failure is propagated towards the end points from either side of the failure, de-allocating resources along the way. The network terminates the old path and the source end point proceeds with the path establishment procedure. Failure to find a new route results in path termination.

PORS is concerned with failures that cause unavailability of a Passport trunk and thus prematurely terminate all the paths carried on the Passport trunk. Passport trunk unavailability can be caused by

- data link failure
- failure of a subsystem which interacts with the Passport trunk within the function processor (such as a link controller)
- failure of the function processor itself
- failure of the entire node

More than one Passport trunk becomes unavailable if the function processor that failed was servicing more than one data link. A failure causing outage of the node makes unavailable all Passport trunks connected to it.

A failure usually initiates a lot of activity outside of PORS. A down event triggers alarm generation, system diagnosis, and network topology updates. The updates are performed by the topology manager components residing in the control processors of each node. These activities are simultaneous with the PORS reactions that minimize loss of data packets and the length of the service disruption. The loss of data packets is minimized by promptly informing the applications at both ends of the logical connections about unavailability of the path and the necessity to stop sending data packets. The length of the service disruption is minimized by promptly initiating the rerouting process and maximizing its effectiveness. Rerouting is initiated by the source end point as soon as it receives the information about path failure.

The effectiveness of rerouting is maximized by quick termination of the paths that were carried on the unavailable Passport trunk. Path termination causes the return of the path bandwidth to the free-bandwidth pool for the Passport trunk. The network topology data base at the source of the path is updated on the released bandwidth.

The termination process is distributed along every path affected by the failed component. Tear down packets are sent on each path from the failure point and propagated in each direction towards the source and destination end points. Software controlling all Passport trunks on the failed paths receive updates on the failure, and then terminate the failed logical channels to free up the reserved bandwidth.

Information about the failure is simultaneously broadcast to all nodes in the network by the topology maintenance system. This broadcast permits prompt updating of the topology database in each node. There are two ways that the end point can be informed of the failure: broadcast notification and reception of a tear down packet.

Rerouting by operator command

An operator can initiate rerouting of an established path that supports the following services:

- Voice Service

- bit transparent data service (BTDS)
- HDLC transparent data service (HTDS)
- AAL1 circuit emulation service (AAL1 CES)
- multiservice cut-through switching (MCS)

Operator-initiated rerouting is achieved through the CAS reroute command. See “Forcing a path to reroute” (page 71).

This command attempts to reroute the service path to a new path as long as the new path has the same or better characteristics for delay or cost, depending on the characteristics that the service is minimizing.

If the reroute command is unsuccessful. the network operator can ask for the path of an existing logical connection to be rerouted by clearing a logical channel of the path. In this scenario, the switch at the far end undertakes the following checks:

1 if the *pathFailureAction* attribute is set to the *disconnectConnection* attribute, does not initiate the reroute

2 terminates the path by transmitting a tear down packet

3 requests a new route from the route selector in a Passport switch; if this request fails, terminates the logical connection

4 if a new route is selected, instantiates the path on the new route; if setup fails, repeats the check in point 3

It is possible that the new path that results from clearing a logical channel is not better (in terms of the cost and/or delay metrics) than the previous path. The operator is taking a chance that a worse route can be chosen given the current network load, or that the path fails to reroute and is terminated.

Path optimization

The objectives of automatic or manual optimization are twofold:

- move a PORS connection to a lower cost or delay route if one exists
- if the path is already on the best route, then balance the bandwidth usage of the links within the link groups along the route supporting the path being optimized

Balancing the bandwidth between the various links is accomplished by moving one or more connections to paths on different links within the link group until load balancing is achieved. The load is balanced when individual links in the same link group each carry approximately the same proportion of PORS traffic to the total bandwidth for the link.

More than one optimization cycle might be required to achieve both objectives described above. For example, one path moving to a better route might introduce some load imbalance that can't be rectified in the same optimization cycle because all the other paths along that route have already been optimized during this cycle. The next optimization cycle re-balances the load if it is possible.

> ***Note:*** A link in this context is defined as a Passport trunk, not a physical link.

The optimization process is administered by PORS connection control which resides on each Passport node in the network.

Benefits of optimization

Optimization offers the following benefits:

- Optimization maximizes total network bandwidth by rerouting the highest bandwidth connections to routes with the least number of hops.
- As new Passport trunks are configured or come back to service, these trunks are used at the next optimization interval.
- A network using optimization has lower bandwidth consumption across the network, freeing up bandwidth for other traffic.
- A manual reroute command is available which allows the user to reroute specific connections.
- Optimization balances link groups along connection paths in terms of bandwidth utilization.

Triggering the optimization process

Optimization of a connection on a node can occur in one of three ways:

- automatic optimization of a node through configuration
- manual optimization of a node by operator command

- manual optimization of a connection by operator command

Automatic optimization on a node is enabled when the operator enters configuration mode and types the command

```
add Routing Pors
```

to activate the optimization feature on that node by creating PORS Connection Control. Thirty minutes after activating the feature, PORS Connection Control begins automatic optimization. Automatic optimization works on paths from the calling end only, provided these paths have been operational for at least one minute, and have their *optimization* attribute for the *PermanentLogicalConnection* set to enabled.

Manual optimization of a node by operator command occurs when the operator types the command

```
optimize Routing Pors
```

This command causes all connections on that local node to be optimized. Manual optimization applies to paths from the calling and called end. Paths do not have to be operational for one minute for them to be optimized, but they must have their *optimization* attribute set to enabled. If optimizing a path from the called end, the *remoteName* attribute must be specified. Also the attributes for the *PermanentLogicalConnection* on the called side must be identical to that on the calling side.

Note: Only one path is optimized per second during automatic and manual optimization of a node.

Manual optimization of a connection by operator command occurs when an operator types the command

```
reroute <servicetype>/ n lco
```

where <servicetype> can be Voice Service, BitTransparentDataService, HdlcTransparentDataService, or Aal1Ces. This command attempts to optimize the service path to a new path as long as the new path is not worse than the current path in terms of delay or cost, depending on whichever attribute the service is attempting to minimize. If the reroute command is

issued on the called end, *remoteName* attribute must be specified. Also the attributes for the *PermanentLogicalConnection* on the called side must be identical to that on the calling side.

> ***Note:*** Reroute command disregards the *optimization* attribute. That is, optimization occurs whether optimization is enabled or disabled. It is possible that the reroute command results in a path selection that is identical to the original path.

This command triggers the rerouting of an established path which runs on a specific service. It attempts to optimize the current connection and result in the change of the current path to a better one if one exists.

Transporting data on the path

The packets transmitted on the path are delivered in order of transmission since all packets follow the same route in the network. Procedures to insure orderly delivery of packets are used during rerouting and optimization. This removes the necessity for end point buffering to re-establish order, or the need for sequencing packets at the routing level, thus requiring less bandwidth for each packet.

The current PORS design is to be used by services which tolerate loss of packets. PORS does not support end-to-end data reliability, windowing, and flow control, but it is possible to add these functions on top of PORS. When adding these functions, take care that configurations do not affect existing PORS performance. Because packet loss is permitted, bandwidth usage is more efficient since packets are never retransmitted. The lack of retransmitted packets also insure that packets cannot be duplicated. Packet loss can be due to

- poor quality facilities losing packets
- packets discarded to alleviate congestion in the network
- packets discarded or lost when a path is being rerouted or optimized

The path-oriented packets make efficient use of the bandwidth by requiring only minimal information in the Passport routing header:

- An LCN is used instead of a lengthy destination end point address.
- There is no need for packet sequence numbers at the routing level.

- Destination of the next hop, emission priority and discard priority are kept at each hop of the route instead of in the packet header.

PORS makes use of packet forwarding procedures that distinguish between the urgency at which packets are emitted, and the importance of emitting those packets. These procedures also provide a mechanism to accomplish bandwidth sharing between connectionless and path-oriented traffic.

The urgency of a packet is reflected by the emission priority. The emission priority associated with a path is an indication of the queue (software or hardware) that a packet waiting to be emitted on a Passport trunk is inserted in, which in turn is a reflection of how often a queue gets serviced by the Passport trunk.

The discard priority of a packet indicates the importance of a packet. This is used in the decision to discard packets in congested situations or for bandwidth sharing purposes. Each packet can also be individually labelled by the end points as being eligible for discard in preference to others in congested situations.

Terminating a route

The path termination is the process of releasing the reserved bandwidth and de-allocating the LCNs used on the Passport trunks on the route of a path. Path termination is accompanied by a topology update with information on the bandwidth released on a particular Passport trunk.

A path is terminated in the following cases:

- The service requests the path termination.
- The path is being rerouted due to bumping by a path with higher setup priority or failure and the rerouting path cannot be established.
- The path is being rerouted due to bumping by a path with higher setup priority or failure and the *pathFailureAction* attribute under the path *PermanentLogicalConnection* is set to disconnectConnection.

The logical connection is also terminated when the path is terminated, except when the termination is due to a reroute attempt. For reroute attempts, the logical connection is terminated only if the reroute fails or is not allowed.

The path termination is done in a progressive fashion. A path termination packet is sent from the end point initiating the termination, all the way to the remote end point. At each Passport trunk along the path, the reserved bandwidth for both directions of transmission is released. The LCN is marked as terminating but it cannot be released until the path termination procedure is completed. This prevents data packets from being transferred on the path.

At the remote end point, a path termination confirmation packet is returned to the initiating end point. It de-allocates the LCN used for the path.

Termination of logical connection - reroute implications

A logical connection whose path is being rerouted terminates under the following conditions:

- A failure has terminated one of the end points of the logical connection.
- For the Passport switch, the Route Selector cannot find a route satisfying the path requirements, even when bumping has been permitted.

  Note: In a Passport switch, a path instantiation failure never causes a termination of the rerouting procedure since it always results in a request to the Route Selector for another route.

- The logical connection has the *pathFailureAction* attribute set to disconnectConnection.

  If this attribute is set to disconnectConnection, no attempt is made to select another route and the logical connection is terminated. This is desirable for applications which cannot tolerate the delay and service disruptions encountered because rerouting process.

Chapter 8
Routing applications

This section describes specific applications for the Passport Path-Oriented Routing System (PORS) outside of the default implementation for services. Information in this section is organized under the following headings:

- "Selecting paths based on cost and delay" (page 121)
- "Restricting traffic" (page 123)
- "Restricting paths" (page 124)
- "Configuring ATM services over PORS on a Passport switch" (page 129)
- "PORS across topology regions" (page 133)

For information on PORS fundamentals, see "Routing fundamentals" (page 89).

Selecting paths based on cost and delay

PORS can select a path based on either the lowest cost or lowest delay. Both cost and delay cannot be minimized. Use the *pathAttributeToMinimize* attribute under the *PermanentLogicalConnection* component to specify cost or delay. The routing system computes a path based on optimization criteria (cost and delay) from

- the values assigned to the Passport trunks (cost)
- measured delay values that are associated with each Passport trunk

This calculation is shown in the figure "Path for cost or delay using trunkAttributeToMinimize" (page 122).

To assign a cost to a Passport trunk use the *trunkCost* attribute under the *Trunk* component. Cost can be an actual dollar value or any parameter that you want to use. If default values are used, cost represents the hop count. Thus the number of hops across the network is minimized. If you use a parameter for cost that reflects the actual cost of facilities, high cost facilities receive less use and reduce the cost of operating the network. This is the recommended method of using this option.

