NORTEL NETWORKS

Passport 7400, 15000, 20000

Networking Overview

241-5701-400

Passport 7400, 15000, 20000

Networking Overview

Publication: 241-5701-400
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

The Passport Networking Introduction document provides background information on Passport networking. After reading this introduction, you will have a general understanding of networking in a Passport network.

The following topics are discussed in this section:

Who should read this document and why

This introduction is intended for people performing one or more of the following tasks:

- traffic engineering
- setting up a network
- incorporating one or more Passport nodes into an existing Passport network

What you need to know

You need a basic understanding of

- general principles of packet switching
- general principles of circuit switching
- general networking principles
- communication products, particularly switching products

How this document is organized

The 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*, is organized as follows:

- "Passport networking architecture" (page 21) explains the features and architecture of the Passport trunking, routing and traffic management systems.
- "Passport trunking and base routing systems" (page 49) describes the Passport trunking and base routing layers including topology manager, and transport resource manager.
- "Dynamic Packet Routing System (DPRS)" (page 85) describes the Dynamic Packet Routing System (formerly referred to as DPN routing and RID/MID routing) architecture and routing mechanisms.
- "Path-Oriented Routing System (PORS)" (page 119) describes the PORS architecture and routing mechanisms.
- "ATM routing system" (page 141) describes Asynchronous Transfer Mode routing architecture and mechanisms.
- "Passport addressing and call routing" (page 183) provides information on addressing and call routing as it applies to the Passport network.

What's new in this document

There were no new features added to this document.

Other changes made to this document include the following:

- Updated the section "MID" (page 95) to show new engineering limitations.

Text conventions

This document uses the following text conventions:

- `nonproportional spaced plain type`

  Nonproportional spaced plain type represents system generated text or text that appears on your screen.

- **`nonproportional spaced bold type`**

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

- *italics*

  Statements that appear in italics in a procedure explain the results of a particular step and appear immediately following the step.

  Words that appear in italics in text are for naming.

- `[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.

- UPPERCASE,lowercase

  Passport commands are not case-sensitive and do not have to match commands and parameters exactly as shown in this document, with the exception of string options values (for example, file and directory names) and string attribute values.

- `|`

  This symbol separates items from which you may select one; for example, ON|OFF indicates that you may specify ON or OFF. If you do not make a choice, a default ON is assumed.

- ...

  Three dots in a command indicate that the parameter may be repeated more than once in succession.

The term absolute path name refers to the full specification of a path starting from the root directory. Absolute path names always begin with the slash (/) symbol. A relative path name takes the current directory as its starting point, and starts with any alphanumeric character (other than /).

Related documents

Passport networking systems are introduced in this document. More detailed descriptions and instructions are provided in the networking-specific NTPs and related service guides.

See the following documents for related information:

General references

For general information, see

- 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*
- 241-7401-200 *Passport 7400 Hardware Description*
- 241-1501-200 *Passport 15000, 20000 Hardware Description*
- 241-5701-445 *Passport 7400, 15000, 20000 Multiprotocol Label Switching Guide*

Passport trunking references

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

DPRS references

For DPRS information, see

- 241-5701-425 *Passport 7400, 15000, 20000 Dynamic Packet Routing System Guide*

- 241-5701-901 *Passport 7400, 15000, 20000 Frame Relay Fundamentals*
- 241-5701-902 *Passport 7400, 15000, 20000 Configuring Frame Relay*
- 241-5701-410 *Passport 7400, 15000, 20000 Call Redirection Server Guide*
- 241-5701-415 *Passport 7400, 15000, 20000 Hunt Group Server Guide*
- 241-7401-110 *Passport 7400, DPN-100 Interworking Guide*

IP references

For IP services information, see

- 241-5701-805 *Passport 7400, 15000, 20000 Understanding IP*
- 241-5701-810 *Passport 7400, 15000, 20000 Configuring IP*

VPN references

- 241-5701-581 *Passport 7400, 15000, 20000 Basics: VPN Fundamentals*
- 241-5701-582 *Passport 7400, 15000, 20000 VPN Configuration Management*

PORS references

For PORS information, see

- 241-7401-750 *Passport 7400 Voice Transport Guide*
- 241-7401-770 *Passport 7400 HDLC Transparent Data Service Guide*
- 241-7401-775 *Passport 7400 Bit Transparent Data Service Guide*
- 241-5701-720 *Passport 7400, 15000, 20000 AAL1 Circuit Emulation Guide*
- 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide*
- 241-7401-440 *Passport 7400 Frame Relay Managed Cut-through Switching Guide*

ATM routing references

For ATM information, see

- 241-5701-420 *Passport 7400, 15000, 20000 Trunking Guide*

- 241-5701-702 *Passport 7400, 15000, 20000 ATM Routing and Signaling Fundamentals*
- 241-7401-440 *Passport 7400 Frame Relay Managed Cut-through Switching Guide*

Traffic management references

For traffic management information, see

- 241-5701-901 *Passport 7400, 15000, 20000 Frame Relay Fundamentals*
- 241-5701-902 *Passport 7400, 15000, 20000 Configuring Frame Relay*
- 241-5701-700 *Passport 7400, 15000, 20000 ATM Overview*
- 241-5701-705 *Passport 7400, 15000, 20000 ATM Traffic Management Fundamentals*
- 241-5701-706 *Passport 7400, 15000, 20000 ATM Traffic Shaping and Policing*
- 241-5701-707 *Passport 7400, 15000, 20000 ATM Queuing and Scheduling*
- 241-5701-708 *Passport 7400, 15000, 20000 ATM CAC and Bandwidth Management*

Call routing references

For call routing information, see

- 241-5701-405 *Passport 7400, 15000, 20000 Call Server Guide*
- 241-5701-410 *Passport 7400, 15000, 20000 Call Redirection Server Guide*
- 241-5701-415 *Passport 7400, 15000, 20000 Hunt Group Server Guide*

How to get more help

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

Chapter 1
Passport networking architecture

Networking systems provide the infrastructure required by application services for forwarding traffic within a network from its source to its destination. Networking systems perform the following functions:

- identify the route or logical destination the traffic follows
- establish a connection
- perform all necessary operations to move the traffic along the route
- provide bandwidth and congestion management

For more information on networking systems, see the following sections:

- “Networking features” (page 21)
- “Networking architecture” (page 25)
- “Passport trunking and base routing” (page 29)
- “Routing methods” (page 34)
- “Traffic management” (page 44)

Networking features

IP, DPN gateway, and frame/cell trunk are available only on Passport 7400 series switches. Passport 15000 and 20000 do not use these services.

Consolidation of different traffic types

The Passport product is distinct in its ability to consolidate different traffic types, such as voice, video, and data, on the same network over common facilities. This integration is possible because of the combination of the Passport trunking, routing, and traffic management systems. The end result is more efficient use of network resources and more flexibility.

Scalable to suit very large networks

The Passport network is designed to be scalable from one node to thousands of nodes to provide flexibility as the network grows. Using a combination of the packaging options, the Passport system handles small and large wide-area network- (WAN) sized network requirements.

Frame and cell switching

The flexible Passport architecture supports both frame and cell switching to accommodate all traffic types. The architecture provides a smooth conversion of different traffic types (frame, cell, bit stream) to and from ATM.

Connectionless and connection-oriented routing supported

The Passport system's unique architecture supports both connectionless and connection-oriented routing for efficient transport of traffic.

Currently Passport supports the following routing systems:

- Connectionless routing systems
  - — Dynamic Packet Routing System (DPRS)
  - — IP routing on Passport 7400 series switches
- Connection-oriented routing systems
  - — Path-Oriented Routing System (PORS)
  - — ATM Routing System
- Hybrid routing systems
  - — frame relay managed cut-through switching (MCS)
  - — multiprotocol label switching (MPLS)

Routing systems to support specific access services

Some routing systems support different types of access services:

- DPRS supports frame relay UNI and NNI, and interworking between Passport and DPN-100 devices.
- PORS supports voice, HDLC transparent data, bit transparent data services, and AAL1 Circuit Emulation Service (CES).
- ATM routing supports ATM bearer service (ABS), FR/ATM interworking, and ATM multiprotocol encapsulation (MPE).

PORS and ATM also support frame relay managed cut-through switching (MCS). MCS provides dynamically-established, point-to-point virtual connections for multiplexed frame relay service traffic.

ATM also supports multiprotocol label switching (MPLS). MPLS is a label-swapping, networking technology that forwards IP traffic over multiple, underlying layer-2 media.

Fully dynamic design

Passport routing systems are designed to dynamically reroute traffic to make use of new facilities and to avoid failures and congestion. These systems are fully automatic and require no operator intervention. This dynamic design simplifies network operator involvement yet maximizes the network efficiency.

Support of applications with varying quality of service parameters

Passport routing systems use quality of service to define the criteria for path selection and for path access. The routing class of service (RCOS) attributes such as cost, delay, and throughput contribute to the decision of the best path. Under congestion situations, access to paths depends on the priority and reliability attributes which define the packet importance and emission status. See “DPRS quality of service attributes” (page 99).

Data traffic, for example, can require high throughput but can tolerate variable delays and uses connectionless routing. Voice and video traffic, on the other hand, uses small cells, requires low delays and minimal delay variances, and uses PORS. PORS is ideally suited to meet these requirements because of the low overhead, the interrupting queue and bandwidth reservation features.

ATM routing supports constant bit rate (CBR), real time and non-real time variable bit rate (rt-VBR, nrt-VBR) and unspecified bit rate (UBR) connections.

Multiple Priority System (MPS)

Each frame/cell entering a Passport node is assigned both an emission (delay sensitivity) and a discard (importance) priority. The emission priority is used during frame/cell forwarding to determine the order in which the packets will be emitted to Passport trunk facilities. The discard priority is used to decide whether the frame/cell will be discarded when congestion is encountered. MPS is also used at service egress points, such as frame relay queues.

MPS offers several key benefits: it provides the basis for flexible network engineering and quality of service (QoS) guarantees to the different applications, supports traffic policing and fairness among network users, and enables the offering of innovative QoS differentiated services by the service provider.

Dynamic bandwidth and congestion management

The Passport bandwidth and congestion management functions constitute the basis of the Passport traffic management system. Bandwidth management involves optimizing the network traffic flow over the links by allocating or releasing PORS connections based on their bandwidth requirements and the bandwidth available. For DPRS, bandwidth management includes grouping of links into link groups based on routing class of service (RCOS), and basing the routing metrics on the bandwidth not renewable by PORS. For ATM routing, bandwidth management includes the allocation of bandwidth to different traffic types and connection types.

Congestion management involves optimizing the network traffic flow when the demand for network resources exceeds capacity. Passport mechanisms such as MPS, overflow routing and load balancing contribute to congestion management. Dynamic control of traffic flow is achieved through adaptive-reactive congestion control mechanisms.

Increased service reliability with Passport 15000 and 20000

With Passport 15000 and 20000, the applications and features that provide services fall into three categories: hot standby, warm standby and cold standby.

Hot standby applications and features can run uninterrupted, even when the hardware providing that service changes. Hot standby applications and features offer hitless services during an FP or CP switchover. During an equipment switchover, hot standby applications incur minimal traffic interruption and established connections are maintained. This ability reduces service down time and increase service reliability.

Warm standby applications and features also provide increase service reliability, although to a lesser extent than hot standby applications and services. During an equipment switchover, warm standby applications and features incur a longer outage of service than hot standby applications, but not as long as cold standby applications. As well, all connections must be re-established.

Support for link groups

Link groups allow cost-effective and flexible Passport trunking options. Link groups support multiple links between nodes, increasing available bandwidth. This function is similar to inverse multiplexing without extra hardware. Link groups also allow the partitioning of different traffic over different links.

Networking architecture

In the Passport system, the networking and application software is organized into three layers. This modular design

- allows the use of different facilities (frame-cell trunks or Passport trunks over ATM) transparent to the applications
- provides a clean architecture so that enhancements or evolution to underlying layers can be introduced without impact to upper layers

- provides optimal delivery of packets that minimizes the segmentation and reassembly required

The layered architecture of a Passport network is shown in the figure "Passport networking architecture" (page 27).

IP routing, DPN gateway, and frame/cell trunk are available only on Passport 7400 series switches. Passport 15000 and 20000 do not use these services.

The layers from the bottom up are

- "Passport switching layer" (page 28)
- "Passport network layer" (page 28)
- "Passport service layer" (page 28)

Figure 1
Passport networking architecture

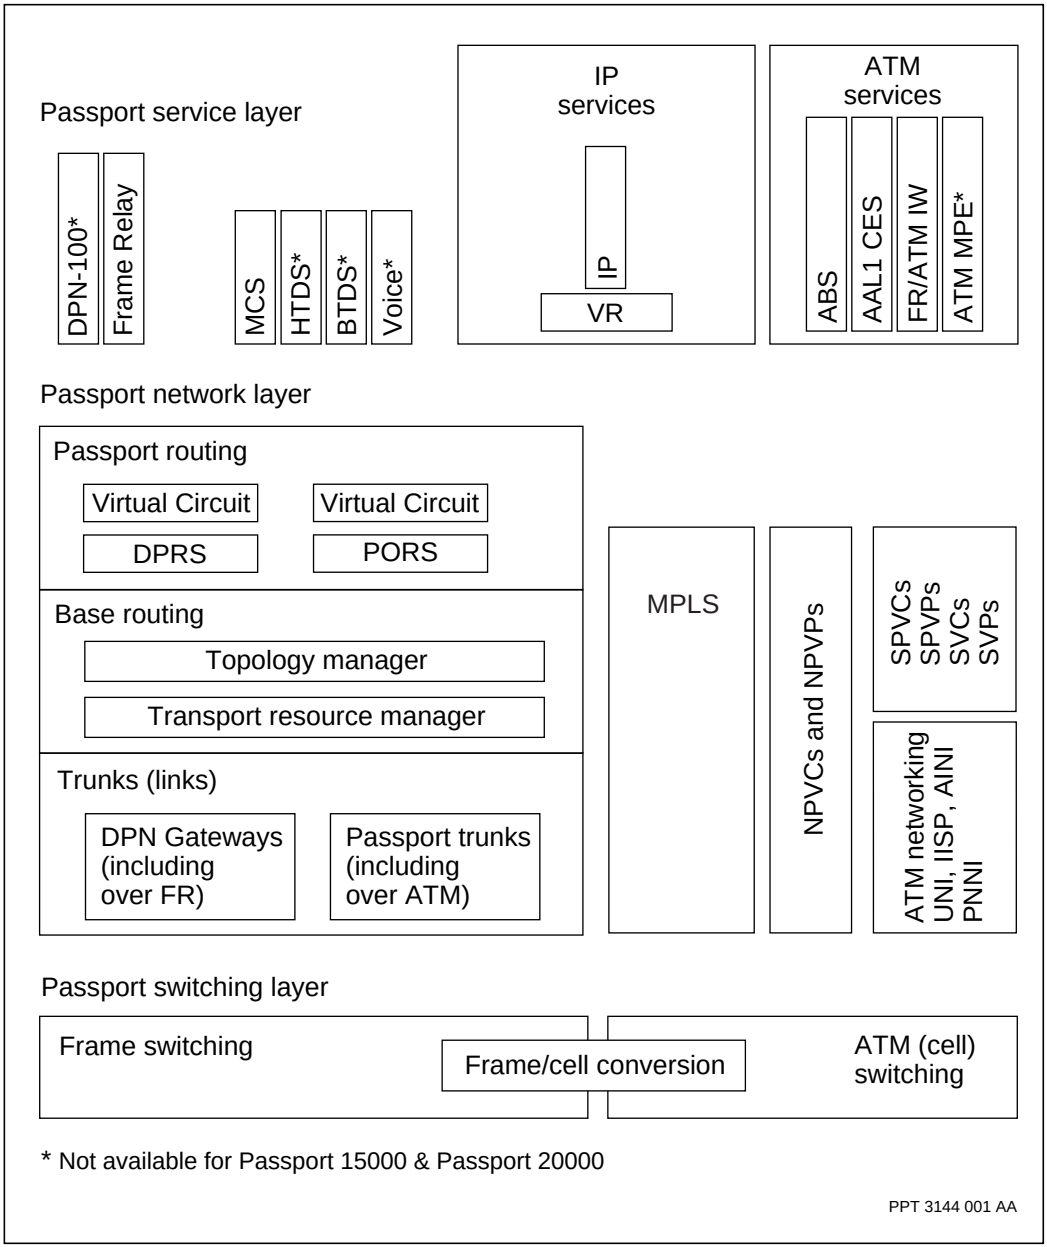

Passport switching layer

The switching layer provides Passport frame/cell and ATM links. This layer is optimized for packet forwarding at very high throughput rates. The software is tightly coupled to the Passport hardware that handles data transfer functions, and frame/cell conversion.

The switching layer involves the actions that transmit raw bits over a communication channel. Examples of physical layer protocols supported by the Passport system are V.11, V.35, DS1, E1, and OC-3.

Passport network layer

The network layer consists of Passport trunking, routing, and traffic management systems. This layer maintains a complete and up-to-date view of the network state, topology and Passport trunk availability. Base routing provides this basic routing information to a number of protocol-based routing systems. These systems determine the best route through the network, based on their specific routing algorithms, and the types of Passport trunks and attributes (such as delay or throughput) associated with the application.

The network layer monitors Passport trunk utilization; detects and reacts to any congestion; and provides a consistent approach to emission priority, discard priority, and alternate routing in the event of new facilities or failures. Bandwidth management provides allocation and sharing of bandwidth resources to all routing systems.

The network layer is divided into core networking and ATM networking. Core networking refers to the two Passport networking systems: DPRS, and PORS. ATM networking for ATM access services utilizes standards-based protocols such as Interim inter-switch signaling Protocol (IISP) and private network to network interface (PNNI).

Also in the network layer is multiprotocol label switching (MPLS). MPLS is a label-swapping, networking technology that forwards IP traffic over multiple, underlying layer-2 media.

Passport service layer

The service layer consists of the application services, some with their own routing protocols. A wide variety of applications with different traffic characteristics and network requirements can be mapped onto underlying

network capabilities and are transparent to the application. Changes in Passport trunking and topology configurations can be made with minimal interruption to the existing applications.

Passport trunking and base routing

Passport trunking and base routing software performs functions common to all routing systems except ATM and IP, which have their own software.

IP is available only on Passport 7400 series switches. Passport 15000 and 20000 do not use this service.

Passport trunking

The Passport trunking system manages the links interconnecting Passport nodes or the links interconnecting Passport nodes to DPN-100 modules. The Passport trunking system includes the Passport trunking protocols (for example UnAcked) from the switching layer and the Passport trunking elements (for example Passport trunks) from the network layer.

This Passport trunking section includes the following topics:

- “Passport trunks” (page 32)
- “DPN gateway on Passport 7400 series switches” (page 32)
- “Link groups” (page 32)
- “Point-to-point and logical links” (page 33)

See the figure “Passport trunking terminology” (page 30) for an illustration of Passport trunking terminology.

DPN gateway is available only on Passport 7400 series switches. Passport 15000 and 20000 do not use this service.