Figure 15
Path for cost or delay using trunkAttributeToMinimize

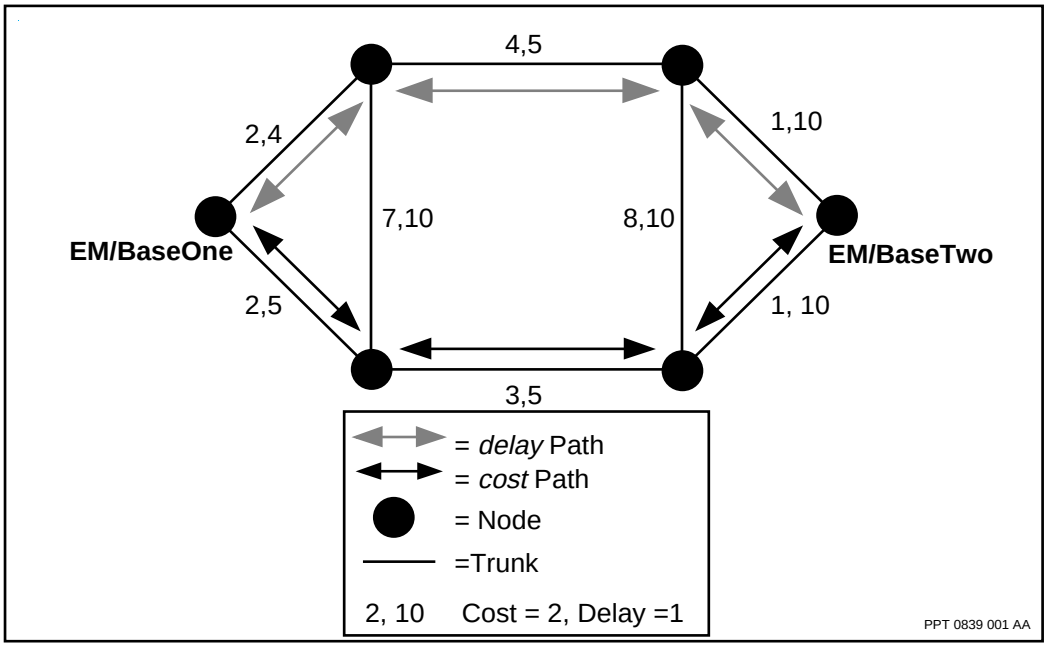

Specifying a maximum cost for a path

Providers of network services can restrict some parameters for a particular circuit (for example the number of hops).

To specify the maximum total cost value of a path you can use the *maximumAcceptableCost* attribute under the *PermanentLogicalConnection* component. Although this value is called cost, you can use it to reflect a variety of considerations, including geographic distance, hop count, or real

dollar value. The sum of the *trunkCost* attribute values of all Passport trunks used in the path is less than or equal to the value specified by the *maximumAcceptableCost* attribute.

Specifying a maximum delay for a path

Delay in PORS is measured for a 512 byte packet in one direction at the time of Passport trunk staging. Over time this measured delay can change to reflect the updated operating delay but does not affect existing paths unless a Passport trunk restages. If Passport trunk staging is the measurement instance, then new or existing paths do not pick up the new delays. Choose a maximum delay with care since it is dependent on the real delay of a cell on the path.

To specify the maximum delay value of a path, use the *maximumAcceptableDelay* attribute under the *PermanentLogicalConnection* component. The sum of the delay values associated with all Passport trunks used in the path are less than or equal to the value specified by the *maximumAcceptableDelay* attribute.

> ***Note:*** This parameter is used when large delays are unacceptable for the service (for example voice and other interactive data).

Restricting traffic

PORS allows you to specify which types of traffic are carried on a given Passport trunk. Use the *supportedTrafficTypes* attribute, under the *Trunk* component to create an individual list of traffic types for each Passport trunk in your network (data, voice, and video for example).

When you configure the connection, use the *requiredTrafficType* attribute to specify which traffic type is transported by the path. PORS chooses Passport trunks that include the *requiredTrafficType* attribute list in their *supportedTrafficTypes* attribute list. The the *requiredTrafficType* attribute list must be a subset of the *supportedTrafficTypes* attribute list or PORS does not select the Passport trunk for the path.

For example if the requiredTrafficType list of a path is data, only Passport trunks with supportedTrafficTypes lists that include data are selected for the path.

Restricting traffic to certain types of Passport trunks

You can create an indicator of the type of Passport trunk that various traffic types use. Examples of Passport trunking facilities might be terrestrial or satellite links. The *trunkType* attribute, under the *Trunk* component allows you to do this for up to eight different types of Passport trunks.

The *permittedTrunkTypes* attribute under the *PermanentLogicalConnection* component allows a set of possible Passport trunk types to be specified for a route. Only Passport trunks with *trunkType* attributes that are found in the *permittedTrunkTypes* attribute list are used to create the path.

Restricting paths

You can restrict certain classes of paths to certain Passport trunks. Most of the commonly used qualifiers are represented in security and traffic type. This is an additional option to be used for any function that you think appropriate.

Security

PORS allows you to define varying security levels for the Passport trunks of the network. This option can, for example, prevent sensitive data from traveling over certain Passport trunks.

PORS has an option that allows you to specify the minimum security level of a path. To do this, configure a security value for the Passport trunks in your Passport network using the *trunkSecurity* attribute under the *Trunk* component. When you configure the connection, enter a value for the *requiredSecurity* attribute under the *PermanentLogicalConnection* component.

The connection only uses Passport trunks that have been assigned security values of an equal or higher level than that of the connection. A lower number represents a higher security level. This is illustrated in the example in the figure “Path determined using a requiredSecurity value of 4” (page 125).

The default value for security is mid-range so that the network administrator can add security with minimal configuration.

> ***Note:*** Over-use of this option can reduce its usefulness. This option can also reduce the number of recovery paths available to high security routes when an outage occurs.

Figure 16
Path determined using a requiredSecurity value of 4

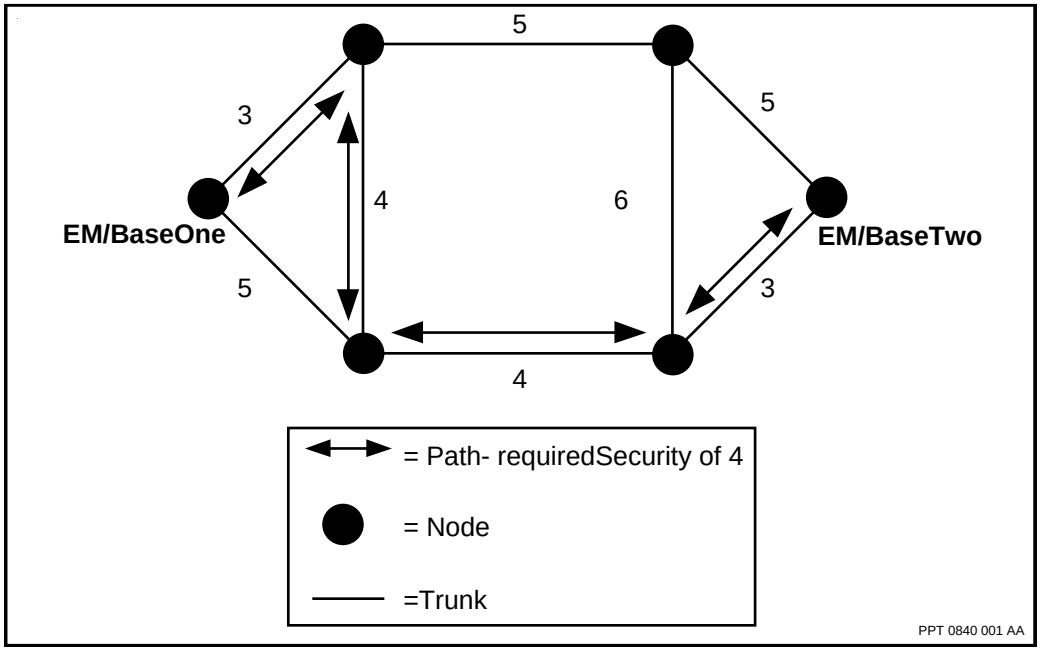

Defining general parameters to restrict paths

PORS allows you to restrict certain paths to certain Passport trunks. This is done in a similar manner to the way that security is configured. Values are assigned to various Passport trunks in the network using the *customerParameter* attribute under the *Trunk* component. When the path is configured, it can be assigned a value using the *requiredCustomerParameter* attribute under the *PermanentLogicalConnection* component.

PORS assigns the path to Passport trunks that have an equal or lower number associated with them. This is illustrated in the example shown in the figure “Path using a requiredCustomerParameter of 4” (page 126).

Figure 17
Path using a requiredCustomerParameter of 4

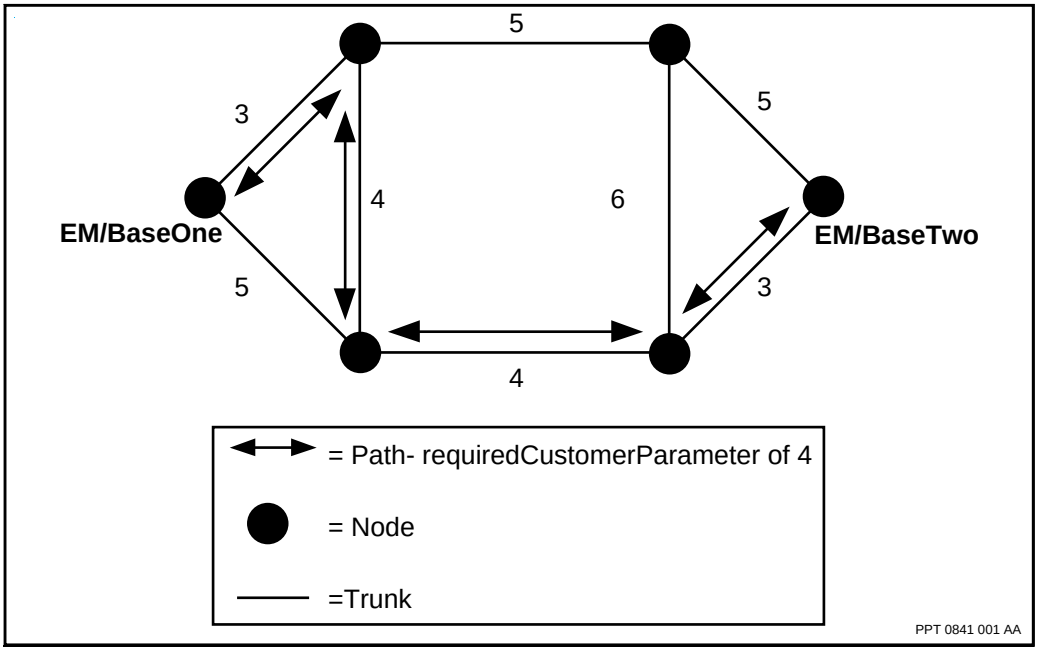

Specifying a path manually

PORS is designed to select an appropriate route automatically. In an exceptional case, however you can define the set of Passport trunks that are used.

The path can be defined at both ends and the two paths do not have to use the same set of Passport trunks. If different paths are defined at each end, however PORS does not guarantee which one is used.

Defining different routes manually at both ends of a connection enhances robustness of the connection against network element failures. If the existing path is subject to a failure, another route is used to reroute the path very quickly without further intervention from network administrator, provided the other route is usable.

If you want to override the automatic selection of a path and specify the Passport trunks manually use the *manualPath* attribute of the *PermanentLogicalConnection* component. Enter the outbound sequence of *Trunk* component names for the path that you want.

> ***Note:*** The path still must satisfy the characteristics specified in the other attributes under the *PermanentLogicalConnection* component, including bandwidth requirements.