Figure 2
Passport trunking terminology

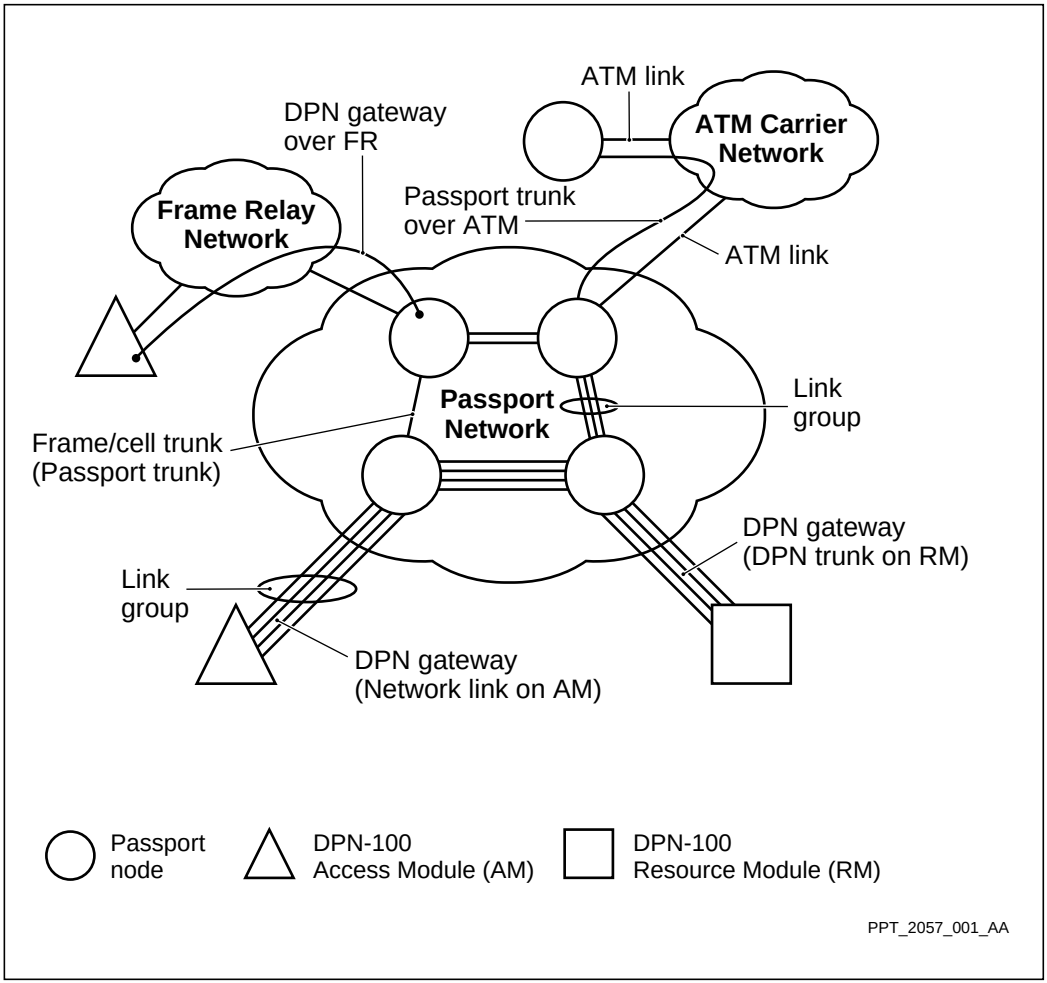

Link is the generic term for the point-to-point or logical connection between devices such as Passport nodes, DPN-100 modules, frame relay access devices, and third party devices. Specific instances of these links are called Passport trunks and gateways (see the figure “Passport links” (page 31)). On Passport nodes, a total of 1,023 links (Passport trunks for all Passport switches and DPN gateways for Passport 7400 series switches) can be supported.

Figure 3
Passport links

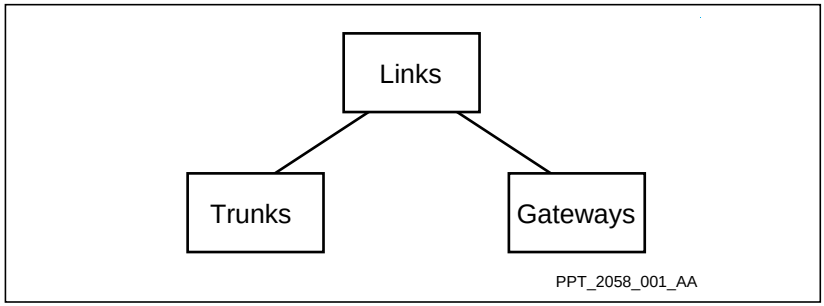

Links support many types of traffic including voice, video, transparent data, frame relay and DPN-100 services. Different traffic types require different transport mechanisms or link types (see the figure “Passport transport mechanisms” (page 31)).

DPN gateway and frame/cell (UNACK) are available only on Passport 7400 series switches. Passport 15000 and 20000 do not use these services.

Figure 4
Passport transport mechanisms

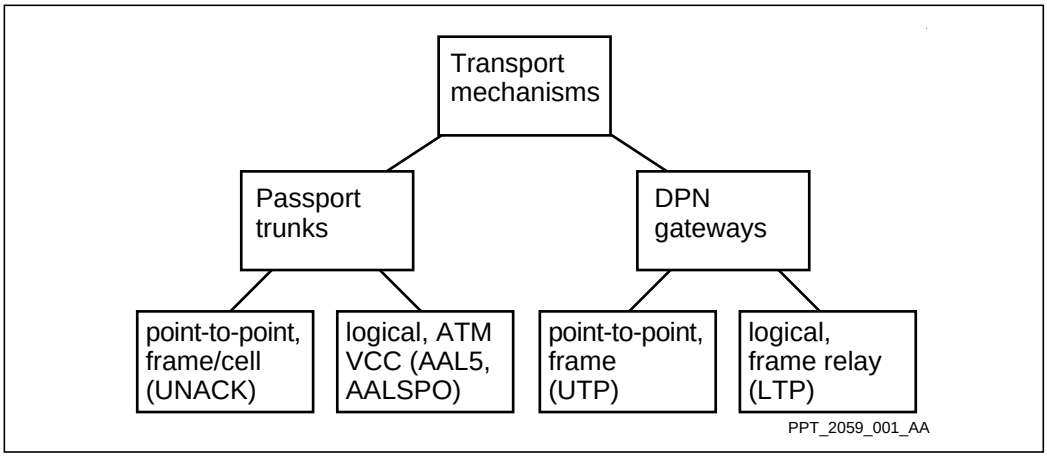

Passport trunks

Passport-to-Passport links are called Passport trunks. A Passport trunk is a point-to-point or logical (ATM virtual channel connection [VCC]) connection between two Passport nodes over which Passport proprietary routing protocols are run. They use two transport mechanisms: frame-cell trunks and Passport trunks over ATM.

- Passport frame-cell trunks transport data traffic (frame relay and DPN-100) as frames, and constant bit rate traffic (voice, video) as cells. The trunking protocol used for transmission is HDLC-based. This trunk protocol is available only on Passport 7400 series switches. Passport 15000 does not use the trunking protocol.
- The Passport trunk over ATM (formerly known as ATM logical trunks) is carried on one or more ATM VCCs and uses standard ATM adaptation layer protocols based on AAL5.

DPN gateway on Passport 7400 series switches

Links to other networks are called gateways. The following terminology exists for Passport to DPN gateways:

- A connection from a Passport node to a DPN-100 module is called a DPN gateway and uses the acknowledged universal trunk protocol (UTP).
- DPN gateway over frame relay uses an unacknowledged version of UTP called light-weight trunk protocol (LTP).
- A DPN gateway can be linked to an Access Module (AM), Resource Module (RM), or an RM configured as a Call Server (CSRM). A gateway connection at the AM side is called a network link. A gateway connection at the RM or CSRM side is called a DPN trunk.

Link groups

Passport trunks can be configured into link groups. A link group represents all the links between two devices. Each Passport node supports a maximum of four links in a link group. A maximum of 255 link groups per Passport node includes link groups between Passport nodes, and link groups to DPN-100 RMs and AMs.

Point-to-point and logical links

The figure “Point-to-point or logical links” (page 33) illustrates point-to-point and logical links.

Figure 5
Point-to-point or logical links

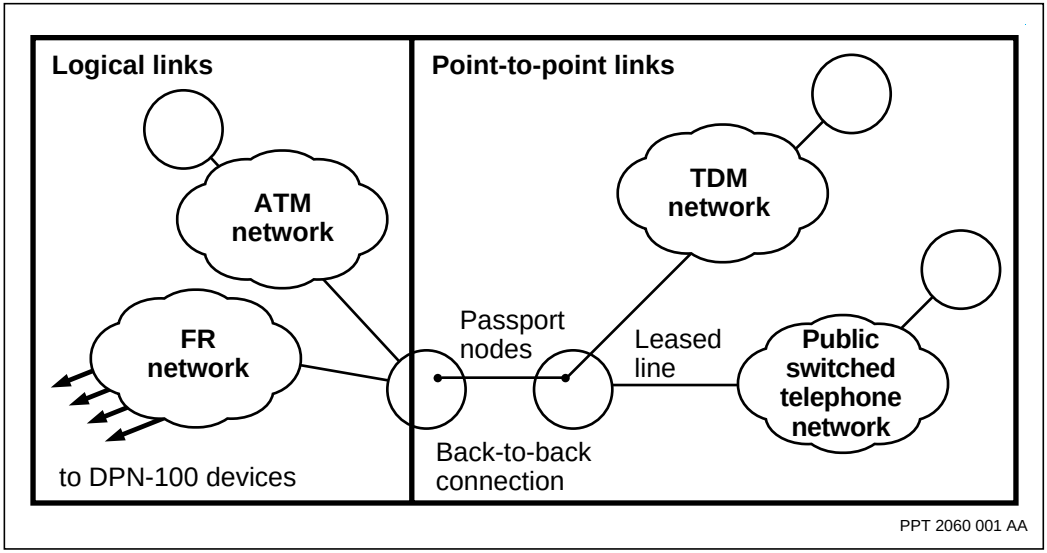

Point-to-point connections are those where the link is dedicated to a Passport trunk or DPN gateway (has a sole user of the end-to-end connection). The network and switching layers of the Passport architecture are managed by Passport trunking and base routing. For point-to-point links, Passport trunks and DPN gateways transfer information over a time-division-multiplexed (TDM) network, leased lines, and back-to-back connections.

Logical connections are those where the link is shared with other users/ applications. The network layer of the Passport architecture is managed by Passport trunking and base routing while the switching layer is managed by other underlying protocols. For logical links, Passport trunks and DPN gateways transfer information over public ATM or frame relay networks.

Base routing

The Passport base routing system, part of the network layer, provides the basic function of transport resource and topology management through

- the transport resource manager (TRM), which maintains a local view of the current status and bandwidth available on the links
- the topology manager, which is responsible for obtaining and distributing the network-wide view of the point-to-point and logical links between all Passport nodes in the network

The base routing system captures and maintains a view of the Passport network topology. It calculates optimal paths based on certain criteria (for example delay, throughput) between all Passport nodes in the network.

The Passport trunking and base routing systems are used by the DPRS and PORS routing systems.

For more information on Passport trunking and base routing, see “Passport trunking and base routing systems” (page 49).

Routing methods

Routing is the process of selecting a route along which to direct traffic from point A to point B.

A possible route through a simple network is shown in the figure “A possible route through a network” (page 35).

Figure 6
A possible route through a network

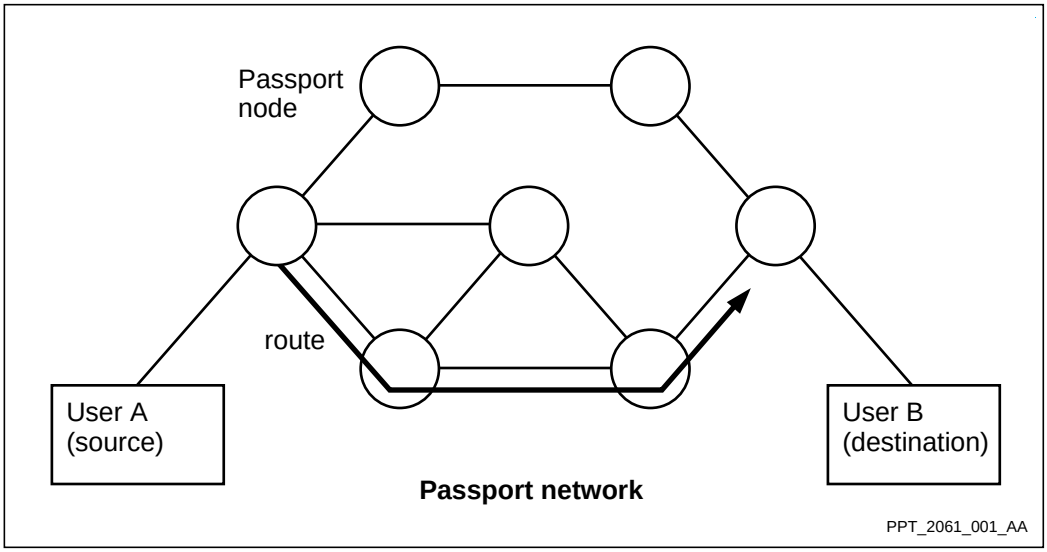

The routing system monitors network connectivity to support routing of data between Passport nodes. The routing system has a network-wide scope, implemented by software distributed on each node of the network. The routing systems on all nodes in the network communicate with one another to propagate routing information.

The routing layer supports the forwarding functions of the different routing systems by providing mechanisms for the following:

- discover the network topology and propagate this information across the network
- on each node, distribute routing information
- create and maintain packet forwarding tables for DPRS and PORS on Passport nodes, and for IP on Passport 7400 series nodes
- forward application data from node to node along the selected routes

The routing system also performs bandwidth management and congestion management functions.

The network layer supports the two basic network routing methods: connectionless and connection-oriented routing.

The routing methods depend on the base routing system to perform the following tasks:

- use topology information to make optimal route decisions between Passports
- efficiently and quickly distribute base routing control information
- rely on underlying transport resource manager (TRM) and Passport trunking systems to monitor and manage resources. This process removes the burden from the routing protocols. For example, introducing Passport cell trunking into the network backbone requires no modification to the routing protocols.

  Note: ATM routing and IP have their own systems to perform base routing and Passport trunking functions.

The following routing methods are described:

- “Connectionless routing” (page 36)
- “Connection-oriented routing” (page 39)
- “Hybrid routing systems” (page 43)

Connectionless routing

Connectionless routing dynamically determines the path of the data through the network. Connectionless routing has the following characteristics:

- A forwarding decision for each individual packet is made at each node (or hop) in the network en route to its destination. The path is based on: the network topology at data transmission time; Passport trunking characteristics (for example, throughput and delay); and the destination node address carried in the packet header. Forwarding tables on each node map the destination address to an outgoing link.

- The packets that make up a transmission or data flow may not always follow the same route. Under normal, stable network conditions, all packets follow the same route and arrive in the correct order. Under congested situations or when topology changes occur, packets may be rerouted instantaneously.
- Traffic load can be spread throughout the network to allow the available bandwidth between two nodes to exceed that of any single route between the nodes.
- If disordering occurs because of rerouting or load spreading, packets are reordered by the virtual circuit, by an access service, or from outside the network.
- Delivery of packets is not guaranteed because lost or discarded packets may not be retransmitted by the service.

The figure “Connectionless routing characteristics” (page 38) illustrates some of the characteristics of connectionless routing.

Figure 7
Connectionless routing characteristics

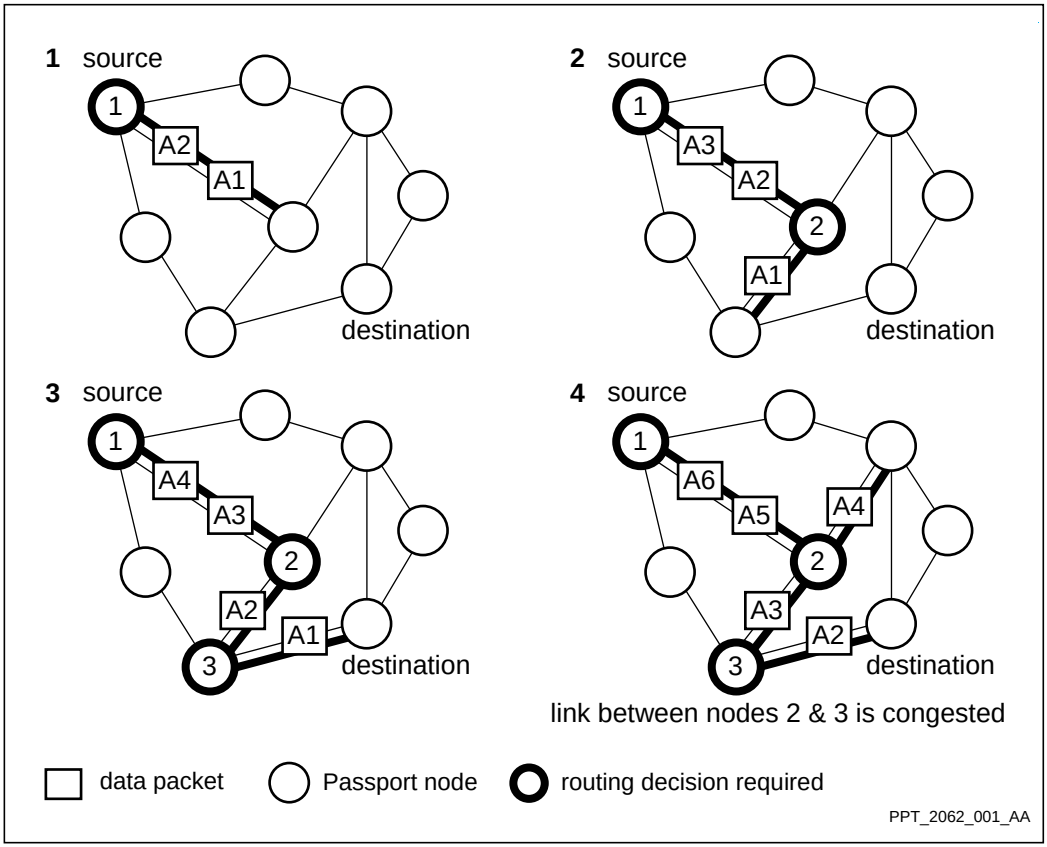

Connectionless routing is ideal for interactive applications such as frame relay applications. The connectionless routing system finds the best route through the network and relieves the applications and the network operator of the burden of handling a changing topology. It reacts quickly to topology changes with negligible impact on services.

Passport connectionless routing consists of two routing systems:

- “Dynamic Packet Routing System (DPRS)” (page 39)

Dynamic Packet Routing System (DPRS)

DPRS provides Passport with the ability to support the Passport frame relay services. DPRS also supports transport of DPN-100 traffic across a Passport backbone network. DPRS implements a hierarchical addressing routing protocol. DPRS uses topology information to find the best routes to every destination address. These routes are placed in the DPRS forwarding tables for use by the packet forwarding function.

Spanning all Passport nodes in the network, DPRS routing also supports Passport interaction with Preside Multiservice Data Manager for network management activities.

For more information on DPRS, see “Dynamic Packet Routing System (DPRS)” (page 85).

IP services on Passport

Passport also supports the native routing of standard IP traffic using connectionless routing.

For more information on IP services available on Passport, see 241-5701-805 *Passport 7400, 15000, 20000 Understanding IP*.

Connection-oriented routing

Connection-oriented routing selects a single path through the network when the connection is set up. Data transfer and path tear down phases follow the call setup phase. Connection-oriented routing has the following characteristics:

- Routing decisions are made at call setup time, based on the network topology at call setup time and routing information in the call setup request packet. The routing information includes the destination address and quality of service (such as cost and delay for PORS; CBR, rt-VBR, nrt-VBR, and UBR for ATM) attributes.
- Bandwidth reservation and path establishment occur next. Bandwidth for a connection may be reserved at each node along the chosen path. Load balancing on links may take place at connection admission time. Routing information is required on each link in the path instead of in each packet.

An indicator or label (LCN for PORS, VCI or VPI for ATM) is associated with each data packet allowing the packets to be switched at tandem nodes rather than being routed.

- Once the path is established, data transfer occurs. The routing decisions at each node on the way to the destination are minimal because of the call setup effort.
- All packets of the same connection are sent along the same path during data transfer and arrive at the destination in the same order as they were sent. For PORS, forwarding tables on each node map the path identification marker to an outgoing link and LCN, and contain the class of service (COS) parameters. For ATM, hardware autonomously forwards the cells.
- Under congested situations, packets are discarded until congestion ceases.
- Upon the failure of a component of a path, all connections through that component must be reestablished over other paths. When the failed component comes back into service, new connections can be established over it. For PORS, in time, some or all of the original connections may also return to using the component.
- Delivery of packets is not guaranteed by the routing system. Lost or discarded packets may be retransmitted by the service.
- The selected path is torn down when the call is taken down.

The figure “Connection-oriented routing characteristics” (page 41) illustrates connection-oriented routing characteristics.