To configure a manual path, see "Configuring optional manual paths for managed cut through switching" (page 31) and "Configuring optional manual paths for a PLC" (page 33).

Tandem suppression

The tandem suppression feature gives the network operator increased control over routing behavior by preventing tandem traffic from travelling through specified nodes. When tandem suppression is enabled on a node, the only PORS traffic that the network routes to the node is traffic destined for the node itself. The tandem suppression option applies to PORS automatic calls. PORS manual paths can still traverse tandem-suppressed nodes when the option is in effect.

Tandem suppression is particularly useful for small access Passport nodes and customer premise equipment (CPE) edge nodes. The feature can be used to control traffic flows by

- preventing traffic from routing through nodes with insufficient link capacity
- preventing traffic from routing through CPE
- preventing traffic from flowing through a newly-deployed node until the network operator is sure that the node is up and running successfully

When tandem suppression is enabled on a node, it affects PORS and DPRS traffic.

In the example in the figure "Tandem suppression" (page 128), PP5 is assigned the tandem suppression option, so traffic is not routed to that node from any of the other nodes unless it is destined for PP5.

PORS tandem traffic is not suppressed immediately after the node is configured for tandem suppression. Instead, it is suppressed either when the tandeming *PermanentLogicalConnection* component is re-established, or when PORS optimization reroutes the *PermanentLogicalConnection* component away from the tandem-suppressed node. In most cases, this is not an issue, as most PORS traffic is local traffic to a node that is typically an edge node.

Figure 18
Tandem suppression

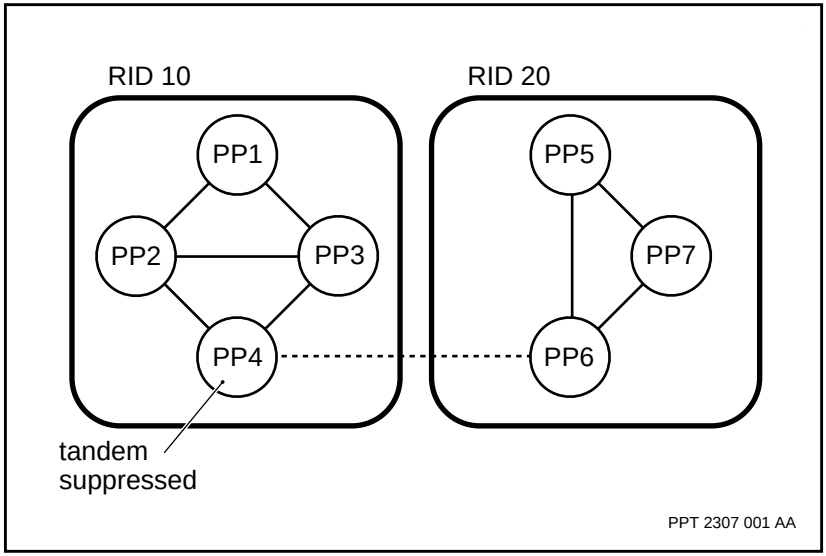

When adding a new node at the periphery of the network, it is important to be aware of the tandem-suppression configuration of the connected nodes. If the new node is connected to tandem-suppressed nodes, it is effectively severed from the network. For example, in the figure "Isolated node" (page 129), node PP4 becomes isolated if nodes PP2 and PP3 are tandem suppressed.

Note: This example shows a non-robust network

Figure 19
Isolated node

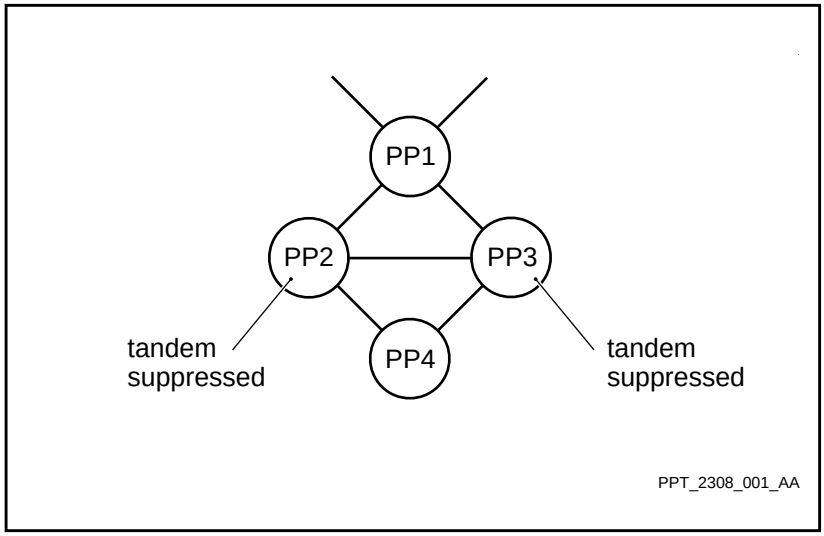

It is important to set up a robust network that can remain operational even when links go down. For example, assume that only PP3 in the figure "Isolated node" (page 129) is tandem suppressed. If the link between PP2 and PP4 goes down, PP4 is still effectively isolated.

Configuring ATM services over PORS on a Passport switch

Routing an ATM service over PORS on a Passport switch has specific configuration requirements. This section describes these requirements under the following headings:

- "PORS profiles for ATM services" (page 129)
- "Path administrator mode for the trunk" (page 132)
- "NSAP addressing" (page 132)

PORS profiles for ATM services

To create PORS routing profiles for ATM services on a Passport switch, do the following:

- create a profile for each service or connection

- customize the profiles to suit each individual service or connection
- associate the profile with the service source end point component (for information on how to create the link, see the appropriate service guide)
- activate the connection (changes to a profile do not take effect until the service launches the next call request)

When a PORS profile is created, attributes are assigned the following default values:

```
setupPriority = 2
holdingPriority = 2
bumpPreference = bumpWhenNecessary
requiredTxBandwidth = 0 bit/s
requiredRxBandwidth = 0 bit/s
requiredTrafficType = data
permittedTrunkTypes = terrestrial satellite tt1 tt2
tt3 ~tt4 ~tt5 ~tt6
requiredSecurity = 4
requiredCustomerParm = 4
pathAttributeToMinimize = cost
maximumAcceptableCost = 1280
maximumAcceptableDelay = 100000 msec
emissionPriority = sameAsUser
discardPriority = sameAsUser
pathFailureAction = reRoutePath
optimization = enabled
```

Customize the profile by changing these default values to values more appropriate for the service. For example, the setup and holding priorities in the profile control the PORS bumping feature. In the default profile, these priorities are set to 2, which specifies average priority. To define an important service that bumps lower-priority services, set the setup and holding priorities to 1.

Another area of the profile that is typically customized to a service is the quality of service (QOS). By default, the emission and discard priorities assigned to a connection are the ATM values signalled by the application. For example, each ATM service connection is defined with a particular service category (such as CBR or VBR), which corresponds to specific emission and discard priorities in Passport switches. By default, the profile does not override the selections made by the application. As a result, PORS assigns

emission and discard priorities along the path to the same values as the ATM routing system. However, PORS provides three different emission priorities and three discard priorities, resulting in nine different types of QOS. Since many of these combinations have no counterpart in ATM, PORS profiles provide a way to override the ATM service category. The table Table 5, “ATM-to-PORS QoS mapping,” (page 131) shows the mapping of ATM service category values to PORS QOS combinations, indicating that four of the nine permutations are available from ATM.

Table 5
ATM-to-PORS QoS mapping

ATM service category	***emissionPriority* attribute**	***discardPriority* attribute**
CBR	0	1
VBR-rt	1	1
VBR-nrt	2	2
UBR	2	3

Overriding the original service category allows two applications to run with different QOS values depending on their relative importance. For example, you can have one connection with the default service category, and another with a high discard priority, making its traffic very discardable. This situation can be accomplished by assigning a profile to the second connection, and setting its *discardPriority* attribute to 3. The next time the connection comes up, it uses the default ATM-assigned emission priority for its service category and the new discard priority of 3. In this way, the profiles allow ATM applications access to the full range of Passport QOS permutations on an instance-by-instance basis.

The profile also allows you to override the ATM bandwidth values. When there is no profile, or the profile bandwidth values are set to zero, PORS chooses an effective bandwidth based on the ATM peak cell rate (PCR). This is a conservative setting, and can result in under-used links.

To change the reservation used by PORS, you can set the requiredTxBandwidth or *requiredRxBandwidth* attributes of the profile. These values are in bits per second, not cells per second, so you must do the appropriate conversions.

For example, a particular service uses 683 cells/second. If you do not use a profile, or if you do not override the bandwidth values, PORS reserves 683x53 octets in both directions on every link the call traverses. If you know that the service is idle 50% of the time, you can reduce the reservations by setting the bandwidth attributes to (683x53)/2. If you combine this under-reservation of bandwidth with a high discard priority (to make the traffic more discardable), substantial bandwidth savings can result.

Path administrator mode for the trunk

At this stage, you must decide which Passport trunk PA ATM mode to use for each Passport trunk in the network. To support AAL1 traffic on the PORS network, the Passport trunks over ATM must be run in mapped mode. A mapped mode Passport trunk can support ATM services routed by PORS, and non-ATM services routed by PORS, both at full hardware speeds. The rerouting speeds are faster because there is more CPU available for the processing of setup messages. The system runs more smoothly because of the physical separation of the data path resources from the control path resources.

> ***Note:*** Mapping can only be used in a directly connected Passport configuration or in a network configuration where a virtual path is dedicated to the PA. Logical trunks cannot be configured in mapped mode.

NSAP addressing

ATM services like CES are routed over ATM links on NSAP addresses, rather than component names. When routing these services using PORS, you must provide an NSAP prefix for each node in the network. In addition, a route must be available between the source and destination that consists of entirely ATM hardware on which *Trunk* and *PathAdmin* components are running in mapped mode. In the sample network, all the Passport trunks are running in PORS mapped mode, ready to handle ATM calls.

The NSAP prefixes must be configured so that all NSAP addresses reachable on that node share the node prefix. However, PORS supports five distinct prefixes on each module, so you can use one prefix for CES services, and others for another ATM service.

Simply assign an NSAP prefix to each node, and then make sure all services reachable on that node share that prefix. An address is reachable by a node if the address of the node is a prefix of the address being routed to. For example, the address 123456789 is reachable by a node with an address prefix of 123456. PORS supports a 20-byte NSAP address, using the first 13 bytes as a node prefix and the last 7 bytes to identify the service.

In the sample network, you must assign four unique 13-byte node prefixes and then configure each node accordingly. Assign the values 00...01 through 00...04 to AM/A through EM/D. The following procedure is an example of this process.

By configuring each node with its own prefix, all the other switches learn about this prefix. You need to configure the prefix only once, and only on the node to which the prefix applies. If you change the prefix, PORS distributes it to the other nodes, updating them accordingly.

NSAP address information stored in the Base Routing Topology subsystem can be displayed through a component administration system (CAS) interface.

> ***Note:*** If an NSAP address can be displayed on a Passport node, it is reachable from the local node.

For more information, refer to “Displaying NSAP address information” (page 72).

PORS across topology regions

A topology region is a set of Passport nodes that are partitioned as an autonomous section of the network. All PORS services can route dynamically within a topology region because the topological information of the region is available to all nodes within the region. PORS across topology regions allows PORS to dynamically route calls across multiple topology regions while maintaining unique capabilities such as bumping, rerouting, and dynamic

path attribute (QOS) adjustment. Trunks that span region boundaries are referred to as inter-region Passport trunks. A call may be comprised of one or more path segments across a region or inter-region Passport trunk.