Figure 8
Connection-oriented routing characteristics

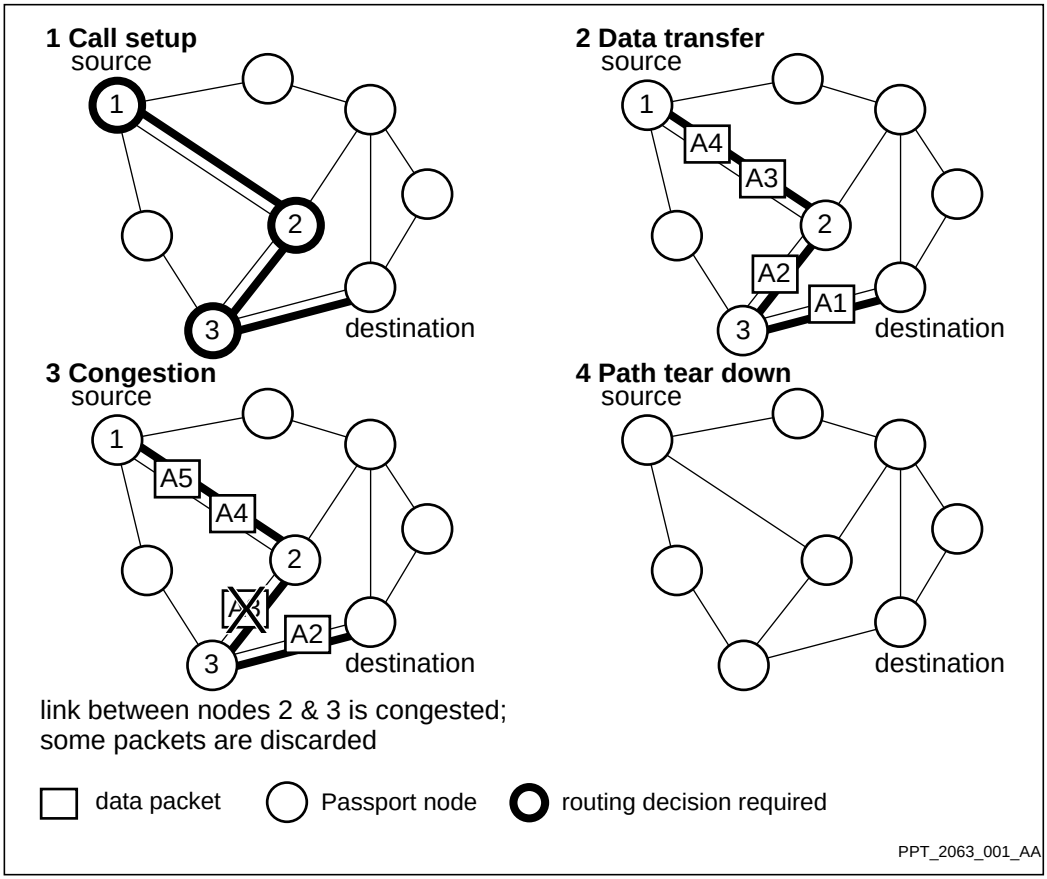

Connection-oriented routing is ideal for delay-sensitive applications such as voice and video applications. The connection-oriented routing system finds the best route through the network and relieves the applications and the network operator of the burden of handling a changing topology. It reacts quickly to topology changes with negligible impact on services.

Passport connection-oriented routing consists of the following:

- "Path-Oriented Routing System (PORS)" (page 42)
- "ATM Routing System" (page 42)

Path-Oriented Routing System (PORS)

PORS provides a connection-oriented routing system that automatically establishes and maintains a connection (reroutes, optimizes, load balances) during the life of the connection. PORS supports switched and permanent connections. PORS provides Passport with the ability to transport traffic that is sensitive to delay variance such as video, voice, and transparent traffic across Passport networks. PORS uses hop-by-hop forwarding based on the path identifier in the packet. At call setup time, a path is established, bandwidth is reserved, then all packets flow along the same path to the destination. The path is torn down after the connection is no longer needed.

For more information on PORS, see "Path-Oriented Routing System (PORS)" (page 119).

ATM Routing System

ATM routing is a connection-oriented system. ATM networking provides dynamic run-time connection setup between Passport nodes, and allows Passport switches to interwork with other ATM switches. ATM networking provides the addressing, signaling, and routing facilities needed to support the following:

- switched virtual connections (SVC)
- switched virtual paths (SVP)
- soft permanent virtual connections (S-PVC)
- soft permanent virtual paths (S-PVP)

Nailed-up PVCs (N-PVC) and nailed-up permanent virtual paths (N-PVP) are also supported. These networking capabilities enable users to set up ATM connections in real time.

For more information on ATM routing, see "ATM routing system" (page 141).

Hybrid routing systems

A hybrid routing system is a system that combines the characteristics of connectionless and connection-oriented routing. Networking includes the following hybrid systems:

- “Frame relay managed cut-through switching (MCS)” (page 43)
- “Multiprotocol label switching (MPLS)” (page 43)

Frame relay managed cut-through switching (MCS)

MCS is a special routing facility. It uses connection-oriented routing systems to transfer Passport services that need aggregate and high-performance switching, such as frame relay. The MCS connection is a switched path that provides many-to-one multiplexing of connections for a data service. Service connections that share the same QoS parameters can be mapped onto a single switched path and transported through the network. MCS takes advantage of the underlying network infrastructure, using either the PORS or ATM routing system for traffic transmission.

For more information on MCS, see 241-7401-440 *Passport 7400 Frame Relay Managed Cut-through Switching Guide*.

Multiprotocol label switching (MPLS)

MPLS is a label-swapping, networking technology that forwards IP traffic over multiple, underlying layer-2 media. Essentially, MPLS forwards a packet by swapping labels at each node in its path. In Passport networks, MPLS transports IP traffic over ATM infrastructure, allowing carriers and large enterprises to send IP data easily across the ATM backbone.

MPLS is an emerging standard for network-layer packet forwarding, based on a number of signaling protocols proposed by the Internet Engineering Task Force (IETF). Among these protocols are the label distribution protocol (LDP) and constraint-based routing using LDP (CR-LDP).

For more information on MPLS, see “Multiprotocol label switching” (page 109).

Traffic management

This section provides a high-level overview of main concepts relating to traffic management on a Passport network For detailed information about the traffic management techniques described here, see the following documents:

- 241-5701-705 *Passport 7400, 15000, 20000 ATM Traffic Management Fundamentals*
- 241-5701-706 *Passport 7400, 15000, 20000 ATM Traffic Shaping and Policing*
- 241-5701-707 *Passport 7400, 15000, 20000 ATM Queuing and Scheduling*
- 241-5701-708 *Passport 7400, 15000, 20000 ATM CAC and Bandwidth Management*
- 241-5701-901 *Passport 7400, 15000, 20000 Frame Relay Fundamentals*
- 241-5701-902 *Passport 7400, 15000, 20000 Configuring Frame Relay*
- 241-5701-920 *Passport 7400, 15000, 20000 Frame Relay to ATM Interworking Guide*
- 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide*

Traffic management incorporates a group of control mechanisms and engineering practices to achieve optimal traffic flow across the network.

Traffic management functions exist in all layers of the Passport networking architecture.

Role of traffic management

The traffic management system maintains the network's performance objectives by ensuring the efficient use of network resources in the transfer of customer traffic, and by managing congestion.

When traffic flows across a network, the primary goal of the traffic management system is to maintain certain levels of performance and prevent information loss. The system must also ensure that network resources are

optimized, and handle any saturated or congested situations that occur. To achieve these goals, the components of the traffic management system monitor connection requests and traffic flow across the network.

These traffic management mechanisms can be characterized in terms of preventive and reactive control. Preventive control includes the mechanisms used to avoid contract violation and congestion conditions. Reactive control mechanisms are applied when a network resource becomes saturated or congested.

Traffic management mechanisms

The traffic management system consists of many mechanisms that work together to provide an end-to-end traffic management architecture. As shown in the figure "Traffic management areas" (page 46), the traffic management mechanisms operate in three main areas of the Passport system: service, routing system, and platform.

The individual services handle the mechanisms that control the access of traffic to the network. This area includes call characterization and traffic policing mechanisms.

The routing systems provide the traffic mechanisms that control the flow of traffic after it has been admitted to the network. This area includes connection admission control, bandwidth management, and route selection mechanisms.

The platform area includes the control mechanisms that apply to all Passport services and routing systems to ensure the appropriate quality of service for the various types of traffic. These mechanisms work at the frame or cell level to implement the Passport queue scheduling and traffic discarding policies.

Figure 9
Traffic management areas

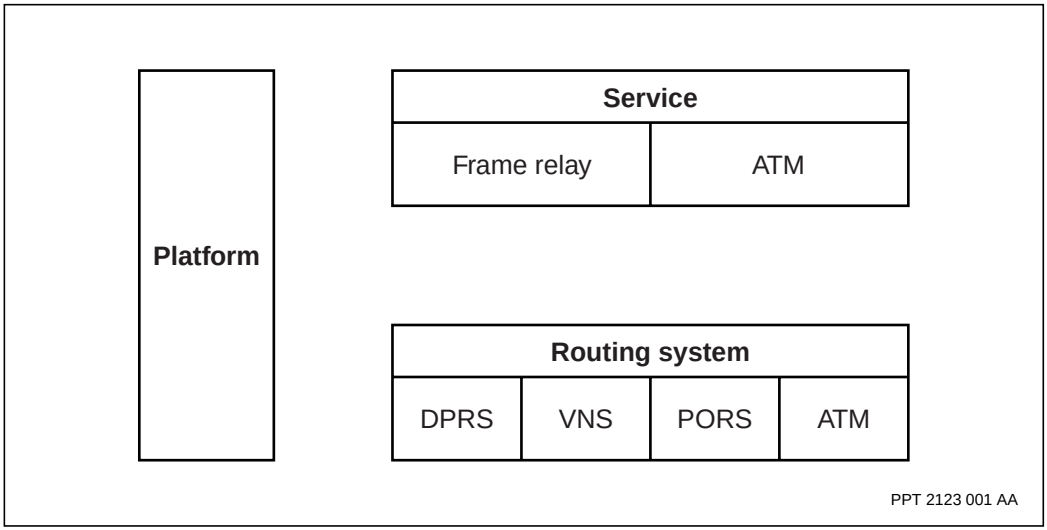

Service mechanisms

The traffic management mechanisms that operate at the level of individual services apply primarily to

- defining the user's service requirements
- ensuring that the system meets these requirements

The table "Service-level traffic management features" (page 46) shows the four main service-level mechanisms and the features through which the mechanisms are implemented for each service.

Table 1
Service-level traffic management features

Mechanism	ATM feature	Frame relay feature
Call characterization	Traffic contract	Negotiated traffic parameters
		Quality of service
		Class of service

(Sheet 1 of 2)

Table 1 (continued)
Service-level traffic management features

Mechanism	ATM feature	Frame relay feature
Traffic policing	Usage parameter control (UPC)	Rate enforcement
Traffic shaping	Traffic shaping	
Adaptive-reactive response		Rate adaptation
(Sheet 2 of 2)		

Routing mechanisms

After traffic has been admitted to the network, the routing mechanisms of the traffic management system control the traffic flow. The routing systems search for paths that conform to the connection requirements, and make sure that calls are spread across the best available routes. The routing systems also provide notification of any congestion encountered across the network. The table "Routing-level traffic management features" (page 47) shows the routing-level mechanisms and the features through which the mechanisms are implemented for each routing system.

Table 2
Routing-level traffic management features

Mechanism	DPRS	PORS	ATM
Connection admission control (CAC)		CAC	CAC
(Sheet 1 of 2)			

Table 2 (continued)
Routing-level traffic management features

Mechanism	DPRS	PORS	ATM
Route selection	Load spreading	Bandwidth reservation	Bandwidth allocation
	Load sharing	Path optimization and bumping	
	Overflow routing		
Adaptive-reactive control	Congestion notification (FCI/BCI), (FECN/BECN for frame relay traffic)	Congestion notification (FCI/BCI)	Congestion notification (EFCI)

(Sheet 2 of 2)

Platform mechanisms

Traffic management at the platform level concentrates primarily on monitoring and controlling the node's resources. The platform-level mechanisms tend to be reactive and instantaneous in character, and operate at the frame or cell level, often responding to transient congestion conditions.

The main mechanism in this category implements the Passport queue scheduling and discarding policies. It is known as the Multiple Priority System (MPS). The MPS applies to all Passport services.

Chapter 2
Passport trunking and base routing systems

This section describes Passport trunking and base routing software systems used by DPRS and PORS. The base routing system is the underlying layer of the routing system which manages the network topology and local links.

Although DPRS and PORS can use Passport trunks over ATM, ATM routing has its own systems to perform base routing and Passport trunking functions. IP also has its own systems.

The topics covered are as follows:

- "Passport trunking" (page 49)
- "Transport resource manager" (page 58)
- "Topology manager" (page 64)

For more detailed information on Passport trunking, see "Passport trunking references" (page 18).

This chapter also includes an overview of the "Topology regions" (page 68) feature and the "Tandem suppression" (page 82) feature.

Passport trunking

This section introduces Passport trunking. It includes the following topics:

- "Role of Passport trunking" (page 50)
- "Passport trunking features" (page 50)
- "Passport trunking architecture" (page 51)

- “Passport trunking mechanisms” (page 52)
- “Passport trunking components” (page 58)

Role of Passport trunking

Trunks physically interconnect the nodes of the network and transport information across the network. The Passport trunking system manages the frames and cells transmitted and received on the links. The Passport trunking system also determines link type, priority, and speed, and manages the Passport trunking protocol used for transport.

Passport trunking features

Passport trunking system has the following features:

Transport mechanisms supported on Passport 7400 series switches

Passport 7400 links support two transport mechanisms: Passport frame-cell trunks, and Passport trunks over ATM. Frame-cell trunks are used for transporting a mixture of both frame and cell traffic on a single connection. Passport trunks over ATM are carried on an ATM VCC.

Transport mechanisms supported on a Passport 15000 and 20000 series switch

Passport 15000 and 20000 series links support the Passport trunks over ATM transport mechanism. Passport trunks over ATM are carried on an ATM VCC.

Supports wide range of interfaces and associated link speeds

Passport 7400 links support physical interfaces from V.35 up to OC-3c with link speeds from 9.6 kbit/s up to 155 Mbit/s rates.

Provides attributes for base routing

The Passport trunking system gathers link information concerning the connection quality. Continual monitoring of link information such as availability, speed, and remote node identification permits instantaneous updates to base routing.

Multiple emission priorities

Three emission priorities manage the multiplexing of data onto the link. For lower speed links, these priorities enable delay-sensitive data, for example voice, to interrupt frames and be transmitted immediately, avoiding being delayed by large data frames.

Passport trunking system transparency

The underlying Passport trunking facilities at the switching layer are totally transparent to the routing systems.

Passport trunking architecture

The Passport trunking system spans the network and switching layers. See figure "Networking architecture" (page 25) to see where Passport trunking fits into the Passport networking architecture.

IP, DPN gateway, and frame/cell trunk are available only on Passport 7400 series switches. Passport 15000 and 20000 do not use these services.

Passport trunking and the Passport switching layer

The switching layer receives frames from and transmits frames to the network layer across point-to-point links and is independent of the transmission facility used.

Passport trunking and the Passport network layer

The Passport network layer is divided into three sublayers:

- Supported at the routing layer are the fundamental data communication routing systems: Passport DPRS and PORS. The Passport trunking system is transparent to the routing system.
- The base routing layer maintains the network topology and computes the best routes through the Passport network.
- The Passport trunk or gateway layer stages point-to-point links and provides a well-defined service interface independent of the Passport switching layer. This layer also reports to the routing system link characteristics such as availability, speed, and remote node identification. For PORS, it also manages all paths set up on the link.

Passport trunking and the Passport service layer

The service layer includes all services that use Passport trunks or DPN gateways. The use of ATM as a switching layer is totally transparent to the service layers.

Passport trunking mechanisms

The Passport trunking mechanisms section includes information on the transport mechanisms, types of trunks, and metrics.

Passport frame and cell switching

A Passport node can perform both cell and frame forwarding. Frame and cell forwarding takes place between nodes on Passport frame-cell trunks. ATM switching between ATM interfaces is supported. The Passport switch can also perform adaptation from frames to ATM cells and from Passport cells to ATM cells (see the figure “Passport is a frame/cell switch” (page 53)).

Figure 10
Passport is a frame/cell switch

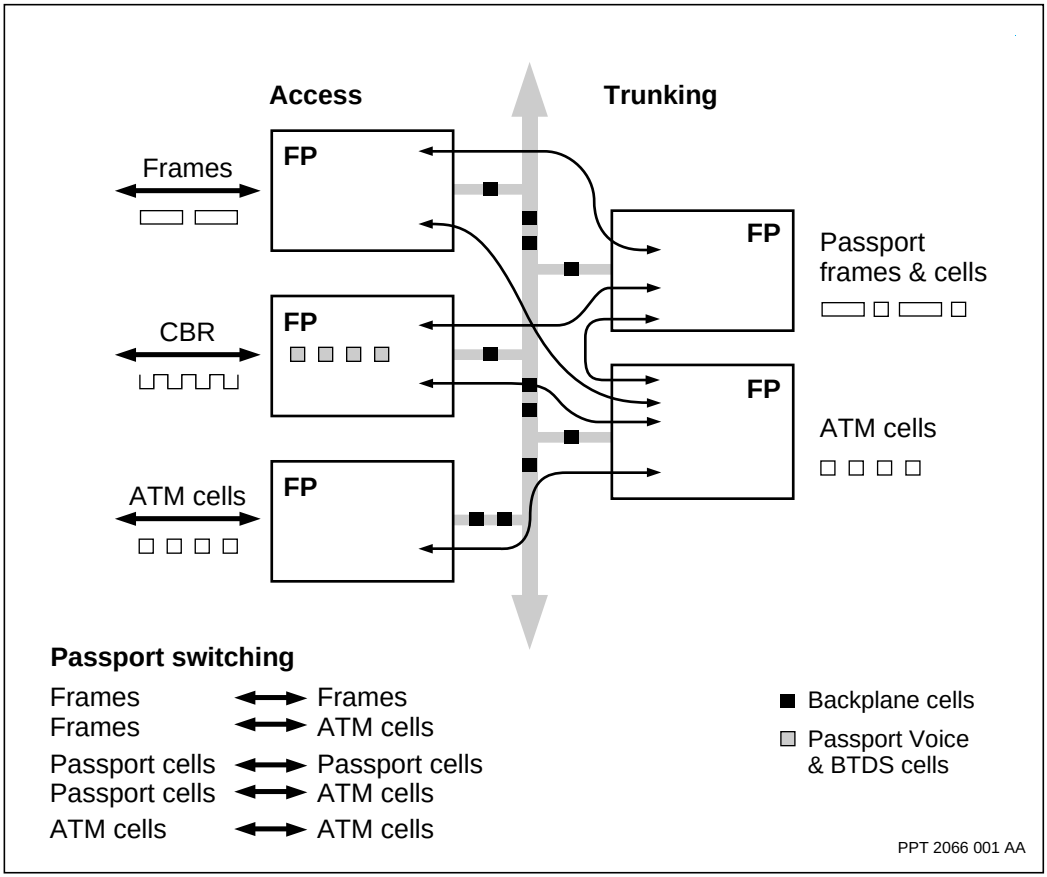

Frame-cell trunks for Passport 7400 series switches

Frame-cell trunks for Passport 7400 series switches use HDLC-based connections. This HDLC-based protocol adopts the framing and error detection scheme of HDLC but not the error recovery scheme.

Passport uses a proprietary frame-cell protocol over frame-cell trunks. The frame-cell protocol provides high performance due to minimum overhead on the link and less processor intervention to perform link layer processing on the transmit path. The major characteristics of the frame-cell protocol are the following:

- HDLC-based protocol for delimiting frames and detecting errors
- classes of service are supported by three hardware based emission priorities
- header consists only of delimiter and error detection bytes

Because of the limited header

- no acknowledgments are provided
- frames are not duplicated
- frame sequence is preserved

There are two types of frame-cell trunks: HDLC and interrupting. Frame-cell HDLC trunks are used when only frame-based services need to be carried on the Passport trunk.

Frame-cell interrupting trunks are used when constant bit rate (CBR) services (voice, BTDS) or multimedia traffic must be carried on the Passport trunk. This frame-interrupting feature of the UnAcked protocol minimizes network delay and delay variance for voice, transparent data, and multimedia traffic by carrying this traffic as cells at the highest emission priority level. See “Emission priorities” (page 57).

Passport trunks over ATM

Passport trunks over ATM are used to transport Passport services, for example frame relay and PORS traffic, and DPN-100 data transparently over ATM. The ATM protocol consists of virtual channel connection (VCC) access through the standard ATM adaptation layer (AAL5). The ATM layer segments frames into cells and transmits them across the VCC. Each Passport trunk is carried using one or more ATM VCCs. PORS cell efficiency can use a single VCC for each voice connection in map mode. In mux mode, a single ATM VCC is used for connectionless traffic while one additional VCC is used for voice.

Passport-to-Passport trunking protocol stacks

The frame-cell trunk (HDLC-based) and Passport trunk over ATM protocol stacks are similar. They differ in that the frame-cell and physical layer of the HDLC-based stack is replaced in the ATM stack with AAL5, ATM, and the physical layer of the Passport trunk over ATM stack. All the layers above the switching layer are the same. See the figure "Comparison of the protocol stacks for frame-cell trunks and Passport trunks over ATM" (page 55).