Path optimization and balancing are only available within a region path segment and not across inter-region Passport trunks. Rerouting and bumping may operate independent of other path segments if an alternate route to the same border node is available within the selected path segment.

Given that a number of inter-region Passport trunks may originate from a single node, bandwidth load spreading is an option available to avoid the over-utilization of a single link. Bandwidth load spreading is applicable for call establishment across inter-region Passport trunks that have been provisioned with identical characteristics by the network operator. Bandwidth load spreading is not applicable in the event of a link failure.

PORS services using NSAP addressing (CES, MCS, and VTDS) are able to route across multiple topology regions because the nodes at the edge of a topology region (the border nodes) are configured as inter-region Passport trunks. This configuration provides a connection to a border node in another topology region. Inter-region Passport trunks are configured with all of the NSAP addresses that they can reach in other topology regions. NSAP addresses that can be reached through an inter-region Passport trunks are called link-reachable addresses. Link-reachable addresses are distributed throughout the topology region of the border node that contains the inter-region Passport trunk, ensuring that all link-reachable addresses are known by all nodes within that topology region.

> ***Note:*** A border node may have more than one inter-region Passport trunk through which a link-reachable address exists. In this case, that address should be specified on all inter-region Passport trunks.

Information in this sections is organized under the following headings:

- “Passport clusters” (page 140)
- “Guidelines for Passport clusters” (page 142)

Overview

Nodes in a topology region know the topology of their own region only, however with PORS across topology regions, inter-region Passport trunks are configured with the NSAP addresses that they can reach in other topology regions, and these addresses are broadcast throughout the topology region of the inter-region Passport trunk.

In order to route across topology regions, a source node routes to a border node which advertises that it can reach the address of the destination node, the border node routes across an inter-region Passport trunk to a border node in the next topology region towards the destination node, and this border node routes to the destination node, or a border node to another topology region, and so on until the destination is reached. The complete path from source to destination is made up of many path segments, and each segment of the path can be balanced, optimized, bumped or rerouted as it is within a topology region.

PORS across topology regions allows for the division of a Passport network into smaller subnetworks which do not exchange Passport topology information with each other. This allows for the building of large Passport networks which are capable of dynamically routing PORS services end to end.

Passport clusters offer further possibilities for network growth, by connecting groups of Passport access switches around a topology region backbone. Passport clusters do not exchange topology information with the backbone nodes, and only a limited amount of routing information is exchanged.

Segmented optimization

The path for a segmented PORS connection consists of $2n+1$ call segments, where n is the number of inter-region links and/or cluster links traversed and the connection does not originate and/or terminate on a gateway node. Segmented optimization ensures that each PORS connection segment can be switched if a better path segment is available. For more information on optimization, refer to “Path optimization” (page 115).

Refer to figure “Segmented optimization across topology regions” (page 137) to see how segmented optimization works. In this example, node B (in Region 1) is the source node, and node F (in region 2) is the destination node. The call comprises three segments: from the source node B to the outbound gateway node C, from the outbound gateway node C to the inbound gateway node D, and from the inbound gateway node D to the destination node F.

Configure segmented optimization first at node B to optimize the path chosen from the source node to an outbound gateway node advertising a lower gateway cost or better address, or to select a better path to a gateway advertising the same gateway cost. Next, configure the outbound gateway nodes A and C to optimize the connection onto the link advertising the lowest gateway cost or address to the inbound gateway nodes, which may involve switching inbound gateway nodes, in this case nodes D or E. Finally, configure the inbound gateway nodes D and E to optimize the path segment to the destination node (node F); however, this is better left to legacy PORS optimization.

To invoke segmented optimization on nodes A, C, or D, first configure the nodes to enable segmented optimization (refer to “Configuring inter-region Passport trunks” (page 55) for details), then issue the following command:

```
optimize -segmented Routing Pors
```

Figure 20
Segmented optimization across topology regions

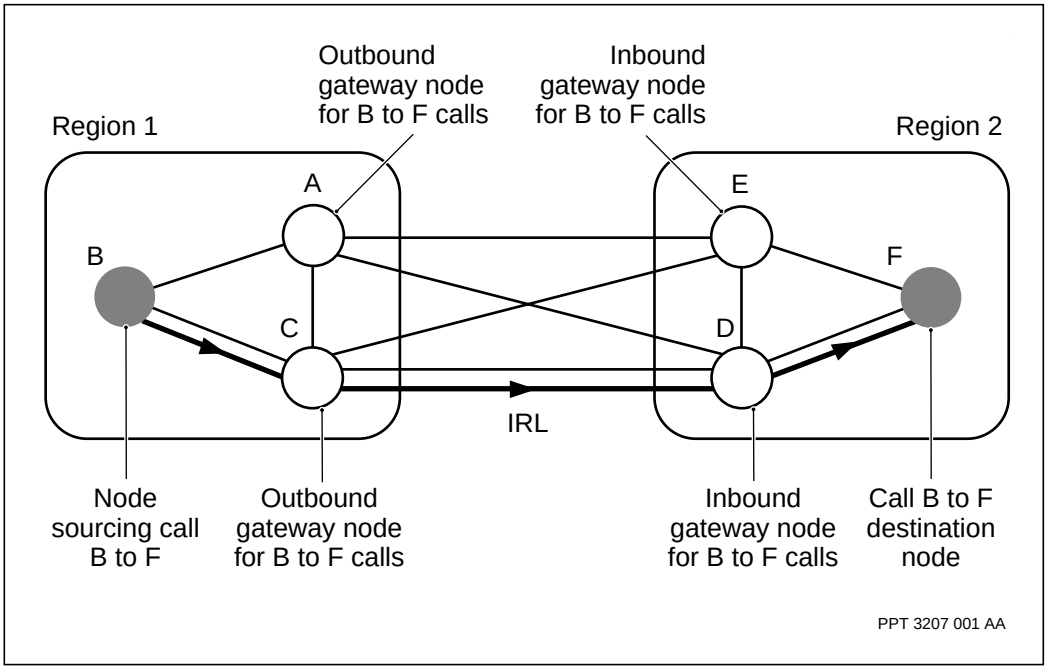

In order to minimize service disruption, segmented optimization should be run from source towards destination, in this order:

1 the source node

2 the outbound gateway nodes

3 the inbound gateway nodes

Note: Segmented optimization, since it is service disrupting, should only be run once a week, after a network failure, or after commissioning new gateway nodes.

For information on configuring each of these three segments for optimization, refer to “Configure segmented PORS optimization” (page 56).

Bandwidth load spreading

If a single border node has several identical inter-region Passport trunks (identical implies that all link-reachable addresses and path costs are the same), then all traffic will be sent to that inter-region Passport trunk with the lowest delay. This will continue until that link is fully booked.

Load spreading allows that traffic to be distributed across all inter-region Passport trunks based on available bandwidth during normal call establishment. To activate the load spreading capability, configure the *overrideTrunkDelay* attribute. Once configured, the *overrideTrunkDelay* attribute will override the measured trunk delay of a given inter-region Passport trunk, with the resulting value being used to determine route selection.

> ***Note:*** The *overrideTrunkDelay* attribute is optional and configurable under the trunk path administrator. Refer to “Display of overrideTrunkDelay attribute” (page 138) for details.

Figure 21
Display of overrideTrunkDelay attribute

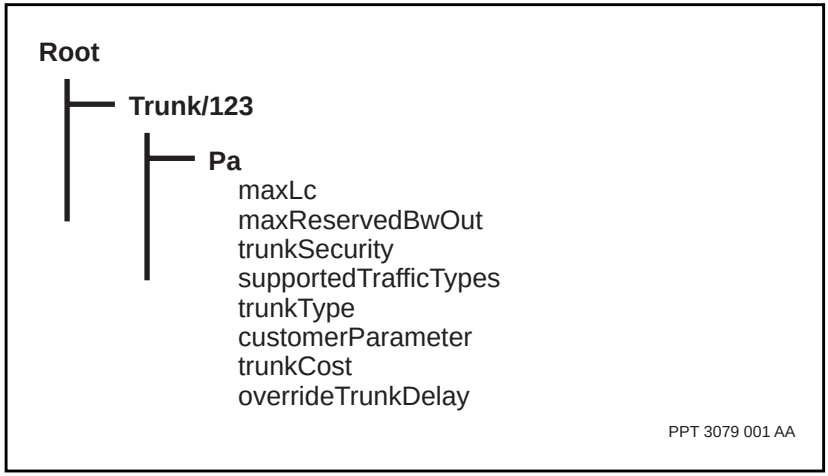

The *overrideTrunkDelay* attribute can be configured on identical inter-region Passport trunks, ensuring that available bandwidth is the final criteria when determining route selection over identical inter-region Passport trunks. This

results in traffic being load shared across all identical inter-region Passport trunks that have been provisioned with identical addresses, region and path costs, and trunk delay attributes.

Load spreading by bandwidth is not possible when a failure causing multiple connections to reroute occurs (for example, a link going down). Due to connection setup latency, any calls established during this interval are rerouted, while the topology database does not receive any bandwidth updates.

Manual path connection between topology regions

Manual paths are defined at each end point of a call. Once a manual path is specified, it bypasses the PORS automatic route selection. Call establishment using a manual path is performed independent of topology regions and can serve as an alternative to dynamically routed calls across multiple topology regions. NSAP addressing is not required when using a manual path to route an inter-region call.

It is recommended to have two independent routes specified for a manual path by specifying a different manual path (non-symmetric) for each end. The first path to come up is used to route the call. The alternate path is used to reroute the call if a failure occurs on the first path.

Figure 22, “PORS inter-region manual path connections,” (page 140) illustrates two manual paths configured between topology regions.

Figure 22
PORS inter-region manual path connections

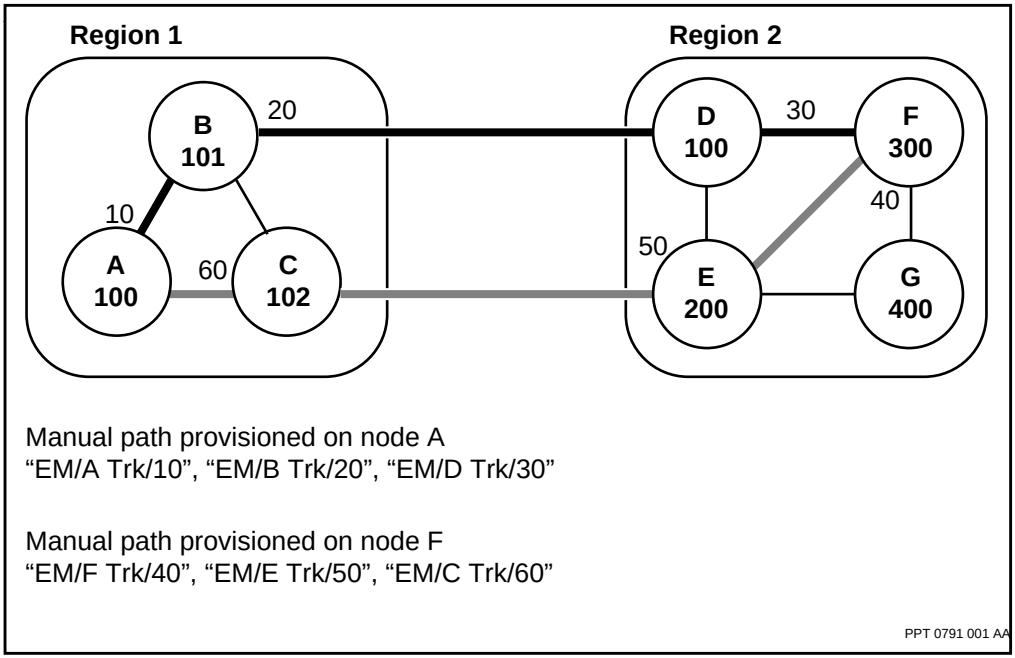

When configuring a manual path connection, remember the following details:

- Configure different paths at each end point because, if both ends use the same path, a single point of failure can occur.
- Select paths that have similar metrics.
- Duplicate nodeIds can be used on a manual path.

To configure a manual path connection, see "Configuring optional manual paths for managed cut through switching" (page 31) and "Configuring optional manual paths for a PLC" (page 33).

Passport clusters

PORS service support for Passport clusters is achieved by deploying NSAP addressing. See "Configuring NSAP addressing" (page 38) for details. PORS support for Passport clusters is achieved by deploying routing gateway

functionality and reachable addresses on cluster links. See “Configuring a routing gateway” (page 53) and “Configuring inter-region Passport trunks” (page 55) for details.

From a PORS perspective, routing to or from a cluster is the same as routing across a topology region boundary. As with multi-topology region routing, path establishment consists of multiple path segments: from the source node within one cluster, across the cluster link, within the backbone of the first region, across the inter-region link (IRL), within the backbone of another region, across the cluster link, and within another cluster to the destination node.

On cluster border node cluster links, configure reachable NSAP addresses to the backbone and beyond. If the NSAP address plan allows for summarization of reachable NSAP addresses from the cluster border node to the backbone and beyond, configure summary addresses on the cluster border node.

On backbone border node cluster links, configure reachable NSAP addresses to the cluster. If the NSAP address plan allows for summarization of reachable NSAP addresses from the backbone border node to the cluster, configure summary addresses on the backbone border node.

Reachable NSAP addresses that are not summarized by summary addresses configured on a cluster border node will be broadcast to all nodes in the cluster. In the same way, reachable NSAP addresses that are not summarized by summary addresses configured on a backbone border node will be broadcast to all nodes in the backbone.

The *TrkReachableAddress* component (under the *Trunk* component) is similar to the *SummaryAddress* component, the difference being that it is link-based as opposed to the *SummaryAddress* component, which is node-based. The *SummaryAddress* component provides a summary of the addresses on the entire node, whereas the *TrkReachableAddress* component specifies the reachable addresses on a specific link on that node. NSAP addresses configured using the *TrkReachableAddress* component on the cluster border node (or backbone border node) links are not broadcast in the cluster (or backbone) if they can be summarized by those addresses configured in the *SummaryAddress* component.

For information on configuring Passport clusters, refer to “Configuring Passport clusters” (page 39).

Guidelines for Passport clusters

Passport clusters are deployed when planning for large network growth, or when approaching the engineering limits for a topology region. Following are some guidelines for Passport clusters:

- Passport trunks must not be configured between clusters.
- Clusters are configured around a topology region backbone.
- Topology regions are interconnected through backbone nodes rather than cluster nodes.
- All planned cluster nodes and backbone nodes to be connected to a cluster node must first be upgraded to a software level supporting the clusters behavior.

Chapter 9
Traffic management

In the Passport Path-Oriented Routing System (PORS), traffic management consists of bandwidth management, adaptive routing, and congestion management.

This section describes these traffic management capabilities. Information is organized under the following headings:

Introduction to bandwidth management

PORS reserves bandwidth for at the connection level for the duration of a logical connection. This approach partitions the network bandwidth among the different services according to priority settings.

The three principal aspects to bandwidth management in a Passport network are:

- path bandwidth reservation, which determines how bandwidth is allocated to an individual path (see the section "Path bandwidth reservation" (page 144))

- trunk bandwidth sharing, which determines how bandwidth is partitioned and shared on a Passport trunk between path-oriented and connectionless traffic (see the section “Passport trunk bandwidth sharing” (page 150))
- dynamic trunk speed changes, which enables Passport trunks and routing to adapt to changes resulting from increasing and reducing bandwidth (see the section “Dynamic trunk speed change” (page 156))

Path bandwidth reservation

Every logical connection has a configured *requiredTxBandwidth* attribute and *requiredRxBandwidth* attribute, which is a pair of numbers indicating the bandwidth in bits per second to reserve in each direction of the connection.

These numbers can represent the peak, average, or an intermediate bandwidth. The exact meaning is controlled by the service reserving the bandwidth.

This section includes the following information:

Maximum reserved outbound bandwidth

Every Passport trunk has a *maxReservedBwOut* attribute indicating how much bandwidth can be reserved on this Passport trunk in the outgoing direction. This reserved bandwidth is a percentage of the total Passport trunk bandwidth, which is indicated by the Passport trunk attribute *measuredSpeedToIf*.

PORS subtracts the values for the *requiredBandwidth* attribute (*requiredTxBandwidth* and *requiredRxBandwidth* attributes) from the available remaining reservable bandwidth in each direction of a Passport trunk.

Example of bandwidth reservation

Consider the network illustrated in part (a) of the figure “Path bandwidth reservation” (page 144), with

- tandem nodes C and D connected by a DS-3 trunk
- peripheral nodes A, B, E, and F each connected by a DS-1 trunk to one of the tandem nodes

Part (b) of the figure “Path bandwidth reservation” (page 144) illustrates the topology and state after VC/1 has reserved bandwidth on route A->C->D->F, with the *requiredBandwidth* attribute set to (1.0M, 1.54M).

Part (c) of the figure “Path bandwidth reservation” (page 144) illustrates the topology and state after VC/1 has reserved bandwidth as described above, and then VC/2 has reserved bandwidth on route E->D->C->B, with the *requiredBandwidth* attribute set to (1.54M, 1.54M).

> ***Note:*** In the statements above, the route specifications are simplified for illustrative purposes. Also, the first number in brackets is the bandwidth from source to destination (*requiredTxBandwidth* attribute), and the second number is the bandwidth from destination to source (*requiredRxBandwidth* attribute).

Figure 23
Path bandwidth reservation

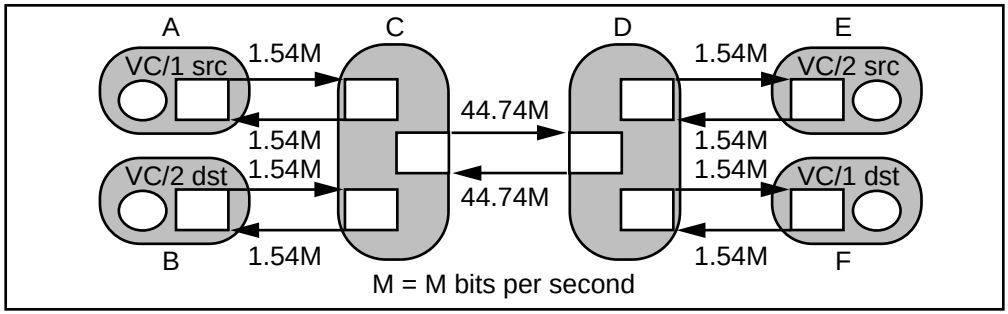

(a) Initial Network State (after VCs initialize but before path is setup)

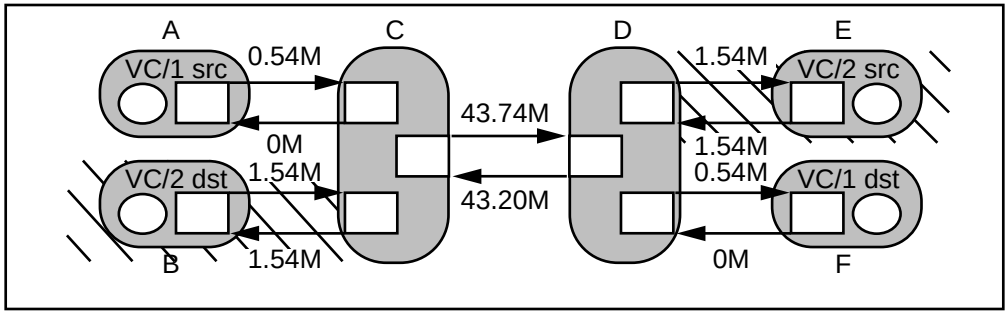

(b) After setup of VC1 with path A->C->D->F and requiredBandwidth (1.0,1.54)

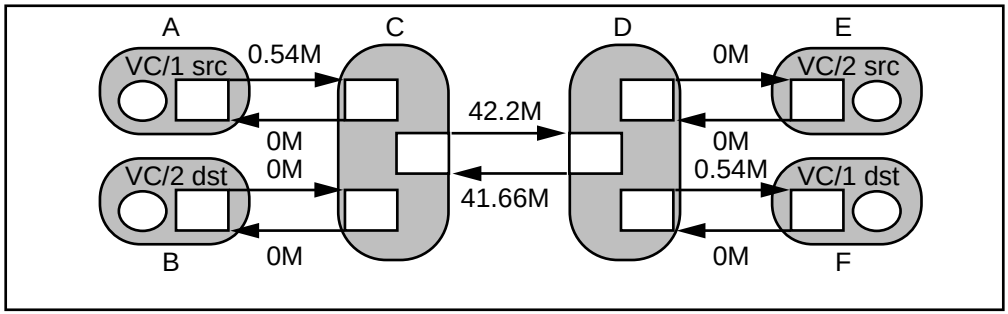

(c) After setup of VC2 with path E->D->C->B and requiredBandwidth (1.54,1.54)

PPT 0805 001 AA

Path instantiation failures: causes

A path can fail to instantiate when it reaches a Passport trunk that does not have any more free LCNs to allocate. The maximum number of LCNs that the Passport trunk can support is specified in the *PathAdmin maxLc* attribute. A path can also fail to instantiate when it reaches a Passport trunk that does not have sufficient bandwidth to support the logical connection. The *maxReservedBwOut* attribute of a Passport trunk indicates the upper limit of the bandwidth that can be reserved for all instantiated paths.

A path can fail to instantiate because of lack of reservable bandwidth caused by unavoidable discrepancies between the Route Selector view of reserved bandwidth and the actual bandwidth reserved on Passport trunks in the network. A Route Selector is unaware of the reservations made for routes chosen by other Route Selectors or bandwidth unreserved when paths terminate. It learns of the changed bandwidth reservation of a Passport trunk either from the regular topology updates, or from path instantiation failures of its own routes going over the affected Passport trunk. Path terminations with significant bandwidth change can also be learned by the RS at the end points.

Path instantiation failures: retry response

In case of failure, path instantiation is terminated and the reserved bandwidth status of (at least) the current Passport trunk is fed back with the failure packets to the end point originating the logical connection. See part (a) of the figure "Learning about bandwidth errors after a setup failure" (page 148). The end point forwards these new bandwidth values to its Route Selector to update its local view of the topology. See part (b) of the figure "Learning about bandwidth errors after a setup failure" (page 148).