Figure 11
Comparison of the protocol stacks for frame-cell trunks and Passport trunks over ATM

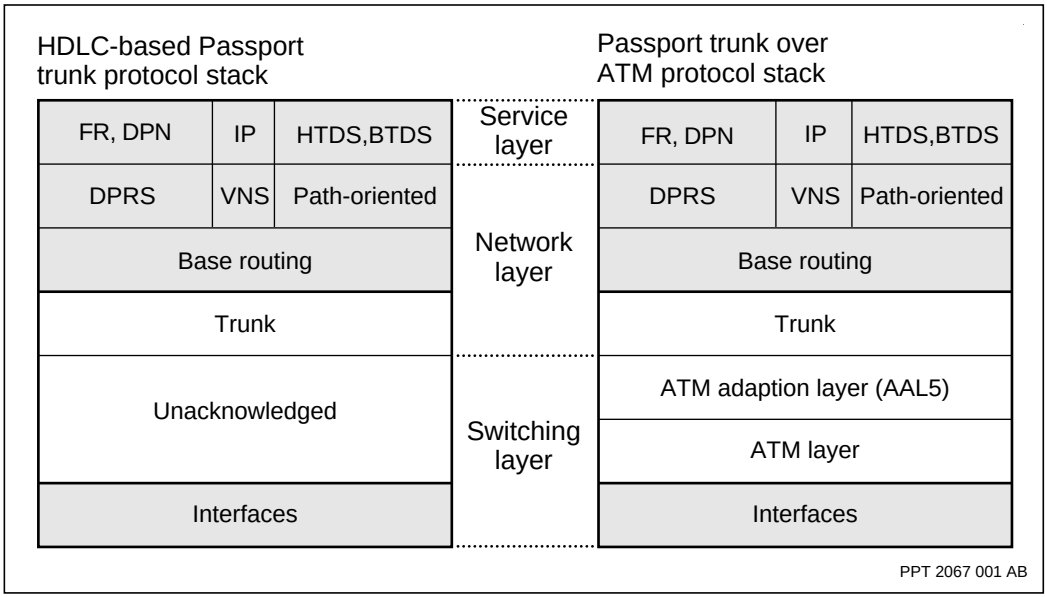

Passport-to-Passport trunking software

The Passport trunk software systems include

- frame-cell trunks: supported on HDLC type FPs and can run from 9.6 kbit/s to DS3/E3 rates of 45 Mbit/s and 34 Mbit/s respectively
- Passport trunks over ATM: supported on ATM types of FPs and can run from 1.5 Mbit/s to OC3 rates of 155 Mbit/s

Passport-to-DPN trunking

Passport-to-DPN trunking (DPN Gateway) is achieved through two DPN trunking protocols called universal trunk protocol (UTP), and light-weight trunk protocol (LTP) over frame relay

Passport-to-DPN trunking is not available on Passport 15000 and 20000 series switches.

Passport-to-DPN trunking protocol stacks

The protocol stacks for DPN Gateway are shown in the figure “Comparison of DPN gateway over UTP and DPN gateway over frame relay protocol stacks” (page 56).

Figure 12
Comparison of DPN gateway over UTP and DPN gateway over frame relay protocol stacks

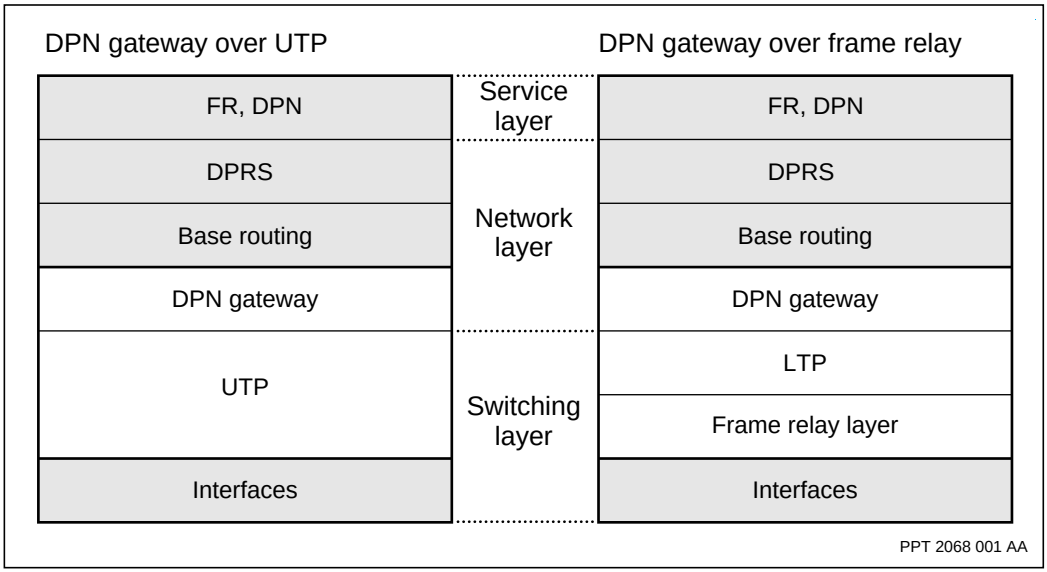

Passport-to-DPN trunking software

The DPN Gateway software systems include

- DPN gateways over UTP: supported on HDLC type FPs running from 9.6 kbit/s to DS1/E1 rates of 1.544 Mbit/s and 2.048 Mbit/s respectively
- DPN gateways over frame relay: supported on HDLC type FPs running from 9.6 kbit/s to DS1/E1 rates

Metrics

Metrics, numbers assigned to links or link groups, are used by the routing system to determine the best route through a Passport network. The best route through the network is the one for which the sum of all the link or link group metrics between the source and destination is a minimum.

The Passport trunking system collects the information required to calculate the metrics such as the bandwidth available, based on link speed or connection size of the transport mechanism being used, round trip delay and remote node identifiers. The base routing system calculates the metrics associated between two nodes.

A link metric is one number for an individual link while a link group metric represents aggregate throughput and delay of multiple links.

Link metrics

PORS uses link metrics when determining the best path between the source and destination endpoints of a connection. A PORS connection can be provisioned to minimize the following:

- cost: a provisioned value for each Passport trunk that carries PORS traffic
- delay: the measured round-trip delay for a 128-byte packet in milliseconds calculated when a link stages

Link group metrics

DPRS uses a link group metric for both the throughput and delay classes of service to determine the best paths from any node to any destination node.

- throughput: a calculated value based on the reported speed of all the links in the link group. The reported speed of the link is the measured speed, or the override speed if it is provisioned.
- delay: a calculated value based on the average delay of the preferred links in the link group or the override delay if it is provisioned.

Emission priorities

Up to three emission priorities are provided to manage the multiplexing of data onto the Passport trunk. This system allows Passport to support both variable bit rate traffic and constant bit rate traffic over the same unchannelized Passport network trunk. The three queues are designated as

normal-priority (data) queue, high-priority (data) queue, and an interrupting queue. Short, high-priority, delay-sensitive cells and frames are directed to the interrupting queue, including voice, bit transparent data cells, and multimedia traffic class (MMTC) frame relay frames.

The queues are serviced in a specific order. If the interrupting queue is empty, data in the high priority queue is transmitted first. When the high-priority queue has been emptied, the low-priority queue is serviced. When a packet arrives in the interrupting queue, the packet is immediately emitted, and may interrupt a packet emitted from the high- or low-priority queues. After the interrupting packet has been transmitted, the remainder of the interrupted frame continues.

Passport trunking components

The *Trunk* component provides a Passport-to-Passport connection. It controls a protocol stack associated with a logical port on a functional processor that provides an interconnection to another Passport with a *Trunk* component. The *Trunk* component requires the existence of the *Routing* component.

Either an *AtmAccess* or an *UnAcked* subcomponent must be added to provide the underlying link protocol and associated hardware that the *Trunk* will use.

For more information on components, see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Transport resource manager

This section introduces the transport resource manager (TRM) subsystem of the base routing system. It includes the following information:

Role of TRM

The transport resource manager is responsible for management and monitoring of link resources (Passport trunks and DPN gateways) on the node. TRM assists in handling link up and link down requests and is the intermediary between the Passport trunking and routing systems. It also manages the grouping of links into link groups.

TRM features

TRM has the following features:

Common point of access to node's view of connectivity

TRM provides a single location that gives a view of all the neighbor nodes, links, and link groups. TRM also makes link characteristics visible, collects statistics for links, and learns which LNNs are supported on a link.

TRM supports the link group concept

TRM builds and maintains the link group tables. Link groups allow cost effective and flexible Passport trunking options. Link groups enable multiple links, and therefore increased available bandwidth, between nodes. This function is similar to inverse muxing without extra hardware.

Supports topology regions

TRM provides the base support for topology regions. Topology regions are used to support very large networks by restricting the topology information that is shared within the network. TRM detects when neighbors are in different regions (region boundaries).

TRM architecture

The TRM system is an element of the base routing system, along with the topology manager system, as highlighted in figure "Passport networking architecture" (page 27). TRM resides in the network layer.

IP, DPN gateway, and frame/cell trunk are available only on Passport 7400 series switches. Passport 15000 and 20000 do not use these services.

TRM and the network layer

TRM interacts with other components in the network layer and in the switching layer. TRM resides on the control processor of each node and handles information as follows:

- receives information from the link elements of the Passport trunking system located on the function processors
- sends information to the topology manager and DPRS component on the control processor

The figure “Passport base routing system - transport resource manager and topology manager” (page 61) highlights the interaction of base routing components.

Figure 13
Passport base routing system - transport resource manager and topology manager

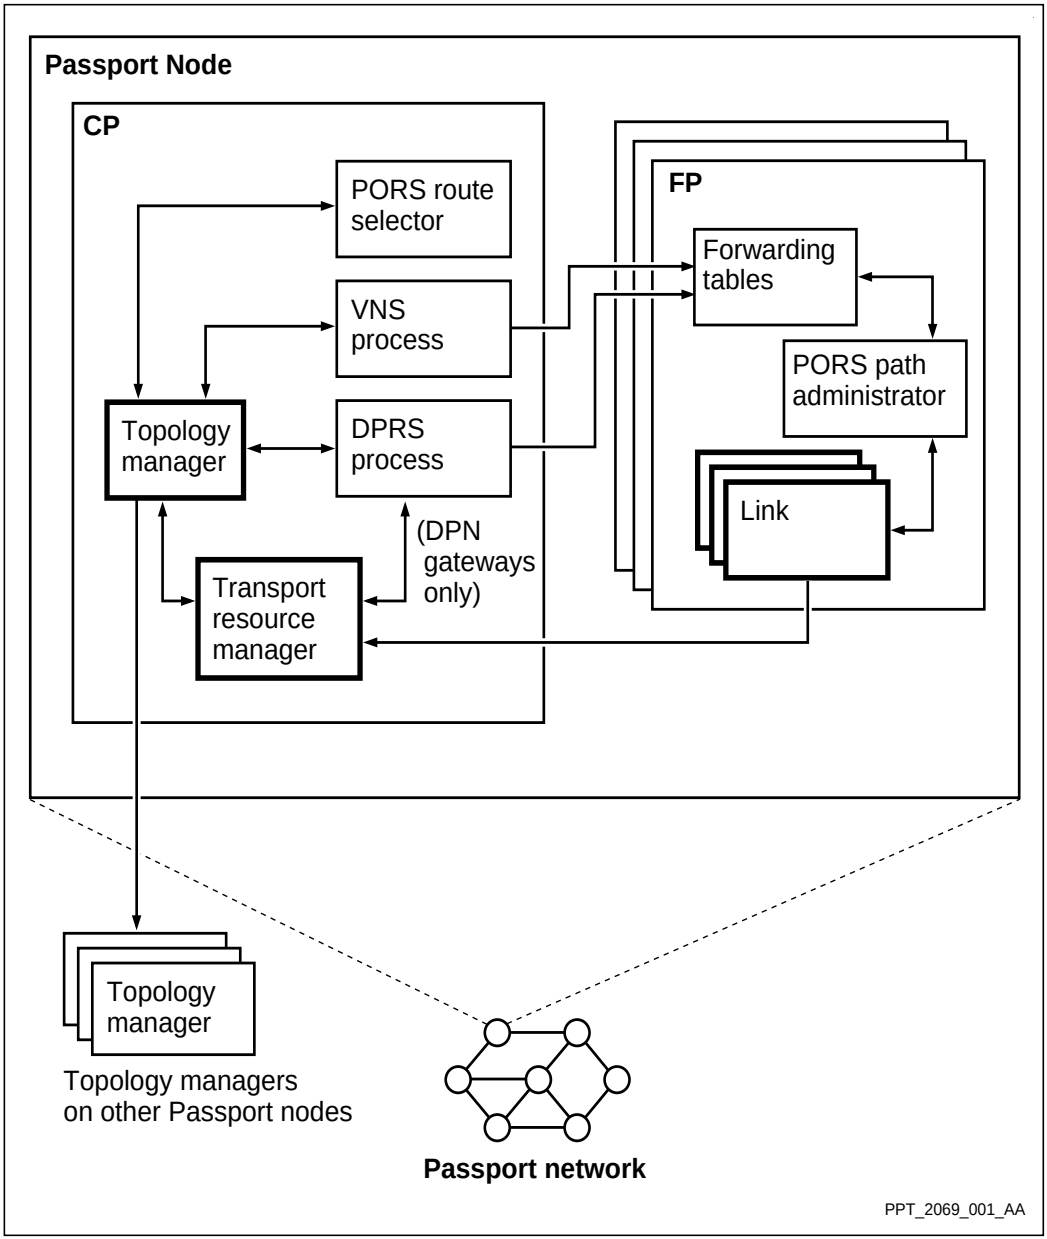

TRM mechanisms

This section describes information exchange for links and link groups, and the role of TRM in bandwidth reservation.

Link information exchange

The transport resource management (TRM) subsystem on a switch interfaces with the local Passport trunking system to learn the status (up or down) of all usable links (Passport trunks and DPN gateways), their attributes to support various classes of routing (such as speed and delay), and the logical network that they support.

This information is required by two sources:

- the topology manager to learn the link state information, and DPRS to learn of gateways
- routing protocols to compute the metrics advertised for each of the available Passport links

TRM maintains the logical network connectivity between neighbor nodes.

Link group information exchange

The TRM subsystem organizes multiple links to neighbor nodes into link groups. TRM builds and maintains link group tables and link MPID tables (see “MPID” (page 96)). TRM broadcasts these tables to all processors on the shelf for use by the packet forwarding systems.

Link groups are used to carry DPRS traffic. PORS can make use of all the links in a link group to carry traffic. However, it treats each Passport trunk individually and does not make use of the link group concept.

There can be up to four links in a link group, and each node can support up to 255 link groups. Link groups, made up of Passport trunks or DPN gateways, are included in this total. Links in a link group may have differing characteristics, such as speed and delay, but support the same set of logical networks and maximum transmission unit (maximum packet size).

Preferred links include some or all links in a link group that carry traffic for DPRS. For example, a high delay satellite link in a link group would not be chosen as a delay preferred link if a terrestrial link also exists in that link group. TRM determines throughput and delay preferred links. DPRS packet forwarding uses both throughput preferred links and delay preferred links.

Bandwidth reservation and TRM

TRM also learns of the reserved PORS bandwidth of each link. The PORS bandwidth reserved on each link limits the number of PORS connections that can be established over that link. The total amount of peak or average reserved bandwidth for the connections cannot exceed the link's PORS reserved bandwidth.

DPRS calculates metrics based on the bandwidth available once the PORS reserved bandwidth has been subtracted from total available bandwidth. For more information see 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide*.

The link override information is also learned by TRM for DPRS. If a link speed (bandwidth) or delay override is provisioned, the override is selected as the prevailing link bandwidth and or delay.

Topology regions concept

A topology region is an autonomous group of Passport nodes, where the complete topology of all Passport nodes in that portion of the network is known to the topology system. The topology of the rest of the network remains unknown to the Passport nodes within the topology region.

Provisioning topology regions within a network will improve network scaling because nodes only need to know the topology of their own region. Topology regions create greater network autonomy and make it possible to increase the size of networks by allowing nodeIds to be reused in different topology regions.

TRM provides the infrastructure to support topology regions by detecting neighbors that are part of a different region. The topology manager is not informed of such links. See "Topology regions" (page 68) for more details.

TRM components

The *TransportResource (Trm)* component resides on all Passport modules and is responsible for managing link resources. This component acts as an intermediary between a link, and all the routing systems that use the links.

For *Trm* component state combination information see the table “Transport Resource Manager component state combination” (page 64).

Table 3
Transport Resource Manager component state combination

Combination (Administrative, Operational, Usage)	Details
Unlocked, Enabled, Active	The *TransportResource* component is operational and is communicating with each of the Passport trunks and gateways configured in the module. It is able to receive and process requests.

For more information on components, see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Topology manager

This section introduces the topology manager subsystem of the base routing system, and includes the following:

- “Role of topology manager” (page 64)
- “Topology features” (page 65)
- “Topology architecture” (page 65)
- “Topology mechanisms” (page 66)
- “Topology components” (page 67)

Role of topology manager

The topology manager subsystem manages the topological view of the Passport nodes in the network.

A Passport node needs full knowledge of the network in which it is working. This knowledge encompasses two dimensions: the connectivity of other nodes, and the characteristics of interconnecting links and link groups between all nodes.

Topology features

The Passport topology system has the following characteristics:

Computes paths

Topology computes efficient paths for connectionless traffic and enables almost instantaneous rerouting when topology changes occur. Topology supports both the delay and throughput RCOS (DPRS delay and throughput).

Supplies link attributes to PORS

Topology distributes link attributes such as cost, delay, and bandwidth utilization to PORS. PORS uses this information to compute paths.

Supports logical networks

Topology provides the base support for logical networks. Logical networks provide the capability to logically partition physical networks.

Topology architecture

The topology manager system is an element of the base routing system, along with the transport resource manager system. See figure “Passport networking architecture” (page 27) for an overview of the Passport networking architecture.

Topology manager and the network layer

Topology manager is located in the network layer. It retrieves information from and distributes information to other elements of the network layer.

The topology manager resides on the control processor of each node and handles information as follows:

- receives information from TRM located on the control processor
- sends information to the routing system components on the control processor and to the other topology managers on other nodes

The figure “Passport base routing system - transport resource manager and topology manager” (page 61) illustrates the interaction of base routing components.

Topology mechanisms

The topology manager supports several mechanisms to implement its functionality.

Local node discovery

The topology manager interworks with TRM to learn local link characteristics and status such as all the link groups connecting to neighbors. TRM also notifies topology of any link state changes such as connectivity, bandwidth and delay. The topology manager updates its database with this local node information.

Neighbor discovery

Passport nodes automatically learn the network topology upon establishing connectivity to the network. When the topology system detects a link to a new neighbor node, it initiates the neighbor staging protocol. This protocol effectively synchronizes the topology databases on both nodes. This process ultimately ensures that up to date topology information is learned immediately and dynamically by new nodes as they join the network.

Network internal topology discovery

On each Passport node, the topology manager is responsible for learning the network topology. It propagates link state information about its own node to topology managers on all other switches within its topology region and maintains its own topological database. Each topology manager learns the complete network topology through this process such that the topology systems on all nodes maintain an identical view of the Passport network. This view includes all Passport nodes in the network (or the topology region if they are deployed in the network), and the links (or link groups) between them.

Path computation

The topology manager computes optimal paths through the network for each connectionless routing system using the shortest path first algorithm. The paths are generated by using a distributed link-state routing protocol based on the IP standard OSPF (open shortest path first). The trees contain network path information. The topology manager provides the summary information to the routing manager for DPRS, and acts as a database for PORS.

Topology also provides quick, transparent recovery for all virtual circuits (VCs) or connectionless flows. Topology reacts quickly to network events and computes new paths to make use of new facilities and to avoid failures and congestion.

Supporting logical networks

Topology knows the logical network numbers (LNNs) for the link groups and carries this information to every node in the network.

Topology components

The *Topology (Top)* component is a subcomponent of *Routing (Rtg)* residing on all Passport modules. The *Topology* component maintains a topological database describing the Passport network topology in terms of nodes and connectivity between them. It computes the reachability costs to each Passport node for the different routing systems taking into account the underlying transport attributes (for example delay, throughput) of interest to them.

For *Topology* component state combination information see the table “Topology component state combination” (page 67).

Table 4
Topology component state combination

Combination (Administrative, Operational, Usage)	Details
Unlocked, Enabled, Active	Topology is exchanging routing protocol data units with Passport neighbors. It is capable of mapping additional Passport nodes in its topological database.
Unlocked, Enabled, Busy	Topology is operational, but its topological database cannot accept more Passport nodes. The number of Passport nodes stored in the topology database has reached the maximum network node number (256).