Figure 24
Learning about bandwidth errors after a setup failure

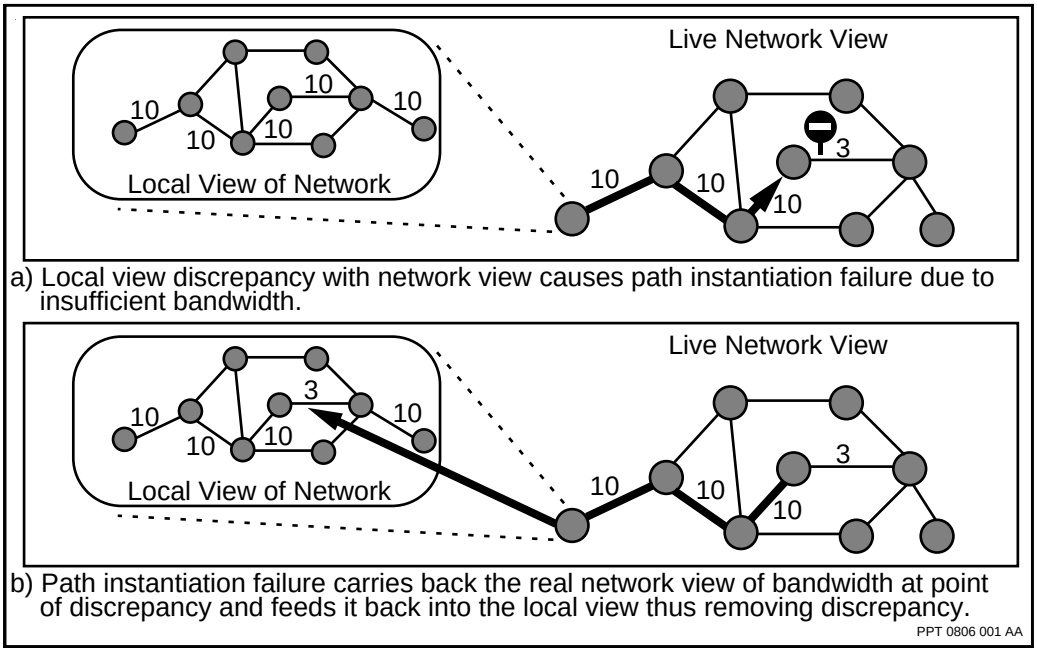

This error feedback mechanism is how the Route Selector learns about its mistakes in order to avoid making them repeatedly. Once the actual reserved bandwidth has been fed back into the local topology, the Route Selector is asked by the end point to choose a new route. If a new route is found, path instantiation proceeds again on the new route. This process continues until either the path successfully instantiates, or the Route Selector is unable to find a route with sufficient bandwidth. If unable to find a route with sufficient bandwidth, the Route Select resorts to bumping of existing paths to obtain their bandwidth. This is explained in the section "Path instantiation failures: bumping response" (page 148).

Path instantiation failures: bumping response

"Path instantiation failures: causes" (page 147) gives the impression that there are only two kinds of bandwidth, free and reserved. PORS actually tracks seven types of bandwidth for each Passport trunk. These types are

- bandwidth currently unreserved

- bandwidth reserved for paths at holding priorities 4 (lowest), 3, 2, 1 and 0 (highest)
- minimum bandwidth reserved for connectionless traffic

These bandwidth types are important. To obtain bandwidth when the network cannot establish a path due to insufficient bandwidth, PORS bumps paths with a lower holding priority than the setup priority for the new path. The perspective of reservable bandwidth depends on the *setupPriority* attribute value assigned to the new path. The bandwidth available for a path at *setupPriority* attribute n consists of the unreserved bandwidth plus all bandwidth reserved by paths at *holdingPriorities* attribute greater than n. See the figure "Bandwidth available at different holding and setup priorities" (page 149).

Note: The highest holding and setup priorities are 0 and the lowest are 4.

Figure 25
Bandwidth available at different holding and setup priorities

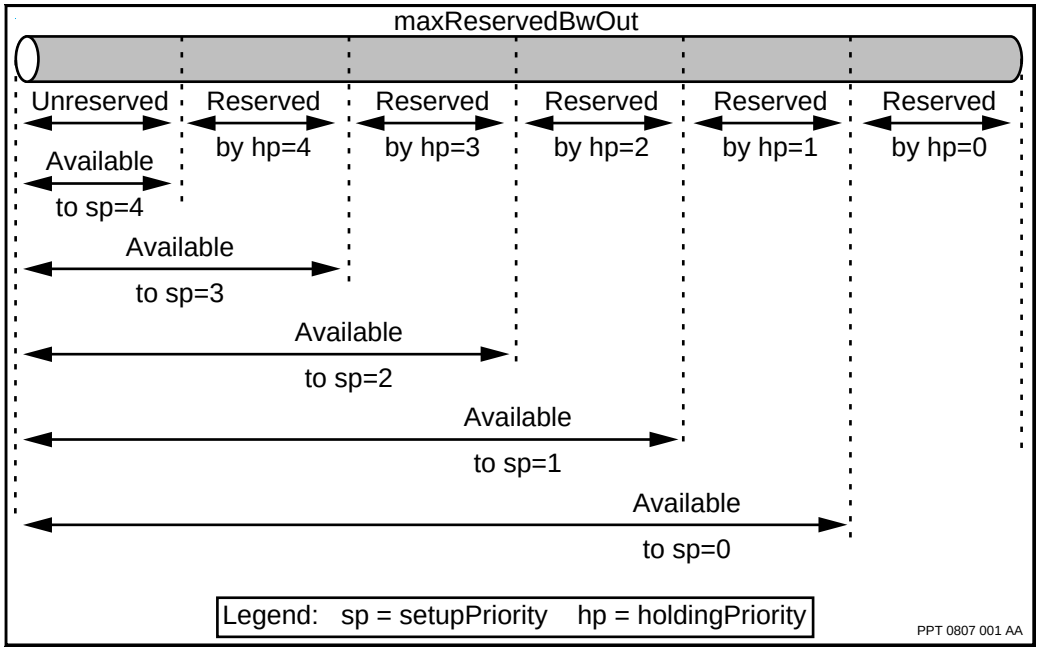

The Route Selector mechanism for selecting a route is governed by the *bumpPreference* attribute under the *PermanentLogicalConnection* component.

When the *bumpPreference* attribute is set to bumpWhenNecessary, the Route Selector always tries first to establish a route through Passport trunks that have enough unreserved bandwidth to satisfy the value set for the *requiredBandwidth* attribute. If unable, it adds to the amount of available bandwidth on all *holdingPriority* attribute 4 bandwidth and tries again, if the new PLC *setupPriority* attribute is 3 or better. This process continues until it reaches a *holdingPriority* attribute equal to the *setupPriority* attribute, at which point it reports that no route is possible.

If adding in bandwidth used by lower holding priority paths results in a successful route, the path instantiation for that route is permitted to bump lower holding priority paths to obtain their bandwidth. Paths that are bumped are setup on alternate routes, if available. If the bumped path has the *pathFailureAction* attribute set to disconnectConnection, the bumped path is terminated instead.

When the *bumpPreference* attribute is set to bumpToObtainBestRoute, the Route Selector adds to the amount of available bandwidth. The amount added equals all bandwidth reserved by PLCs with a holding priority that is numerically greater than the setup priority for the new PLC. This is because when the *bumpPreference* attribute under the *PermanentLogicalConnection* component is set to bumpToObtainBestRoute and its *setupPriority* attribute is set to 2 then this PLC can bump all paths whose *holdingPriority* attribute is 3 or 4 in order to obtain the best possible path.

> ***Note:*** The Route Selector does not bump lower *holdingPriority* attribute PLCs unless it is absolutely necessary even when the *bumpPreference* attribute is set to bumpToObtainBestRoute.

Passport trunk bandwidth sharing

A Passport trunk can carry two broad classes of traffic, connectionless and path-oriented. Connectionless traffic travels through a Passport trunk without making any advanced reservations as to its bandwidth requirements, unlike

path-oriented traffic which travels along pre-setup paths that have individually reserved some portion of the total Passport trunk bandwidth for their use on every Passport trunk they traverse.

This section addresses how Passport fairly and efficiently shares Passport trunk bandwidth between connectionless and path-oriented traffic and includes the following information:

- "Passport trunk bandwidth allocation" (page 151)
- "Passport trunk bandwidth sharing under congestion" (page 155)
- "Guidelines for setting Passport trunk bandwidth attributes" (page 155)

Passport trunk bandwidth allocation

The network engineers determine fair and efficient sharing of resources on a Passport trunk. The Trunk PathAdmin component provides configurable parameters that allow different enforceable policies.

The *maxLc* attribute defines the maximum number of individual logical channels (or paths) that traverse the trunk. When this number is reached, no new paths can instantiate over this Passport trunk until some existing paths clear.

> ***Note:*** The value set for the *maxLc* attribute should be added to the *connectionPoolCapacity* attribute under the *Lp Eng Arc Ov* component. This ensures that PORS connections will establish on map mode trunks. For details, see 241-5701-710 *Passport 7400, 15000, 20000 ATM Configuration Guide*.

The *maxReservedBwOut* attribute defines the percentage of total Passport trunk bandwidth that PORS can allocate among individual Logical Channels. Once this percentage is reached, the Passport trunk has no more reservable bandwidth. No paths can instantiate over this Passport trunk until some existing paths clear. For example, on a DS-1 trunk at 1.536 Mbit/s, a value of 65% for this Passport trunk attribute makes this Passport trunk capacity appear to be 0.9984 Mbit/s, and hence PORS never reserves more than 0.9984 Mbit/s of this Passport trunk. The actual utilization of path-oriented traffic can exceed the reservation limit since bandwidth utilization is not enforced against the bandwidth reservation. This allows for statistical utilization of the bandwidth on the Passport trunks.

When configuring a Passport trunk, the apportioning of the total bandwidth between the path-oriented reservation (*maxReservedBwOut* attribute) and connectionless (the remainder) is flexible. The range for the *maxReservedBwOut* attribute is 0 - 100% however only 80% maximum is recommended at this time if the traffic is purely interrupting (Emission Priority 0). A small amount of bandwidth is reserved for connectionless so that the metric is not infinity. This is illustrated in the figure “Passport trunk bandwidth reservation limit control options” (page 153).