For more information on components, see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Topology regions

This section describes the concept of topology regions, which allow for large network growth. The following information is described:

Network impact of topology regions

As the size of a network approaches 1 000 Passport nodes, the network administrator may consider deploying topology regions. Up to 127 topology regions can be configured within a network. Each topology region can support up to a recommended engineering limit of 1 000 Passport nodes.

Region identifiers (regionIds) are used to indicate which topology region a node belongs to. The *regionId* attribute is a provisionable attribute of the *ModuleData* component. Nodes exchange their *regionId* attribute values during Passport trunk staging, in order to determine whether or not they belong to the same topology region. A neighbor node is assumed to be part of the same topology region if the region identifiers of the two nodes match.

In networks where topology regions are not required, all nodes can be left with their region identifier attributes set to the default of zero.

A RID subnet (see "RID subnets" (page 97)) is a collection of interconnected Passport nodes. Passport trunks between different topology regions, or inter-region Passport trunks can only be defined at RID subnet boundaries; therefore, all nodes within one RID subnet must belong to the same topology region. A topology region can consist of more than one RID subnet.

An example of topology regions is illustrated in the figure "Topology regions" (page 69). The figure demonstrates that several topology regions can exist in a single Passport network but that each RID subnet must be contained within a single topology region.

Figure 14
Topology regions

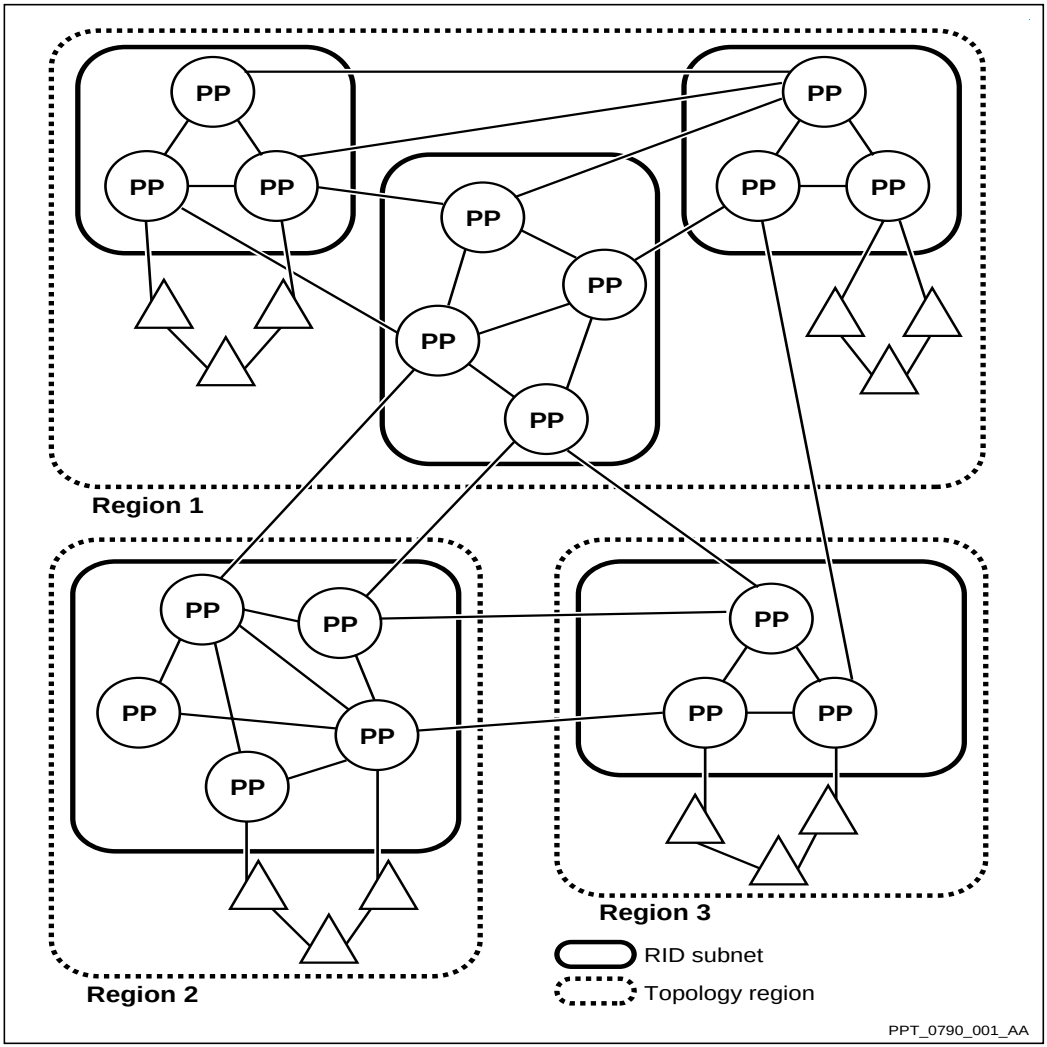

Topology regions impact on Passport routing systems

Topology regions impact each Passport routing system differently.

There is no impact on DPRS when topology regions are deployed within a network. Traffic flows remain unchanged because topology regions are deployed along RID subnet boundaries. A RID subnet must be configured in order to create alternate region configurations. The Passport Multiple Priority System (MPS) and congestion management are fully supported across inter-region Passport trunks.

Dynamic routing of PORS services that use NSAP addressing (CES and VTDS) is possible within and between topology regions or clusters. PORS services that do not use NSAP addressing can be dynamically routed between nodes within a topology region or nodes within a cluster only. For all services, a manual path can be defined in order to make a PORS connection between topology regions or clusters.

Topology regions and nodeId reuse

Topology information is not exchanged between separate regions; therefore, nodeId reuse is possible under certain conditions. The implications of nodeId reuse should be analyzed however, before deploying it in the network.

In order to reuse nodeIds, both regions must possess independent network management systems. The nodeId must be unique across nodes that are managed by the same network management administration. If Passport nodeIds are reused in different topology regions, the network management and accounting systems must be regionalized also. If nodeIds are reused, network tools that provide data to systems (such as the Management Data Provider that produces Bulk Data Format files) will not be able to assimilate data from more than one network management region.

The nodeIds are independent of topology region boundaries. Therefore, nodes from different topology regions can connect, even if their nodeId values are the same.

Note: The reuse of Passport node names is not supported.

Deploying topology regions

Before deploying a topology region, all RID subnets must be configured to allow the deployment of region boundaries. That is, all topology region boundaries must fall between subnet boundaries.

The region identifier value can be set to a value between 0 and 126 inclusive. If topology regions are not deployed, the recommended value for the region identifier is zero.

CAUTION

Software upgrade issue

In a network where topology regions are partially deployed (that is, some nodes within a region have a region identifier of 0), the *regionId* value of all nodes in the topology region must be leveled (changed to the same value) prior to the software upgrade of any node in the topology region. Failure to do so results in unexpected behavior on the Passport trunks.

All nodes interior to regions should have region identifier values assigned. This assignment is essential when performing a split of a topology region.

> ***Note:*** It is a recommended practice to also clear the *splittingRegionIds* attribute upon completion of splitting a topology region. This eliminates the possibility of a topology region merge from occurring.

If Passport trunks are deployed over ATM, identifying border nodes can be complicated. Nodes that appear to be interior topology region nodes, based on their physical topology, are effectively border nodes.

Topology regions can be established by configuring the *regionId* attribute on a set of nodes within a network. Refer to the following sections for details:

- “Splitting topology regions” (page 71)
- “Merging topology regions” (page 78)

Splitting topology regions

Splitting a topology region introduces a new region into the network, with its own region identifier.

Prerequisites

- Ensure that all nodes have the same *regionId* value. In a case where topology regions are not deployed, a *regionId* value of 0 is recommended for all nodes.

- Ensure the Routing IDs (RIDs) of the nodes in the current topology region are different from the RIDs of the nodes in the future topology region. Upon splitting a region, if the nodes of different topology regions share common RIDs, the inter-region Passport trunks will not stage properly and will remain down.
- 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 all border nodes comply to RoutingGateway (RGty) performance recommendations.
- All services crossing proposed region boundaries must have NSAP addressing provisioned. Any services using the *remoteName* attribute must be converted to the *addressToCall* attribute prior to starting this procedure. This step is only applicable if PORS traffic will exist across the planned region border. Refer to 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide* for details.
- For a set of future inter-region links (IRLs) between any two regions, it is recommended that the *overrideTrunkDelay* attribute under the *trunk pa* component be configured to an identical value. This step is only applicable if PORS traffic will exist across the planned region border. Refer to 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide* for details.
- The network must be stable prior to splitting topology regions.

Procedure steps

1 Configure routing gateways on all border nodes. Refer to 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide* for details.

Note: This step is only required if PORS traffic will exist between the planned region border.

2 Configure the *splittingRegionIds* attribute value for all nodes changing regionId:

```
set rtg splittingRegionIds <old_region_id>
<new_region_id>
```

3 Configure the *regionId* attribute value for all non-border nodes belonging to the future topology region:

```
set mod regionId <new_region_id>
```

4 Configure the *regionId* attribute value for all border nodes belonging to the future topology region. Once the last border node is changed, the existing region is split into two regions:

```
set mod regionId <new_region_id>
```

5 Configure NSAP link reachable addresses on all IRLs. Refer to 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide* for details on configuring and verifying reachable addresses from the local topology database.

Note: This step is only required if PORS traffic exists between topology regions.

6 After one hour:

a. verify the topology has aged out:

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

b. verify multi-region routing. Use the RTG RS to verify that the selected route is the same as the expected path.

7 Segment all services traversing the new inter-region Passport trunk. This is achieved by either issuing a lock -force command on all inter-region Passport trunks, or by restarting the service(s).

Note: This step is only required if PORS traffic exists between topology regions.

8 Clear the value of the *splittingRegionIds* attribute for all nodes in the new topology region:

```
set rtg splittingRegionIds !
```

Variable definitions

Variable	Value
<old_region_id>	The region identifier of the current topology region. The value can be any number from 0 to 126 inclusive.
<new_region_id>	The region identifier of the future topology region. The value can be any number from 0 to 126 inclusive.

Example of splitting topology regions

The following is a step-by-step example of how to split a topology region.

Note: This example is based on a DPRS network. It does not include the additional steps required for PORS traffic.

Before you split a topology region, confirm that all nodes in the network have the same *regionId* value. In figure "Splitting topology regions—one regionId" (page 74), a network consisting of two RIDs exists with one region identifier associated with all nodes (in this case, the *regionId* value is 60).

Figure 15
Splitting topology regions—one regionId

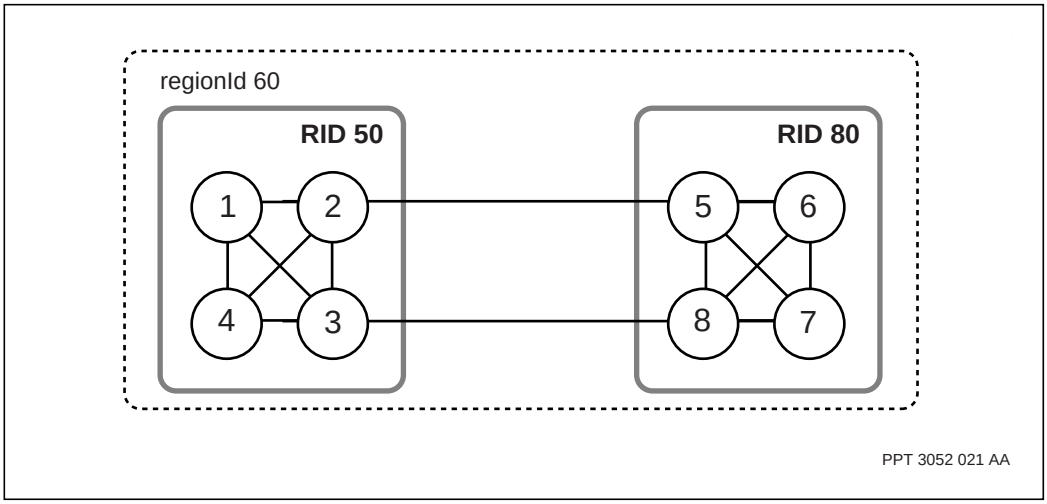

The first step is to configure the *splittingRegionIds* attribute value for all nodes that will belong to the new topology region. In figure "Splitting topology regions—one regionId, configure splittingRegionIds" (page 75), each node in the subset (5, 6, 7, and 8) is assigned a *splittingRegionIds* value of 60 and 90 (60 is the *regionId* of the current topology region, 90 is the *regionId* assigned to the future topology region). If only a subset of these nodes were assigned the *splittingRegionIds*, then those nodes would become isolated.

```
set rtg splittingRegionIds 60 90
```

Figure 16
Splitting topology regions—one regionId, configure splittingRegionIds

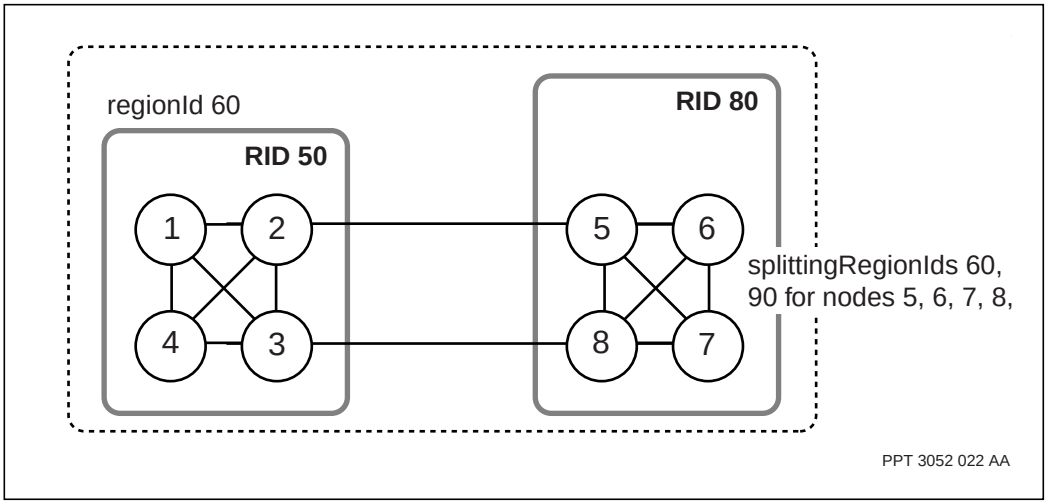

The second step is to configure the non-border nodes in the new topology region. Referring to figure “Splitting topology regions—configure non-border nodes” (page 76), nodes 6 and 7 are non-border nodes in the new region. For each of these two non-border nodes, enter the following command:

```
set mod regionId 90
```

Figure 17
Splitting topology regions—configure non-border nodes

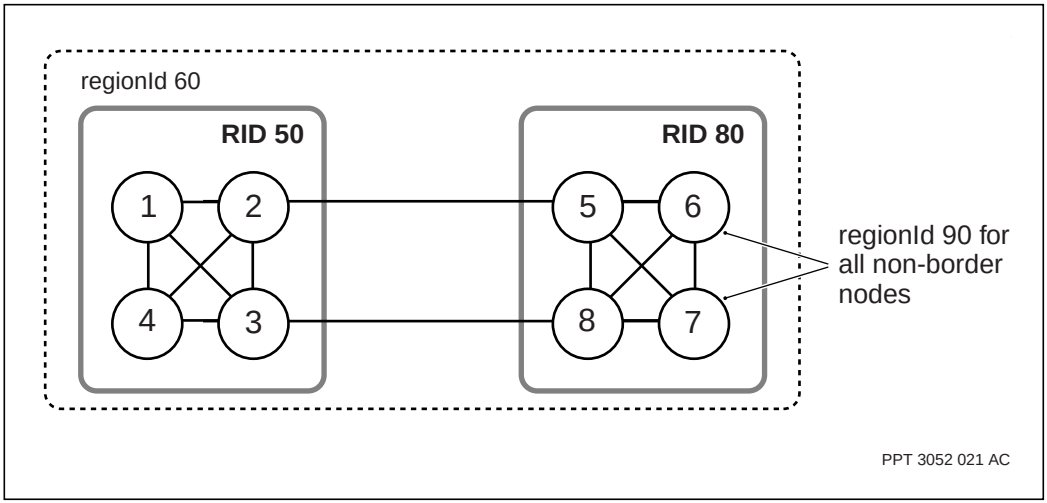

The third step is to configure the border nodes in the new topology region. Referring to figure “Splitting topology regions—configure border nodes” (page 77), nodes 5 and 8 are the border nodes to be configured. For each of these two border nodes, enter the following command:

```
set mod regionId 90
```

Figure 18
Splitting topology regions—configure border nodes

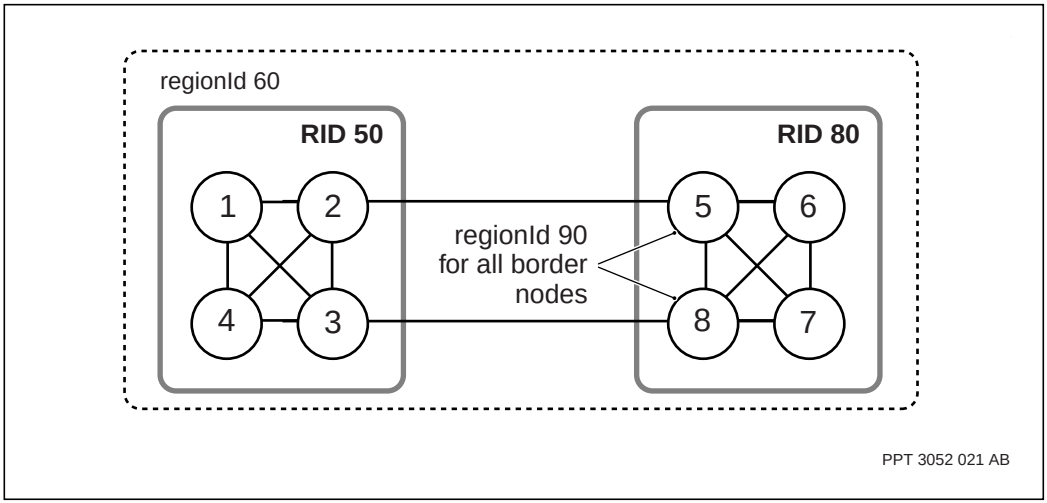

The instant the last border node is changed from 60 to 90, a topology split occurs. After allowing an hour for the topology database to clear, verify that multi-region routing is functioning. To check to see if the topology database has cleared, enter the following command:

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

In this case, if the command is issued on node 1, expect to see nodes 1, 2, 3, and 4 displayed, which means that nodes 5, 6, 7, and 8 have aged out, and are now part of the new topology region.

The final step is to clear the value of the *splittingRegionIds* attribute for the nodes in the new topology region. In this case, the *splittingRegionIds* attribute value of 60 and 90 on nodes 5, 6, 7, and 8 needs to be cleared:

```
set rtg splittingRegionIds !
```

After being cleared, a multi-region environment exists, as shown in figure “Splitting topology regions—two regionIds” (page 78).

Figure 19
Splitting topology regions—two regionIds

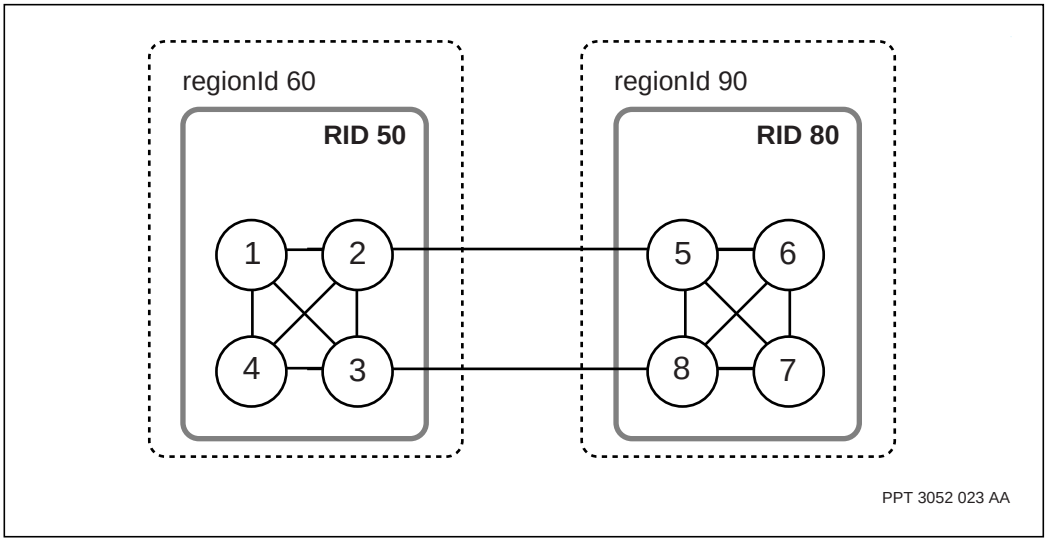

Merging topology regions

Merging topology regions reduces the number of topology regions in the network. A merge takes two regions with different region identifiers and combines them into one topology region with one region identifier. This migration method is the reverse of splitting a region. For more information on splitting topology regions, refer to “Splitting topology regions” (page 71).

Prerequisites

- Ensure that the two regions being merged already have connectivity.
- Ensure that the two regions have different *regionId* values.
- Ensure that there are no duplicate *nodeId* values. Nodes within the two topology regions must all have different *nodeId* values.
- Ensure that the number of nodes in the two regions, when combined, do not exceed the maximum number of recommended nodes in a topology region.

Procedure steps

1 Configure the *splittingRegionIds* attribute value for all nodes belonging to the topology regions that is being merged:

```
set rtg splittingRegionIds <old_region_id>
<new_region_id>
```

2 Configure the *regionId* attribute value for all non-border nodes belonging to the topology regions that is being merged:

```
set mod regionId <old_region_id>
```

3 Remove NSAP link reachable addresses from the IRLs. Refer to 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide* for details.

Note: This step is only required if PORS traffic exists between topology regions.

4 Allow five minutes for the network to stabilize.

5 Configure the *regionId* attribute value for all border nodes belonging to the topology region that is being merged. Once the first border node is changed, the two regions are merged into one region:

```
set mod regionId <old_region_id>
```

6 Unsegment all services that were traversing the inter-region Passport trunk. This is achieved by either issuing a lock -force command on all inter-region Passport trunks, or by restarting the service(s).

Note: This step is only required if PORS traffic was routed between the former region borders.

7 Clear the *splittingRegionIds* attribute value for all nodes belonging to the topology regions that is being merged:

```
set rtg splittingRegionIds !
```

8 Optionally, remove all routing gateways from all border nodes. Refer to 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide* for details.

Note: This step is only required if PORS traffic was routed between the former region borders.

9 Optionally, remove the routing gateway feature from the feature list of the *logicalProcessorType* attribute for that logical processor (LP). This causes the LP to reset. Refer to 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide* for details.

Note: This step is only required if PORS traffic was routed between the former region borders.

Variable values

Variable	Value
<old_region_id>	The region identifier of the topology region that another topology region is merging into. The value can be any number from 0 to 126 inclusive.
<new_region_id>	The region identifier of the topology region that is being merged into the other topology region. The value can be any number from 0 to 126 inclusive.

Passport clusters

Passport clusters offer further possibilities for network growth. Passport clusters further partition a topology region by connecting groups of Passport access switches around a region's backbone. Passport clusters do not exchange topology information with the backbone, and only a limited amount of routing information is exchanged between the backbone and clusters. As a result, memory and CPU requirements to maintain the backbone or cluster's topology database are reduced.

Refer to the figure "Passport clusters" (page 81) for an illustration of Passport clusters.

Figure 20
Passport clusters

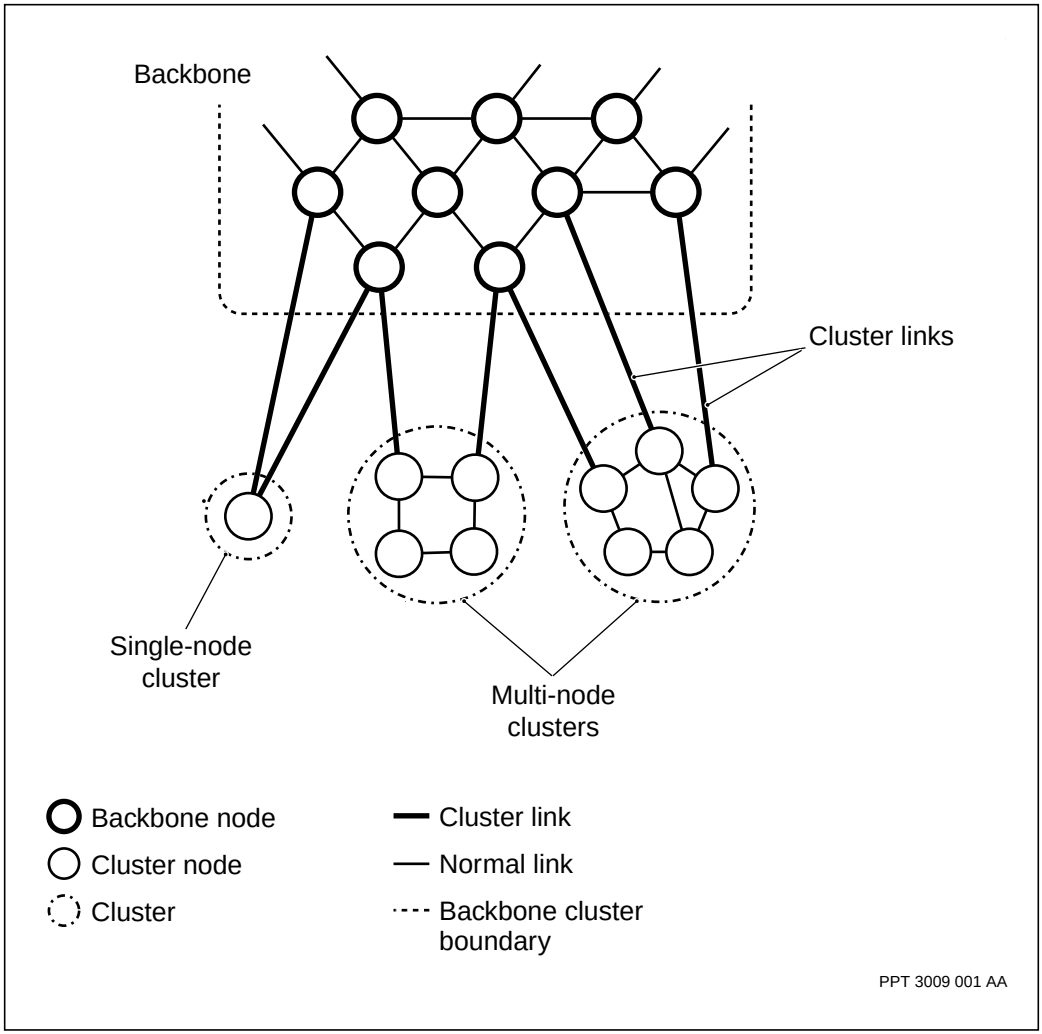

Passport clusters allow RID subnets and topology regions to grow beyond existing engineering limitations. The number of Passport cluster nodes, backbone nodes, and access modules (AMs) continues to be limited by the

number of MIDs in an RID subnet. The combined total number of Passport cluster nodes and backbone nodes configured in a topology region is limited by the node ID range.

For more information on Passport clusters see 241-5701-425 *Passport 7400, 15000, 20000 Dynamic Packet Routing System Guide* or 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide*.

Tandem suppression

The tandem suppression feature gives the network operator increased control over routing behavior by preventing tandem traffic from travelling through selected nodes. When tandem suppression is enabled on a node, the only traffic routed to the node is traffic destined for the node itself.

Tandem suppression is particularly useful for small access Passport nodes and 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

The default value of the tandemTraffic attribute is allowed, permitting tandem traffic to pass through the node. When this attribute is set to denied, DPRS and PORS traffic is normally not permitted to tandem through the node.

There are some exceptions with DPRS gateway nodes and backbone cluster border nodes where traffic can tandem through a node even with the *tandemTraffic* attribute set to denied:

- On a RID gateway node (the nodes that are connected to another Passport RID subnet or to DPN-100 RMs) or a MID gateway node (the nodes that are connected to DPN-100 AM clusters), traffic that flows to or from the subnet is allowed to tandem through the node. Traffic between nodes in the same subnet is not allowed to tandem through the gateway node when this attribute is set to denied.

- On a cluster border node, traffic that is destined for outside the cluster (to the backbone or to another cluster) or comes into the cluster is allowed to tandem through the node, even if this attribute is set to denied. Traffic between cluster nodes in the same cluster is not allowed to tandem through the cluster border node when this attribute is set to denied.
- On a backbone border node, traffic that is destined for outside the backbone (to a cluster) or comes into the backbone (from a cluster) is allowed to tandem through the node, even if this attribute is set to denied. Traffic between backbone nodes in the same RID subnet is not allowed to tandem through the backbone border node when this attribute is set to denied.

For PORS calls, changing the *tandemTraffic* attribute has no impact. The calls can be manually forced to reroute using the new setting by clearing the calls, or by locking the trunks on which they are routed.

Chapter 3
Dynamic Packet Routing System (DPRS)

This chapter describes the Passport Dynamic Packet Routing System (formerly referred to as DPN routing and RID/MID routing), and includes the following information:

For more detailed information on DPRS and DPRS services (frame relay, and DPN-100), see "DPRS references" (page 18).

A RID subnet is an important DPRS concept to understand. A RID subnet is a collection of interconnected Passport nodes. See "RID subnets" (page 97) for more details.

Role of DPRS

DPRS provides efficient connectionless routing for delay-sensitive and high-throughput variable bit-rate (VBR) traffic. DPRS is ideally suited to carry data traffic such as frame relay and the supported DPN-100 services such as X.25.

DPRS routing features

This section describes the DPRS routing features.

Dynamic reaction to topology changes

DPRS uses the information provided by systems such as the base routing component to maintain forwarding tables at each node. Any topology changes, such as new nodes or facilities, failures, and congestion situations, are incorporated in the tables within milliseconds. Topology changes are also considered in automatic packet forwarding and rerouting decisions. No provisioning is required.

Multipath and multilink traffic balancing

Multipath traffic balancing enables the use of different paths through different link groups that exist between source and destination nodes. Multilink options balance the traffic across links within one link group. “Load sharing” (page 101) is a multilink process while “Load spreading” (page 100) is both a multipath and multilink process. These traffic management mechanisms allow for increased application throughput.

Quality of service attributes

DPRS supports

- delay, throughput, and multimedia RCOS (routing class of service) attributes
- priority and reliability attributes

These attributes provide applications with the forwarding characteristics required (for example, use minimum delay for delay-sensitive traffic).

Scaling ability

DPRS offers network scaling characteristics that include the following:

- hierarchical addressing requires a minimal amount of routing information to be shared between hierarchies
- architectural separation of a connection-oriented application from a connectionless routing backbone reduces the overhead in path selection. Less resource consumption such as memory and computation time is required to support additional connections to a particular destination.

Deterministic

Paths to each destination are selected based on optimal link characteristics. However, because multiple equal paths are available, paths are selected deterministically. This approach ensures that routes through the network are fully predictable.

Efficient connection establishment

Routing calculates the best paths to each node once. No extra routing calculation is required to establish a new connection to the same destination. These paths can be used for all connections to each destination. There is no overhead for a connection on tandem nodes.

Dynamic trunk speed change

The dynamic trunk speed change feature enables Passport trunking and routing to adapt to bandwidth changes without taking Passport trunks out of service. For details, see 241-7401-110 *Passport 7400, DPN-100 Interworking Guide*.

Advanced services

Passport provides two advanced routing services for DPRS applications: call redirection, and hunt groups.

A call redirection server (CRS) provides call redirection for DPRS services, such as frame relay, in Passport-only networks. Call redirection servers direct failed call attempts to alternative destinations. A CRS improves the frame relay service's availability by redirecting a call attempt that would otherwise fail when the destination cannot be reached. The RID redirection capability of the CRS can be used for splitting a large RID subnet into smaller RID subnets. For more information on call redirection, see 241-5701-410 *Passport 7400, 15000, 20000 Call Redirection Server Guide*.

A hunt group is a single data network address (DNA) that represents a group of service DNAs. When users call the DNA assigned to the hunt group, the server forwards the call to one of the hunt group members. For more information on hunt groups, see 241-5701-415 *Passport 7400, 15000, 20000 Hunt Group Server Guide*.

DPRS architecture

This section describes the DPRS architecture. The topics include the following supporting architecture layers:

- "Passport trunking system for DPRS" (page 88)
- "DPRS and base routing components" (page 89)
- "DPRS virtual circuit (VC)" (page 91)
- "DPRS services" (page 93)
- "DPRS traffic management features" (page 93)

"Passport networking architecture" (page 27) illustrates where DPRS fits into the Passport network layer.

Passport trunking system for DPRS

The Passport trunking system supports DPRS by providing link information such as availability, speed (bandwidth), and remote node identification. DPRS uses this link information for packet forwarding across Passport nodes and links.

The Passport trunking system manages the links interconnecting Passport nodes or interconnecting Passport to DPN-100. A connection between two Passport RID subnets (see "RID subnets" (page 97)) is referred to as an Internal Gateway for DPRS purposes. All other routing systems simply treat the link as a regular Passport trunk.

DPRS uses two link group metrics:

- delay metrics based on the average delay of the preferred links in the link group
- throughput metrics based on the reported speed of, or total bandwidth available across, all the links in a link group

For more information on Passport trunking, see "Passport trunking and base routing systems" (page 49).

DPRS and base routing components

DPRS packet forwarding across Passport nodes and links uses information from the DPRS and base routing systems. The figure “DPRS and base routing architecture” (page 90) shows the DPRS and base routing components and the information flows.

Figure 21
DPRS and base routing architecture

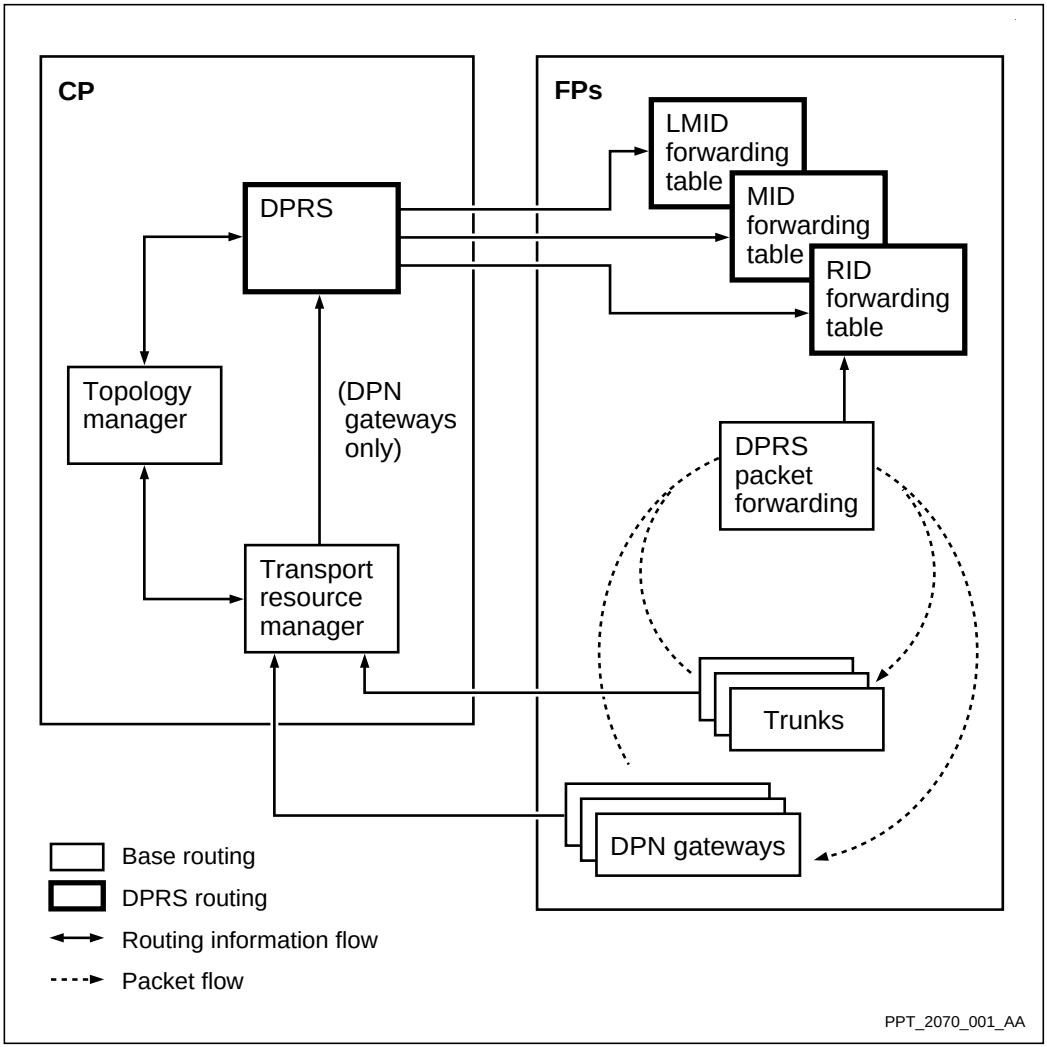

Transport resource manager supporting DPRS

Each link on the shelf registers with TRM, to provide information such as neighbor information, throughput, round trip delay, and percent of bandwidth used by PORS. The links also report link status (up or down) to TRM.

The TRM system assimilates the link information it receives to form link groups. TRM determines the preferred links for both throughput and delay routing classes of service for DPRS packet forwarding. Only preferred links will be used to carry traffic.

After TRM has determined the preferred links in the link group, it passes the link information to the topology routing system and, when required, to the DPRS system. Links between RID subnets, or between Passports and DPN-100 modules are passed on to the DPRS routing system. Information about links to DPN-100 modules is not propagated to the topology system.

Topology manager supporting DPRS

DPRS uses the Passport topology system. It exchanges topology information between Passport nodes and creates a subnet-wide view of the Passport topology.

DPRS

DPRS interacts with the TRM and topology managers. The DPRS functions are to

- learn and exchange RID routing information between RID subnets
- learn MID routing information from the subnet nodes and DPN-100 AMs connected to the subnet
- learn from the topology system the view of the RID subnet interconnectivity
- use the information learned above to compute and maintain RID, MID, and call server forwarding tables

DPRS virtual circuit (VC)

A virtual circuit (VC) uses network resources such as the routing and Passport trunking systems to establish, maintain, and terminate logical connections across the Passport subnet for an application service. The VC provides reliable and ordered data delivery to DPRS applications such as frame relay. The call establishment process sets up the VCs (see “DPRS call establishment” (page 103)).

VC functions

Virtual circuits can perform many functions, with some VCs performing a subset of all potential VC functions. The VC functions are as follows:

- Define the endpoints of the connection.
- Establish and terminate connections through the network.
- Interface to the access service to provide network connectivity with the required quality of service.
- Maintain statistics on the quality of the logical connection.
- Take action upon failure of component on the connection's path.
- Check integrity and ordering of packets within a connection.
- Collect accounting information related to the connection.
- Segment and reassemble frames/cells if necessary.

DPRS VC connections

The following two mechanisms for establishing connections are used:

- permanent virtual circuits (PVCs): These connections have their network endpoints provisioned and are automatically established by the network. PVCs are used by frame relay and for network management connections (IPIVC and IPIFR for connecting Passport to Preside Multiservice Data Manager).
- switched virtual circuits (SVC): These connections are established through user signaling and can therefore be used to access a variety of destinations. The network establishes the connection on demand.

DPRS VC types

DPRS supports full- and light-weight VCs, providing different qualities of packet transport based on the access service requirements:

- Full-weight VCs perform end-to-end acknowledgment to guarantee delivery of packets through the network. DPN-100 and Passport interworking, and IPIVC use full-weight VCs.
- Light-weight VCs do not acknowledge the delivery of packets through the network. This delivery method is known as best effort. Frame relay (including IPIFR) use light-weight VCs.

The table “DPRS virtual circuits: Full-weight, light-weight” (page 93) details the DPRS VCs characteristics.