Figure 26
Passport trunk bandwidth reservation limit control options

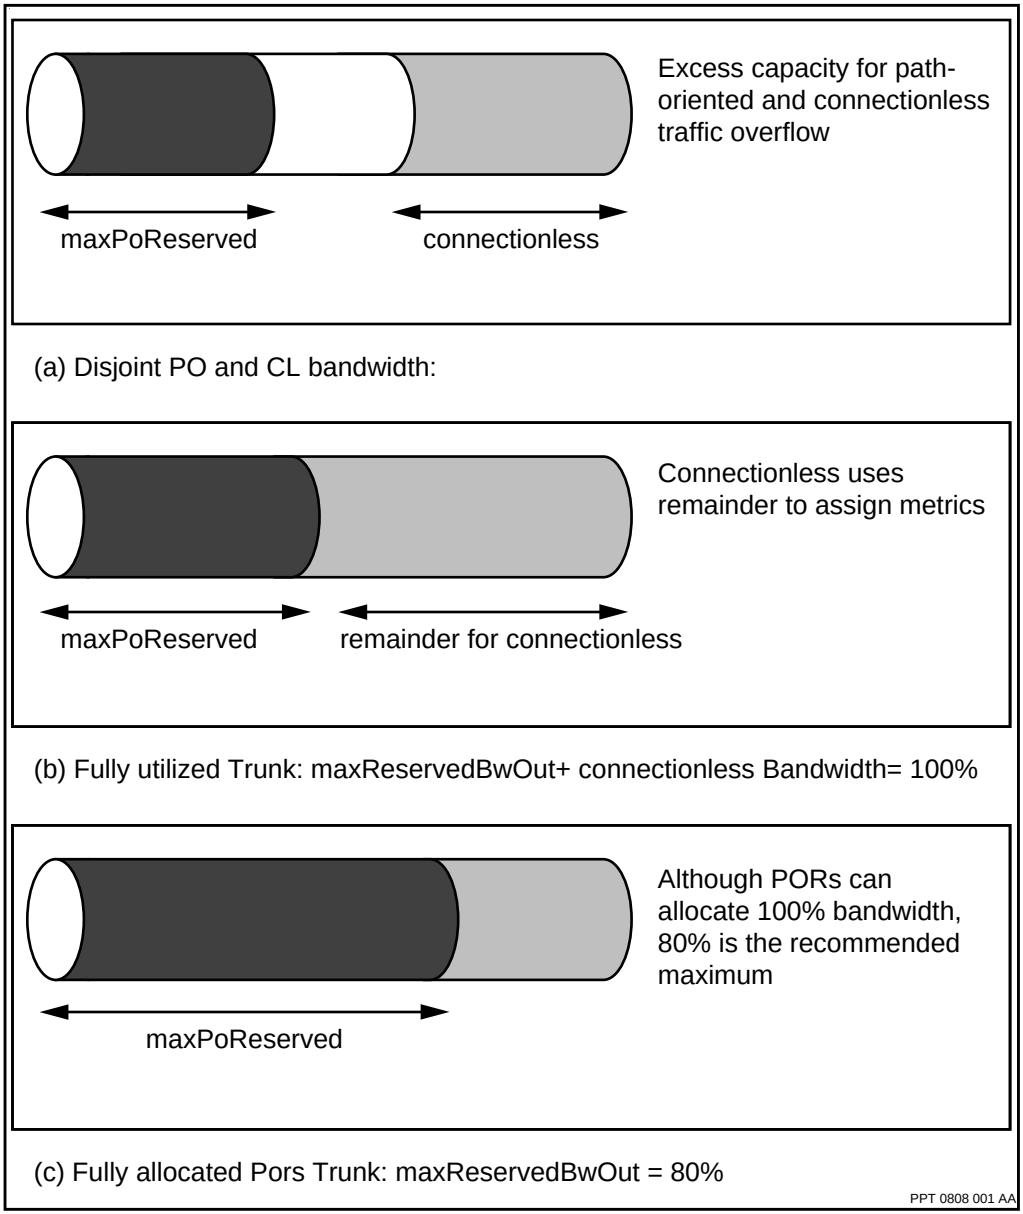

If the path-oriented to connectionless boundary for a Passport trunk is non overlapping, the path-oriented traffic is less likely to use the connectionless share of the bandwidth, especially if the value set for the *requiredBandwidth* attribute for each path has been specified as a peak value. See part (a) of the figure “Passport trunk bandwidth reservation limit control options” (page 153). This represents a conservative allocation of Passport trunk bandwidth. If the path-oriented traffic is not using its entire reservation, the throughput routes chosen by connectionless traffic cannot advertise the unused path-oriented bandwidth. However, if connectionless traffic exceeds 100 - the *maxReservedBwOut* attribute it can take advantage of the unused path-oriented reserved bandwidth. Similarly, the path-oriented traffic exceeding the *maxReservedBwOut* attribute can take advantage of the unused connectionless bandwidth.

If the traffic utilizations are disjoint, there is a safety zone where connectionless and path-oriented traffic can overflow. See part (a) of the figure “Passport trunk bandwidth reservation limit control options” (page 153). This gives a lower link utilization, but has a lower probability of contention between path-oriented and connectionless traffic, and thus has minimal packet loss.

The *maxReservedBwOut* attribute is a fairly coarse control of the maximum percentage of a Passport trunk which can be occupied by path-oriented traffic. It is coarse because nothing prevents the applications using the paths from over utilizing or under utilizing their reservations. For continuous bit rate applications, the reservations are very accurate, but for bursty applications, the reservation can be under-shot (if reservation is peak) or over-shot (if reservation is statistical). The accuracy with which the *maxReservedBwOut* attribute represents path-oriented Passport trunk utilization is a function of how many individual paths are using this bandwidth and the variance of each of these paths from its individual reservation. Typically, a larger number of small statistical reservations have a more predictable sum than a small number of large statistical reservations. See the figure “Passport trunk path-oriented bandwidth fluctuation about its reservation” (page 155).

Figure 27
Passport trunk path-oriented bandwidth fluctuation about its reservation

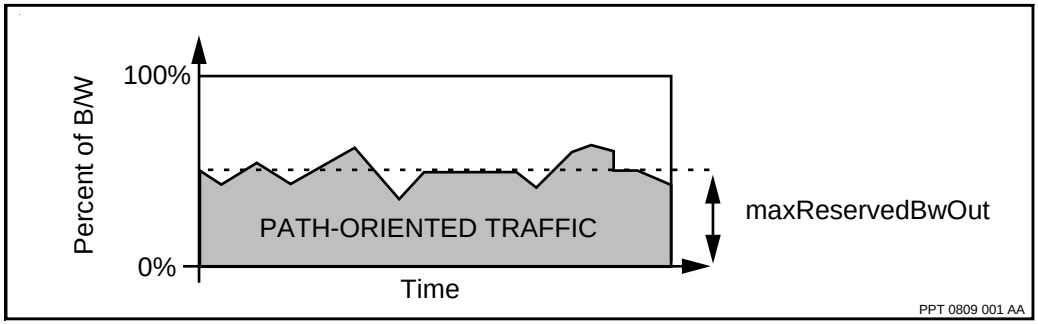

At some time, the actual percentage utilizations of path-oriented and connectionless traffic can add up to greater than 100%, and we get a congested Passport trunk. The frequency and duration of these congestion periods are directly related to the error in predicting the offered loads of each class of traffic.

Passport trunk bandwidth sharing under congestion

Packets are not discarded by Passport trunks when they are under no congestion, regardless of how much a particular application or class of traffic is exceeding its reservation. When congested, packets must be discarded while still providing fair sharing of the Passport trunk bandwidth between path-oriented and connectionless traffic. If there is congestion the different packets are discarded by respecting the discard priorities of the individual data packets not with regard to the reservation limits.

Guidelines for setting Passport trunk bandwidth attributes

Given a number of configurable attributes all of which affect the bandwidth used by various classes of traffic on a Passport trunk, how do we set them to achieve a desired result?

There is a trade-off between the efficient usage of a Passport trunk and the probability of loss due to short term congestion on that Passport trunk. This trade-off is highly dependent on the mix of traffic on the Passport trunk. Loss occurs when bursty traffic peaks align, and inefficiency occurs when extra bandwidth is allocated to make room for the small bursts to reduce the loss. The figure “Efficiency versus loss rate on a Passport trunk carrying bursty

traffic" (page 156) gives an approximation of this, but the steepness and the shape of the curve are traffic dependent, and it is impossible to give a precise graph.

Figure 28
Efficiency versus loss rate on a Passport trunk carrying bursty traffic

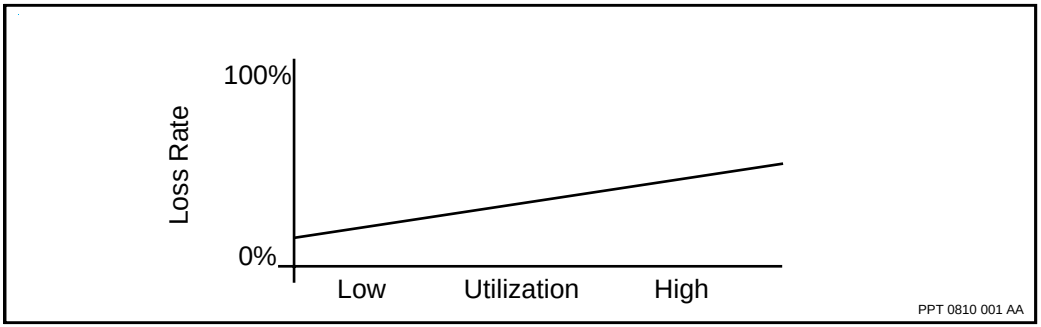

An important decision to be made is how much loss can the network inflict on its users to ensure that the Passport trunk is being used efficiently.

If the objective of the network operator is to give the lowest possible loss rate to its subscribers, the *maxReservedBwOut* attribute is assigned a conservative setting and the individual services accurately reflect the average (peak in some cases) bandwidth utilization.

If the operator wants maximum efficiency of Passport trunk bandwidth, connectionless traffic must be able to take advantage of any lulls in path-oriented traffic, and vice versa. To achieve this more statistical bandwidths can be configured on the individual services.

Dynamic trunk speed change

The dynamic trunk speed change feature enables Passport trunking and routing to adapt to changes that result from bandwidth deltas without taking Passport trunks out of service. This feature also supports Passport trunking on the integrated Passport Inverse Multiplexing for ATM (IMA). For details on IMA, refer to 241-5701-730 *Passport 7400, 15000, 20000 Inverse Multiplexing for ATM Guide*.

The dynamic trunk speed change feature enables a Passport system to interwork with third-party equipment to offer the following functionality:

- ISDN dial backup—resumes connectivity between two Passport nodes when a leased line failure occurs
- ISDN dynamic bandwidth on demand (BWoD)—enables a Passport node to adapt to additional bandwidth when the capacity of the dedicated facility is exceeded
- inverse multiplexing—enables a Passport trunk to support more than four physical connections between adjacent Passport nodes

For details on these functionalities of the dynamic trunk speed change feature, refer to 241-5701-420 *Passport 7400, 15000, 20000 Trunking Guide*.

The dynamic trunk speed change feature offers the following benefits:

- PORS configurable attributes enable selective bumping of connections when bandwidth decreases (see “Path instantiation failures: retry response” (page 147)).
- Changing the *maxReservedBwOut* attribute of the Passport trunk bandwidth is a non-critical operation and affects the current PORS connections on a Passport trunk.
- Changing the *trunkCost* attribute of the Passport trunk bandwidth is a non-critical operation and can attract new PORS connections to a Passport trunk.
- All existing PORS features, such as bumping and optimization, are maintained.

Speed change reporting mechanism

A set of Passport trunk attributes enable and disable the reporting of dynamic trunk speed changes to the routing mechanism. By default, the speed change reporting mechanism is disabled on the Passport trunk, and Passport trunk speed variations are not propagated to PORS.

When the speed change reporting mechanism is enabled, significant Passport trunk speed variations are automatically propagated to PORS through the PA. A speed decrease is deemed more significant than a speed increase. PORS then makes a decision (based on the holding priority) whether or not the speed change is propagated to other nodes in the network.

Dampening the generation of speed change reports

To prevent the routing system from an overload of speed change updates, the speed change reporting mechanism can be customized (such as the granularity of the reported speed change and the frequency of the speed change update) through configuring *Trunk speedReportingThresholds* and *speedReportingHoldOff* attributes.

> ***Note:*** The *speedReportingHoldOff* attribute is not applied to speed decrease updates. Speed decreases are always reported immediately to discourage traffic from using the Passport trunk.

For configuration information, see 241-5701-420 *Passport 7400, 15000, 20000 Trunking Guide*.

PORS trunks

If a PORS trunk is used, a percentage of the Passport trunk speed is allocated to Path-Oriented Routing traffic. A *PathAdmin* component under each *Trunk* component specifies the bandwidth sharing characteristics of the Passport trunk.

Voice services (a part of PORS traffic) uses a Short Path-Oriented Multiplexing (SPO-mux) feature. This PORS efficiency feature uses a separate VCC with no ATM adaptation layer for increased efficiency. The rest of the PORS traffic uses AAL5 VCC. The speed for the SPO-mux feature is included in the PORS portion of the reserved bandwidth Passport trunk. When the allocated bandwidth for the Passport trunk decreases, the allocated bandwidth of the PORS traffic is decreased in proportion to the Passport trunk bandwidth.