Table 5
DPRS virtual circuits: Full-weight, light-weight

Characteristic	Full-weight	Light-weight
Guaranteed delivery of packets	Yes	No
Ordering of packets	Yes	Yes
Elimination of duplicate packets	Yes	Yes
Segmentation and reassembly	Yes	Yes
Flow control of packets	Yes	No
Congestion notification	Yes	Yes
Adaptation to network changes	Yes	Yes

DPRS services

DPRS is the Passport routing system used to support the transport of data traffic for the following services:

- frame relay UNI and NNI access services
- DPN-100 services such as X.25, SDLC, asynchronous protocols (X.3/X.28/X.29), point-of-sale when routed through a Passport network

DPRS traffic management features

DPRS traffic management mechanisms operate primarily in the areas of route selection, bandwidth management, and congestion control. DPRS traffic management features include

- route selection and bandwidth management features: grouping links into link groups (for which DPRS determines metrics), selecting preferred links in the link group based on RCOS, subtracting percent PORS bandwidth reservation
- load balancing features: multipath load spreading, multilink load sharing or load spreading, emission prioritization

- congestion management features: discard priority levels for all traffic types, forward and backward congestion notifications, overflow routing around congestion

For more information on traffic management, see the following documents:

- 241-5701-705 *Passport 7400, 15000, 20000 ATM Traffic Management Fundamentals*
- 241-5701-706 *Passport 7400, 15000, 20000 ATM Traffic Shaping and Policing*
- 241-5701-707 *Passport 7400, 15000, 20000 ATM Queuing and Scheduling*
- 241-5701-708 *Passport 7400, 15000, 20000 ATM CAC and Bandwidth Management*

DPRS routing mechanisms

This section includes the following information:

DPRS addressing system

Addresses are used to name various elements of a network. They are stored in the header of each packet and are used by the packet forwarding system to route the packet to its destination. DPRS addresses are hierarchical. See "DPRS call establishment" (page 103) and "DPRS call routing" (page 194) for determination of the addresses specified in the packet header.

RID

At the top of the hierarchy is the routing identifier or RID. A single RID is assigned to a collection of interconnected Passport nodes known as a RID subnet (see “RID subnets” (page 97)). When forwarding packets between RID subnets, the forwarding system needs to examine only the RID field of the packet address. The RID is a number from 1 to 126. The RID is also used to identify DPN-100 RMs.

MID

The second level is a module identifier or MID. A MID is assigned to every module in the network (AMs, RMs, and Passports). A RID/MID pair can always be used to uniquely identify every module in the network. Within the destination RID subnet, packet forwarding is performed on the MID address to locate the destination module within the RID subnet. The MID is a number from 1 to 1,909. Nortel Networks engineering guidelines support up to 300 Passport nodes in a RID subnet where Passport clusters are not deployed. MIDs can be reused in other RID subnets.

LMID

A packet contains either a MID or an LMID value following the RID. Logical call services, such as source call routing (SCR), are addressed by their logical MIDs or LMIDs. Logical services exist on Passports or CSRMs, or both. CSRMs are only available on Passport 7400 series switches. Passport 15000 and 20000 do not use CSRMs. Some examples of the logical services are

- on Passports: source/destination call routing, call redirection (on Passport-only networks)
- on CSRMs: source call routing, destination call routing, gateway source and destination call routing, access to call redirection and access to NDI for network user identification (NUI) or Global Mnemonic translation

The LMID services are provided to the RID subnet by designated RMs called Call Server RMs or CSRMs. Although the CSRMs have their own RIDs, which are different from the RID subnets, they can still support call services for the RID subnet. This support is possible because routing information about call services supported on the CSRMs is distributed to the Passport 7400 series nodes in the RID subnet.

MPID

The third and lowest level in the DPRS addressing hierarchy is the module process identifier or MPID, also referred to as PID. The MPID identifies the end point process, such as a frame relay PVC, SPVC, or SVC, to route packets to within a module.

The figure "DPRS RID/MID routing hierarchy" (page 96) illustrates the DPRS routing hierarchy.

Figure 22
DPRS RID/MID routing hierarchy

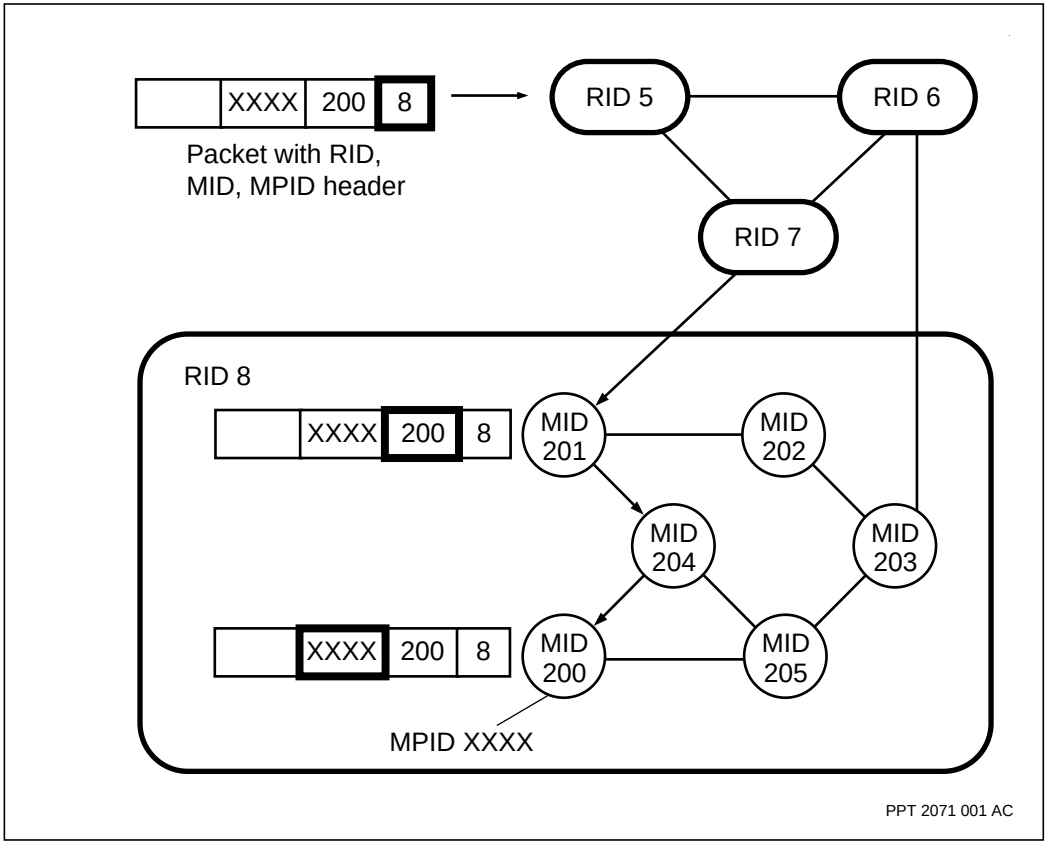

RID subnets

Grouping Passports into RID subnets allows for very large network growth due to the efficient use of addressing space, the advantages of hierarchical routing, and the efficient propagation of network routing information.

Passport RID subnets can be directly interconnected using all Passport trunk types. DPRS propagates routing information both within and between RID subnets as follows:

- distributed link-state routing: within a RID subnet, MID information is propagated to all nodes (except cluster nodes, if present in the subnet)
- distributed distance-vector routing: between RID subnets, only RID information is exchanged

Because detailed MID routing information is not propagated to neighbor subnets, the propagation and the size of routing updates within the network are minimized. The efficient use of bandwidth by DPRS results. Also, efficient use of CPU and memory is achieved for the construction and storage of DPRS routing tables.

Within a RID subnet, DPRS dynamically determines the best paths to other RID subnets. No special configuration is required. Upon establishing connectivity to another RID subnet, DPRS dynamically learns network RID topology information and distributes it within the RID subnet. This assimilation and distribution of information enables Passports within the RID subnet to determine their best paths to other RID subnets in the network.

Passport clusters

Passport clusters further partition a RID subnet by connecting groups of Passport access switches around a subnet's backbone. Passport clusters exchange only limited routing information, and no topology information, with the backbone and other clusters.

Passport clusters are not included in the backbone's topology view. This results in an improvement in the engineering limitation on the total number of Passport nodes in a topology region.

See 241-5701-425 *Passport 7400, 15000, 20000 Dynamic Packet Routing System Guide* for more information on Passport clusters and DPRS.

DPRS packet header

The DPRS packets have a Passport common header, followed by a DPRS header, and then the actual data. The DPRS forwarding system uses the information in the header to forward the packet toward its destination. See the figure "DPRS packet format" (page 98).

Figure 23
DPRS packet format

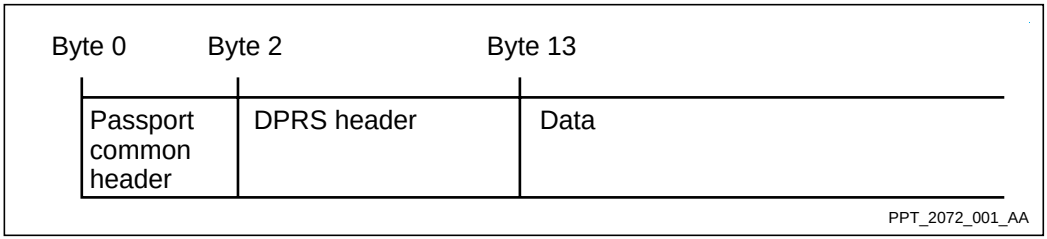

Passport common header

A small Passport common header is appended to the front of all packets routed between Passport nodes. This header identifies that a DPRS packet follows, whether congestion was encountered by the packet, and the packet discard priority. The following header information is included:

- The routing mode identifies the packet type and its header format (example, value of 1 for DPRS).
- The logical network number (LNN) field determines the logical network to which the packet belongs. DPRS uses an LNN of 1.
- The packet congestion indicators identify congestion encountered by the packet. The forward congestion indicator (FCI) identifies congestion encountered on its way to the destination node. Congestion can also be signalled back to the source node of traffic with the backward congestion indicator (BCI).
- The packet discard priority is compared to the link discard level (used to indicate the degree of congestion on the link) in the forwarding tables to determine the importance of forwarding the packet within the link group. The packet discard priority is derived from a combination of the packet type and the priority attribute (see "DPRS header" (page 99)).

- If the packet contains multimedia data for a frame relay service, there is an emission priority indicator in the header. This practice ensures that delay-sensitive multimedia traffic is guaranteed a high quality of service in transport throughout the network. See “DPRS quality of service attributes” (page 99).

DPRS header

The DPRS header provides addressing information, quality of service attributes, and packet type. The following header information is included:

- The destination address (RID, MID, MPID) selects an outgoing Passport trunk or DPN gateway or a service. DPN gateway is only available on the Passport 7400 series switches. Passport 15000 and 20000 do not use DPN gateway. See “DPRS addressing system” (page 94) for details on the routing, module, and process identifiers in the address.
- The quality of service attributes assist in path selection and describe how a packet is treated while it is being routed. The DPRS attributes include delay and throughput RCOS, priority, and reliability. See “DPRS quality of service attributes” (page 99).
- The packet type identifies the type of information in the packet. Examples include call request, call accept, data, and data acknowledgment packets.

DPRS forwarding systems

Packet forwarding is a function that routes packets in the Passport network. A DPRS connectionless packet is forwarded to its destination through nodes, link groups, and links on each node. DPRS forwarding performs link group and link selection through link group tables and link tables built by DPRS routing. Each packet contains the information required by packet forwarding to access the tables and make a next hop route selection

DPRS quality of service attributes

DPRS supports throughput, delay, and multimedia RCOS attributes for path selection, and priority and reliability attributes for path access.

The routing system calculates and assigns a set of metrics (delay, throughput pair) to each link group. Metrics are numbers that determine the best path through a Passport network. The metric is calculated based on the throughput or delay of the links in a link group. The higher the throughput or the lower the delay, the smaller the metric.

The best path to a specific destination is the one for which the sum of all the link group metrics between the source and destination is a minimum. The DPRS system uses link group metrics to determine the best paths from any node to any destination node for the delay and throughput classes of service. The multimedia RCOS traffic is forwarded through the delay-preferred links.

Under congestion situations, access to paths depends on the value of the priority and reliability attributes specified in the DPRS packet header. Priority indicates the packet's importance. Higher priority packets are less likely to be discarded due to congestion situations. Reliability indicates the packet's access to alternate paths. Higher reliability packets are allowed to use alternate paths, if available, under congestion situations.

The classes of service supported by Passport routing systems are described in 241-5701-705 *Passport 7400, 15000, 20000 ATM Traffic Management Fundamentals*.

DPRS forwarding policies

DPRS supports load spreading, load sharing, and load spreading fast packet forwarding policies.

Load spreading

Load spreading has two distinct aspects: multilink and multipath load spreading.

Multilink load spreading occurs within a link group to

- select a link within a group for one VC's traffic
- keep that VC's traffic (individual packets) on that link, assuming no congestion and stable topology
- distribute VCs over links in a group
- try other links in a group if link congestion is encountered

Multipath load spreading occurs between link groups to

- assign VCs to next hop link groups, when multiple paths to a destination exist
- keep that VC's traffic (individual packets) on that link group, assuming no congestion and stable topology

Multilink load spreading, simply referred to as load spreading, works well for many VCs of similar size. In this case, the maximum bandwidth available to a VC is the bandwidth of one link in a group.

Load sharing

Load sharing has a single behavior called multilink load sharing.

Multilink load sharing occurs within a link group to

- distribute traffic (individual packets) from a single VC across preferred links within a link group
- overflow traffic for a VC to an uncongested link within the link group, if link congestion is encountered. If no uncongested link in the group exists, then overflow high reliability traffic to an alternative path in another link group. Normal reliability traffic will be discarded.

Load sharing works well for large frame relay VCs and mixed traffic. The maximum bandwidth available to a VC is that of the entire link group.

Multilink load spreading or load sharing behavior can be provisioned under the *Rtg Dpn* component.

See 241-5701-705 *Passport 7400, 15000, 20000 ATM Traffic Management Fundamentals* for more detailed load spreading and load sharing information.

LoadSpreadFast

When the forwardingPolicy attribute is set to LoadSpreadFast, traffic is forwarded in an order-preserved manner along a randomly selected route. The traffic is routed over one particular link within a link group, and will not spill over to other links. Choosing this value may improve throughput, but traffic may be discarded in the presence of congestion.

Increased Multipath using Variance for DPRS

The Increased Multipath (Variance) feature improves traffic spreading across the Passport backbone by providing more balanced link usage across the network through the safe use of unequal-metric paths. In addition to the better spreading of traffic, which helps avoid congestion situations, Variance provides more possible paths for high-reliability traffic overflow when congestion does occur.

Increasing multipaths

In the default behavior of the Passport routing system, every path used to route packets to a destination must have a minimum metric. So, if two paths are to be used in multipath mode, they must have equal minimum metrics. If two such paths exist, traffic is evenly spread across the two paths, but if only one minimum metric path is found, all traffic to the destination is sent along that one path. In many networks, this requirement for equal minimum metrics results in few multipaths.

The Variance feature increases the frequency of multipaths to a destination by allowing the metrics of the two paths to differ by some provisioned amount. Variance is provisioned on a nodal basis, and is applied to path calculations for all RID and MID destinations.

Variance values can be provisioned for both the delay and throughput RCOS. The variance value for an RCOS is defined as the maximum allowable difference in metric between the best path and a second path. The value is specified as a percentage of the best path metric.

When Variance is active, the second path chosen is the lowest metric path that meets the following conditions:

- The metric of the second path must be within the provisioned range of the best path metric.
- The second path must be loop-free (that is, must not allow traffic to return to a switch previously traversed by the packet).

If no possible alternative path meets both these conditions, only the single minimum-metric path is used for routing to the destination.

Traffic distribution

When Variance is in effect, traffic flows (VCs) are statistically spread across the two paths in inverse proportion to their metrics, with a resolution of 6.25% of the number of traffic flows.

For example, if two paths have throughput metrics of 1000000 and 1500000, the first path should receive 3/5 of the throughput traffic flows. The resolution of 6.25% results in 62.5% of the traffic flows being sent along the better path. (This may not correspond to 62.5% of the actual total volume of traffic, since the traffic flows (VCs) may be of different sizes.)

Similarly, delay traffic is divided in inverse proportion to the delay metrics, allowing the majority of the traffic to be sent on a path with less bandwidth, if the delay for that path is better.

For more information on Variance, see 241-7401-110 *Passport 7400, DPN-100 Interworking Guide*.

DPRS call establishment

Before users of an access service on a Passport subnet can exchange data packets, a virtual circuit must be set up. Setting up the virtual circuit is referred to as call establishment.

Although the virtual circuit system and the routing system are necessary for call establishment, a third system, the call routing system, is also required. This system is supported by call routers.

Two types of call routers exist: those that run locally on Passport and those on Call Server Resource Modules (CSRMs) (only on Passport 7400 series switches). In Passport-only networks, call routing on each Passport node is sufficient to route the calls.

In networks with DPN-100 access modules, CSRMs (only on Passport 7400 series switches) perform call routing services.

The DPRS call routing system

The access service initiates a call by routing a call request packet to the call routing system. The call routing system determines the RID/MID location of the called address (that is, E.164 or X.121, the data network address [DNA],

an address assigned to each link in the network). The call routing system then forwards the call request packet to the called access service at that location. The called access service returns a call accept packet to the caller to establish the virtual circuit connection.

At the network level, the call routing system can be viewed as a network-wide database that translates destination DNAs to the RID, MID, and MPID addresses required by the forwarding system.

Example of call establishment

The following sequence describes call establishment by call routing in a DPRS configuration when using a CSRM (see the figure "Call routing and DPRS—call establishment" (page 105)). CSRM is only available on Passport 7400 series switches. Passport 15000 and 20000 do not use CSRM.

1 The call request packet includes the data network address (DNA), RID (15), MID (3) and MPID for the source (A) and RID 0, source call router LMID, and the DNA for the destination (B). The call request packet is forwarded from the source to the closest Call Server Resource Module (CSRM) to determine the destination RID.

   Note: By default, call request packets are forwarded to the closest CSRM, however you can provision the CSRM routing on a node as shared. CSRM loadsharing allows the two CSRMs with the lowest RID values to evenly share the call request packets. For more information, see 241-7401-110 *Passport 7400, DPN-100 Interworking Guide*.

2 To determine the destination address, the source call router of the CSRM (only on Passport 7400 series switches) maps the DNA to the destination's RID (16). The CSRM (only on Passport 7400 series switches) and Passport nodes in the source RID (15) use their forwarding tables to send the call packet along the best path to the destination RID (16).

3 The call request is forwarded to the closest CSRM (only on Passport 7400 series switches) in the destination RID (16) to determine the destination MID.

4 The call request is forwarded to the destination MID (7).

5 At destination B, the call is sent to the called DNA.

6 A call accept is sent back to the source, carrying with it the RID (16), MID (7) and MPID of the destination, by the shortest available route.

7 Data transfer can begin because both ends know each others RID, MID and MPID addresses.

Figure 24
Call routing and DPRS—call establishment

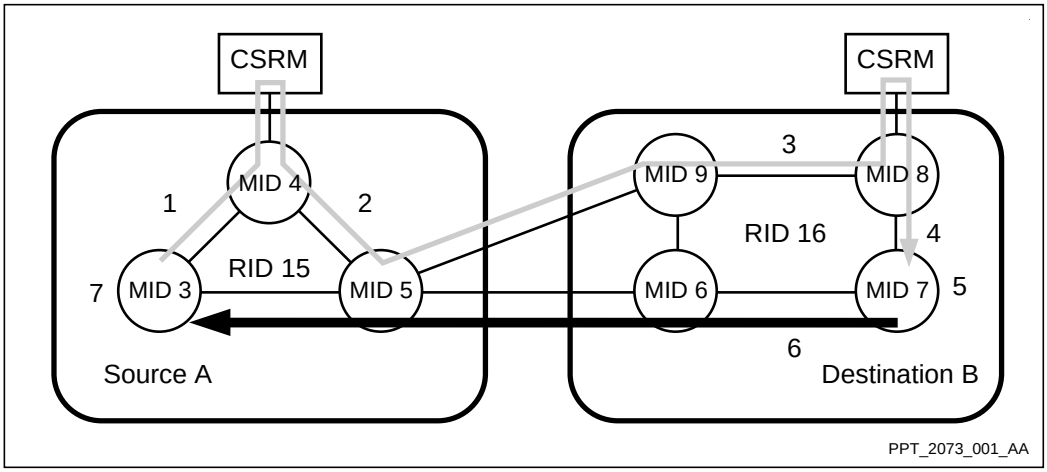

For more information on DPRS call routing, see “DPRS call routing” (page 194).

For a DPRS call establishment example highlighting addressing, see “Passport addressing and call routing” (page 183), “DPRS call routing” (page 194).

DPRS data transfer

After the virtual circuit is established, users of an access service can transfer data packets on a Passport subnet, for example, frame relay. The packet forwarding function, invoked on each node, routes packets. It uses the DPRS header information and the forwarding table information to forward a packet to its destination.

Routing of a packet through a Passport network is performed through a network entrance-point (ingress) Passport node, any additional (tandem) Passport nodes on the path, and an exit-point (egress) Passport node.