If the Passport trunk speed change reporting mechanism is enabled when a dynamic speed change occurs on a PORS trunk, the new Passport trunk bandwidth is passed from the Passport trunk to PORS through the *PathAdmin* component. PORS then recalculates a new PORS trunk bandwidth and propagates to other nodes (if required).

For more information on the PORS efficiency feature, see 241-5701-420 *Passport 7400, 15000, 20000 Trunking Guide*.

Congestion and re-routing traffic

If the available bandwidth decreases, the available bandwidth can be overcommitted to the applications. Certain PORS connections are rerouted or bumped to other paths when congestion occurs.

When the PORS bandwidth returns to its original value, PORS (via optimization) can reroute paths that were bumped back to the original Passport trunks.

PORS percentage

Changing the *maxReservedBwOut* attribute of the Passport trunk bandwidth is a non-critical operation (can be modified without disrupting calls); it can be used to affect the current PORS connections on a Passport trunk. Similarly, changing the *trunkCost* attribute of the Passport trunk bandwidth is a non-critical operation and can be used to attract new PORS connections to a Passport trunk.

The PORS bandwidth percentage applies to existing calls, while the Passport trunk cost applies to new calls.

Delays and delay variations

Delays and delay variations are added to the traffic in the transition period (that is, from the time that the leased line fails until the time the backup lines are functional depends on the third-party device used).

These delays and delay variations lead to quality of service (QOS) downgrading and can violate QOS guarantees on voice and video traffic on the PORS trunk. Depending on the functionality provided by the third-party ISDN device, the network can be configured to avoid using the backup lines for PORS traffic.

Congestion management

Congestion occurs whenever the data to be transmitted or processed by a server exceeds the servers capacity. The severity of the congestion is proportional to the persistence and volume of the excess data. An example of congestion is illustrated in the figure "Congestion situation" (page 160).

Figure 29
Congestion situation

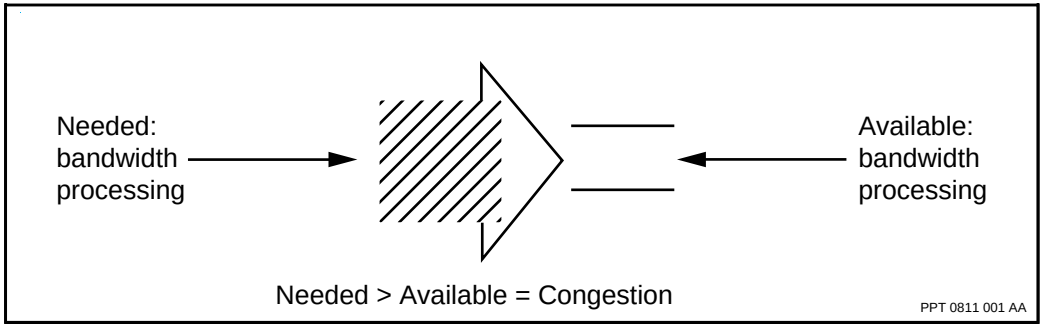

While congested, a system exhibits undesirable behavior. For example

- packets are delayed
- packets are discarded
- path instantiation is delayed or fails

PORS congestion management responds to congestion and maintains the committed grade of service.

This section includes the following information on congestion:

- "Points of congestion" (page 161)
- "Passport trunk congestion" (page 161)
- "Congestion notification" (page 163)
- "Congestion management by packet discard" (page 163)
- "Network management to prevent congestion" (page 164)

Points of congestion

Congestion occurs where processing and transmitting servers interact. Types of congestion include

- processor congestion—when too many packets are received from the bus or Passport trunk
- bus congestion—when too many packets are sent to the bus
- Passport trunk congestion—when too many packets are sent to the Passport trunk

Processor and bus congestion are difficult to proactively identify.

Response to congestion

PORS congestion management focuses on avoiding and resolving Passport trunk congestion with minimal disruption to high priority data. To accomplish this action, PORS

- notifies applications of the onset of congestion
- applies bandwidth sharing rules and discards packets, if required, based on their relative importance

Passport trunk congestion

There are two types of Passport trunk congestion:

- Passport trunk bandwidth congestion—where the offered load exceeds the Passport trunk bandwidth capacity
- Reservation congestion—where the amount of bandwidth reserved for path-oriented traffic is at (or close to) the reservation limit (*maxReservedBwOut* attribute). It can also be caused by reaching the maximum number of paths that can be supported on the Passport trunk (*maxLc* attribute).

These two types of Passport trunk congestion are related as shown in the figure “Types of Passport trunk congestion” (page 162).

Figure 30
Types of Passport trunk congestion

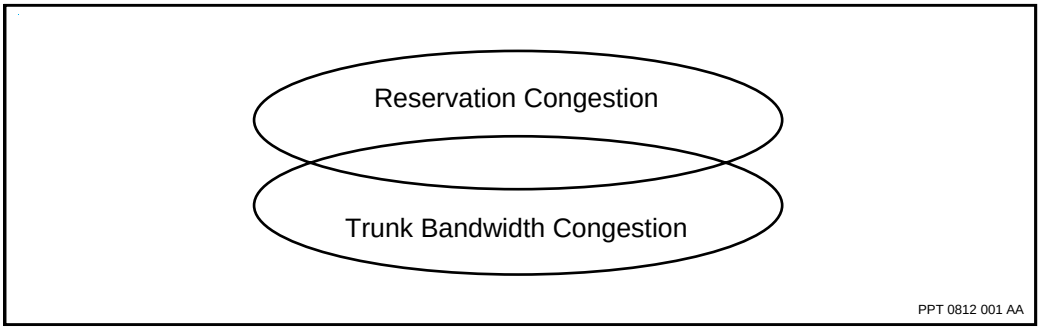

Each type of congestion can occur alone or together with the other type. If reservation congestion occurs simultaneously with Passport trunk bandwidth congestion, this means that the entire *maxReservedBwOut* attribute portion of the Passport trunk bandwidth has been reserved. At the same time, all the Passport trunk bandwidth is being used by the path-oriented traffic, the connectionless traffic, or both. In general, a Passport trunk congestion has impact on paths, on path-oriented data packets, or on both, as depicted in the figure “Impact of congestion types” (page 162).

Figure 31
Impact of congestion types

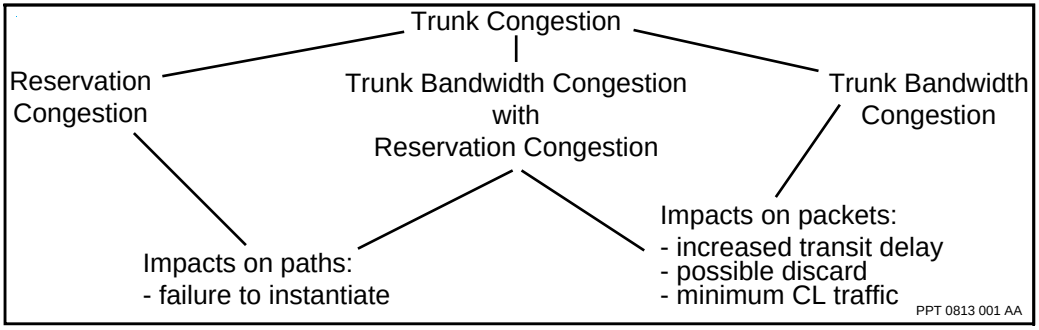

Impact on paths is manifested as a failure of path instantiation for Passport trunks. New paths cannot be instantiated over the Passport trunk experiencing reservation congestion. Impact on packets of already instantiated paths is

manifested by packet delay due to a buildup of packet waiting line, and eventual packet discards if the number of waiting packets grows beyond certain limits.

Congestion notification

When a network becomes congested, packets can be discarded. To avoid packet discard, application end points reduce the amount of traffic they send into the network. Receipt of congestion indication messages triggers this reduction.

PORS transparently transports congestion indication messages to application end points. The congestion management system detects congestion and sets the Forward Congestion Indication (FCI) bit in the packet header of packets travelling in the direction of a congested path. The end point that receives these packets is responsible for notifying the source end point that the path is congested. (In Frame Relay, notification is achieved through the Backward Congestion Indication (BCI) bit on packets sent back to the transmitting end point.) The source end point then reduces the number of packets sent into the network.

HTDS, BTDS, and Voice Transport rely on PORS to transport data in a Passport network. Voice Transport uses FCI to signal congestion and then Voice Transport sends an explicit downspeed or rate control indication to the transmitting side to dynamically reduce the number of packets sent in the network.

Congestion management by packet discard

The Passport trunk bandwidth congestion is volatile since it depends on the process of packet arrivals and packet lengths. Packet discard is invoked when the congestion persists long enough for the emission queues to grow beyond a certain threshold.

The packet discarding mechanism implemented by PORS reduces the rate at which new packets are put into the emission queues. This mechanism is a quick system response to prevent emission buffer overflow which in turn results in uncontrolled packet discarding. Packet discards are based on

- severity of the congestion situation
- packet discard eligibility

- path discard priority
- Passport trunk bandwidth sharing mechanism between connectionless and path-oriented traffic

Packets are discarded beginning with packets marked as Discard Eligible (DE). The packet DE bit is set by an application when emitting a packet that can be preferentially discarded in times of congestion. The intent is to discard first the non essential packets or packets transmitted by an application above the requested bandwidth for the service.

If the congestion becomes more severe, non DE-bit packets are discarded, based on the discard priority of the path. The path discard priority is configured through with the *discardPriority* attribute under the *PermanentLogicalConnection* component. This attribute defines the importance of packets which in turn reflects the probability of discard. A packet with a high discard priority is only discarded during severe congestion, while a low discard priority packet is discarded under lower congestion.

Packet discards are done to respect the configured bandwidth sharing mechanism between connectionless and path-oriented traffic. This mechanism is only invoked under a congestion situation since otherwise, there is enough bandwidth for both types of traffic. The section “Passport trunk bandwidth sharing under congestion” (page 155) explains this mechanism in detail.

Network management to prevent congestion

Congestion is an indication that the collective bandwidth reservation on a Passport trunk does not match the actual bandwidth usage. By increasing individual path bandwidth reservation requests (*requiredBandwidth* attribute) to more closely reflect usage, you more evenly disperse traffic among the Passport trunks. Decreasing the amount of reservable bandwidth on a Passport trunk (*maxReservedBwOut* attribute) reduces the number of paths that can be instantiated. Decreasing the number of instantiated paths (*maxLc* attribute) is also an alternative, but this configuration cannot produce the desired result since a fewer number of paths can still produce congestion if collectively they reserve the same amount of bandwidth as many more paths.

Decreasing the path-oriented bandwidth reservation limit (*maxReservedBwOut* attribute) reduces the bandwidth contention between the two traffic types. Decreasing the *maxReservedBwOut* attribute reduces the path-oriented traffic load (as explained in the previous paragraph), but does not affect the connectionless traffic load. Also by configuring the individual path-oriented services to reflect accurately the bandwidth, in some cases the peak bandwidth for continuous bit rate services such as BTDS, the PORS traffic can be kept within its reservation limits.

As a long term network management reaction to congestion problems, more bandwidth can be provided on frequently congested routes. Increased bandwidth is achieved by one of three methods:

- replacing an existing Passport trunk with one of greater capacity
- adding another Passport trunk along side a congested Passport trunk
- adding Passport trunks between other nodes to increase the probability of choosing another route

Passport 7400, 15000, 20000

Path-Oriented Routing System Guide

Release 5.2

Publication: 241-5701-435
Document status: Standard
Document version: 5.2S1
Document date: November 2003
Printed in Canada

NORTEL NETWORKS