When a packet is received from a link on any node in the path to a destination node, the packet forwarding function examines the packet's headers. The RCOS criteria is used to select the appropriate routing table. The forwarding function extracts the destination address from the packet (RID/MID) and uses the routing table to determine the appropriate next hop link for the packet. The packet is then sent to the link. This process is repeated at every node (hop) in the network until the packet reaches its destination address. At the destination, the packet is forwarded to the MPID (from the packet) of the destination service.

Examples of DPRS data transfer

The figure "DPRS data transfer" (page 107) shows the data transfer process for a network with CSRMs.

> ***Note:*** The DNA is only used during the call setup. Once the data path is established, the call is routed via the RID, MID, and MPID address information.

Figure 25
DPRS data transfer

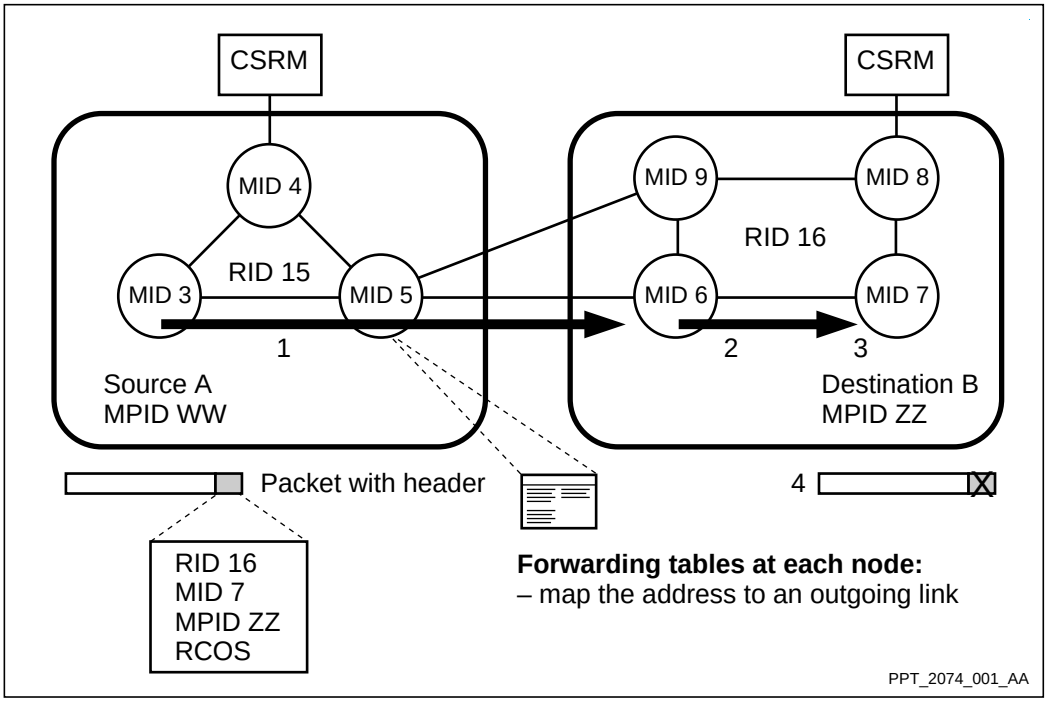

The transfer of data between Source A and Destination B involves the following steps:

1 DPRS packet forwarding examines the packet at the RID of destination B. If the RID does not match the one supported by source A, the routing system's RID forwarding tables are used to determine how to forward the packet towards the destination RID.

2 When the RID is reached, the MID field is checked. If the MID is not that of the destination module, the routing system's MID forwarding tables are used to determine how to forward the packet towards the destination MID.

3 If the MID matches that of the module, the MPID field is checked to determine which process on the module will forward the packet to destination B.

4 The packet is then decapsulated (headers removed) and passed up to the corresponding higher layer protocol.

DPRS route failure

Under normal operating conditions with no congestion and no failures, all packets of a call follow the same path. During congestion, or failures of hardware elements, packets are either rerouted or discarded. Rerouting is done simply by updating the forwarding tables on some, or all, nodes in the path to the destination with new next hop link group information. The routing system calculates new paths.

The decision to reroute or discard a packet depends on the information contained in the packet header and the availability of alternate network resources.

Rerouted packets may not arrive in the order they were sent. The DPRS virtual circuit layer will reorder the packets before passing them to the application.

DPRS components

Each Passport node has a single *Routing* component that maintains a current view of the node and network topology. For DPRS, the *Routing* component has subcomponents *Dpn* and *Topology.*

Routing Dpn is the component residing on all Passport nodes that carry DPRS or DPN-100 traffic. *Routing Dpn* maintains the routing information required to route this traffic. The *Routing Dpn* component also supports provisioning of the load sharing or load spreading forwarding policy. It also enables users to find the RID, MID, and LMID paths in use as well as the forwarding statistics.

For more information on components, see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Chapter 4
Multiprotocol label switching

This chapter describes multiprotocol label switching (MPLS), and includes the following information:

For more detailed information on MPLS, see 241-5701-445 *Passport 7400, 15000, 20000 Multiprotocol Label Switching Guide*.

Role of MPLS

MPLS is a label-swapping, networking technology that forwards packet traffic over multiple, underlying layer-2 media. This technology integrates layer-2 switching and layer-3 routing by linking the layer-2 infrastructure with layer-3 routing characteristics. Layer-3 routing occurs at the edge of the network, and layer-2 switching takes over in the MPLS network core. (See figure "MPLS technology" (page 110).)

Essentially, MPLS forwards a packet by swapping labels at each node in its path. MPLS makes it possible to create new label formats without having to change routing protocols. For example, MPLS traffic can include internet protocol (IP), frame relay, ATM, and even optical waveforms. In Passport networks, MPLS transports IP traffic over ATM infrastructure, allowing carriers and large enterprises to send IP data easily across the ATM backbone.

Figure 26
MPLS technology

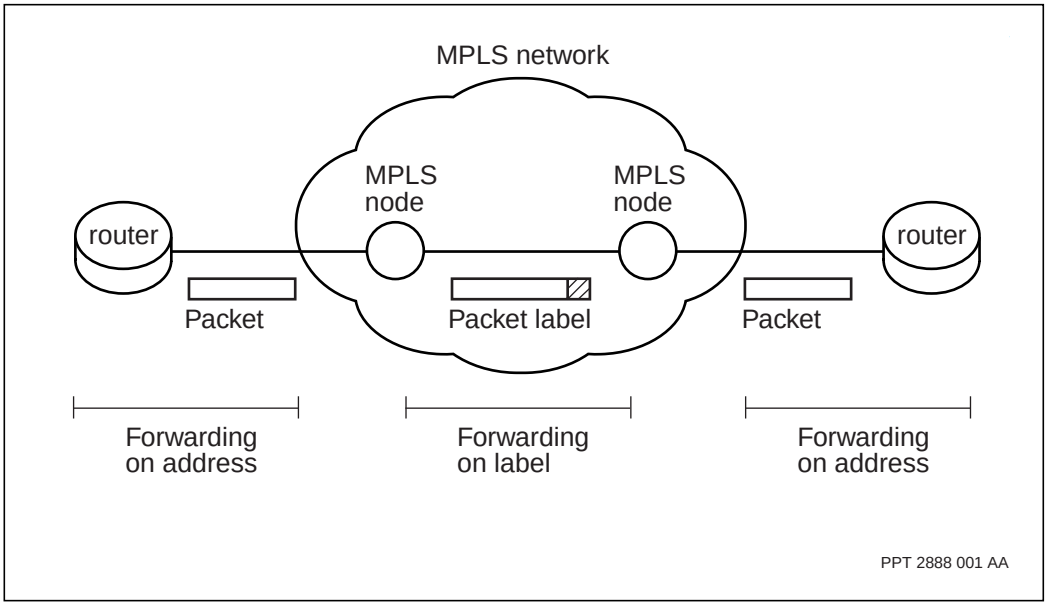

MPLS features

Carrier organizations and large enterprises typically use MPLS in their backbone networks to improve network resource usage. As the key to the future of large-scale IP networks, MPLS provides many benefits:

- independence of function—in MPLS, the forwarding plane is separated from the routing protocol control plane, so that the MPLS core performs a simple forwarding function completely independently of the packet content; this practice allows policy and routing decisions to be applied only once at the network edge
- traffic engineering—MPLS channels the operation of IP routing so that traffic can be steered to achieve efficient network resource usage and optimal performance. The grouping of LSPs into LSP groups, allows for multiple LSPs, of different quality of service, to the same FEC destination.

- resource control—MPLS allows you to control valuable resources precisely, for example, through the definition of different classes of service
- network evolution—MPLS is developing into a robust network in which a single, unified protocol operates over multiple, underlying layer-2 technologies
- standards support—MPLS is an emerging standard for network-layer packet forwarding, based on a number of signaling protocols proposed by the Internet Engineering Task Force (IETF) that are used to distribute labels and forward MPLS traffic; among these protocols are the label distribution protocol (LDP) and constraint-based routing using LDP (CR-LDP)

MPLS architecture

This section describes MPLS architecture. The topics include:

- "MPLS in the Passport networking architecture" (page 111)
- "Functional elements" (page 111)

MPLS in the Passport networking architecture

Figure "Passport networking architecture" (page 27) illustrates where MPLS and its related systems fit in to the Passport networking architecture layers. MPLS supports IP services, but does not use the Passport base routing and trunking capabilities. Instead, MPLS has its own routing and trunking capabilities, and relies on the ATM cell switching infrastructure.

Functional elements

MPLS is a forwarding mechanism that works by applying a label to IP traffic entering the network. The label acts as a shorthand representation of the IP packet header. As the traffic moves through the network, MPLS swaps the label at each node on the route, according to a pre-defined label database at that node. At the egress side of the MPLS network, the packet is decapsulated, and continues under the IP routing protocol.

Figure "MPLS network" (page 112) shows an MPLS network with a number of Passport nodes. The nodes at the edges of the network are called label edge routers (LER). The LER nodes provide ingress and egress functions for IP traffic in the MPLS network. The core nodes are called label switched routers

(LSR). The LSR nodes provide the high-speed switching functions for the network. The path of data between the MPLS nodes is called a label switched path (LSP). An LSP is a unidirectional tunnel through the network.

Figure 27
MPLS network

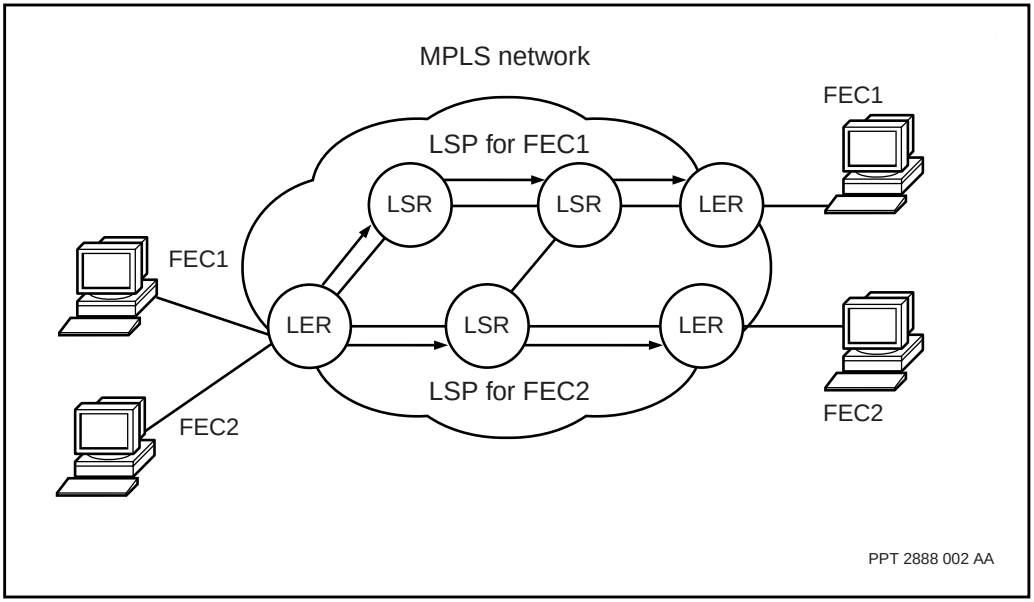

When IP traffic arrives at an LER, MPLS applies the label for the first time. To do this, the LER analyzes the information in the IP packet header, and classifies traffic according to its destination and class of service characteristics.

At the LER, MPLS uses the concept of a forwarding equivalence class (FEC) to map incoming traffic to an LSP. Essentially, a FEC defines a group of packets that are forwarded over the same path with the same forwarding treatment. This means that all the packets with the same FEC can be mapped to the same label.

For each FEC, the LER sets up an LSP through the network to the destination defined by the FEC. After the traffic is assigned a FEC, the LER applies a label based on the label information base (LIB). The LIB maps each FEC to an LSP label that defines the next-hop link. Because the underlying layer-2 media is ATM, MPLS uses the VCI of the ATM VCC as the label.

To forward the packet, the LER looks up the FEC in the LIB, and then encapsulates the packet with the LSP label. The LER then sends the packet out on the next-hop interface defined in the LIB.

When a labeled packet arrives at an LSR, the LSR extracts the incoming label and uses it as an index into the LIB. When the LSR finds the appropriate LIB entry, it extracts the corresponding outgoing label and swaps it with the incoming label in the packet. The LSR then sends the packet on the outgoing interface to the appropriate next hop specified in the LIB entry.

Eventually, the packet reaches the edge of the MPLS network. At that point, an LER removes the encapsulating label, and the packet continues to its destination according to conventional IP routing methods.

MPLS routing mechanisms

For information on MPLS routing mechanisms, see the following sections:

- “Hop-by-hop LSPs” (page 113)
- “Explicit routes” (page 115)
- “MPLS protocols” (page 116)
- “MPLS using ATM media” (page 117)

Hop-by-hop LSPs

The basic LSP is called a hop-by-hop LSP. A hop-by-hop LSP is part of a tree from every source to a particular destination. (See figure “Hop-by-hop LSP” (page 114)). For these LSPs, MPLS builds a set of trees that duplicate the destination-based trees that IP uses to forward traffic. MPLS converts the destination-based trees into label-switching trees.

At each node, MPLS creates the tree by allocating a label for every next-hop MPLS destination and exchanging the labels with those peers. The exchange is accomplished through LDP request and mapping messages. For example, in figure "Hop-by-hop LSP" (page 114), LSR G maps incoming labels 14, 23, and 24 to outgoing label 30.

With IP routing, each router along a path examines a packet's destination and chooses a new link. With MPLS, the packet follows the same path it would take with IP routing, but the packet is assigned a label and link at the ingress LER. When the packet arrives at the next hop, its label is replaced with the next label along the tree toward the destination and sent to the corresponding link. In this way, the packet follows the IP routing path, but its IP header is never checked along the route.

Note: In the diagrams in this section, the node labelling (A, B, C) represents the actual IP addresses used by the software.

Figure 28
Hop-by-hop LSP

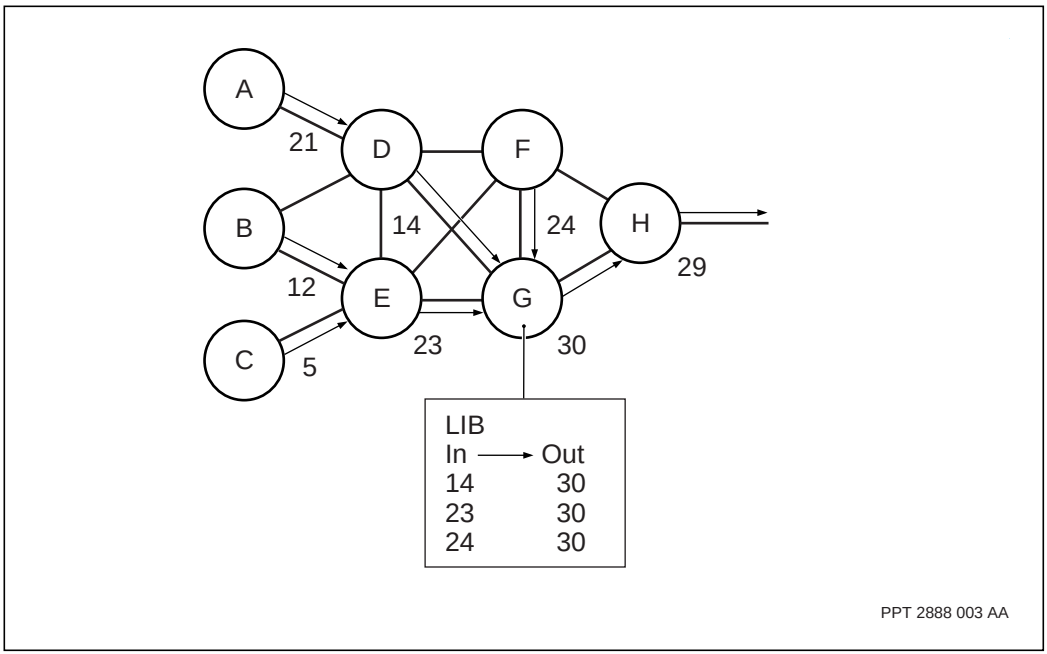

Explicit routes

One of the major advantages of MPLS is its ability to steer traffic, for example to avoid congestion or to meet the QoS of the traffic. MPLS allows the network operator at the source node to determine an explicit route LSP (ER-LSP) that defines the path the traffic will take. Multiple LSPs, with different quality of service, can be configured to the same destination.

Unlike the hop-by-hop LSP, the ER-LSP does not have to follow the IP tree. Instead, the ER-LSP builds a path from source to destination. (See figure "ER-LSP" (page 115).) To build this path, MPLS embeds the explicit route into the label request message using the protocol called constraint-based routing using LDP (CR-LDP).

Figure 29
ER-LSP

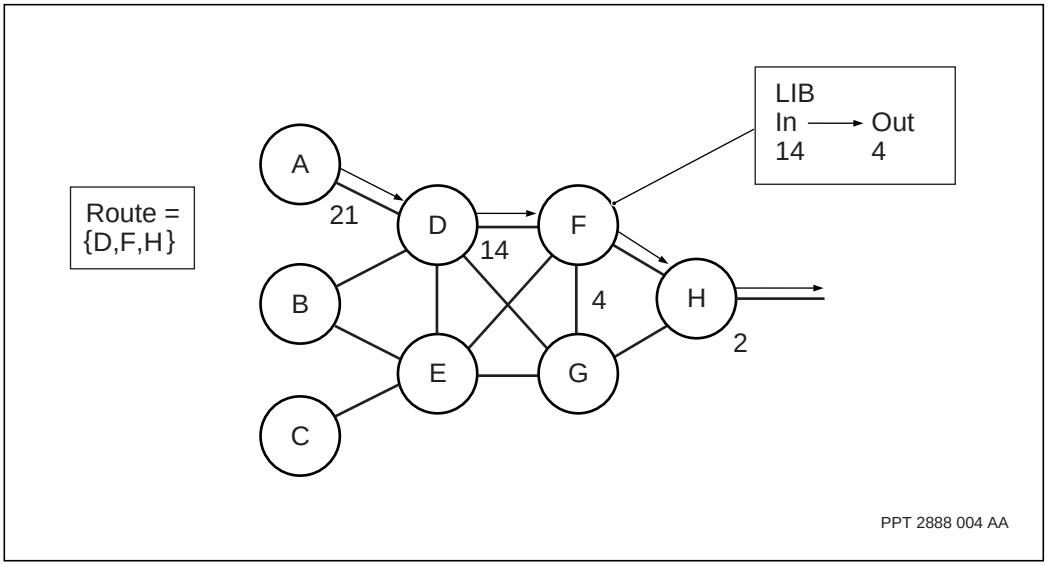

There are two types of ER-LSPs: strict ERs and loose ERs. A strict ER specifies the exact path a packet will take. MPLS at the source node explicitly indicates all the hops along the path between the end points. A loose ER specifies some, but not all, of the hops a packet must traverse on the way to its destination.

Figure “Strict and loose ER LSPs” (page 116) shows the difference between strict and loose ERs. The strict ER in the illustration is specified at LER A as {LSR B, LSR D, LSR E}. In establishing a strict ER, MPLS does not need to check the IP routing tables, since the exact route is defined.

The loose ER in figure “Strict and loose ER LSPs” (page 116) is specified at LER A as {LSR E}. In the illustration, the complete path could be either {LER A, LSR B, LSR D, LER E} or {LER A, LSR C, LSR D, LER E}. In the loose segment between LER A and LSR E, MPLS checks the IP routing tables during call setup to determine the best next hop to the next specified ER hop in the route.

In Passport MPLS, loose ERs are pinned ERs. This means that, after the route is set up, it does not change unless failure occurs, even if the IP routing tables change.

Figure 30
Strict and loose ER LSPs

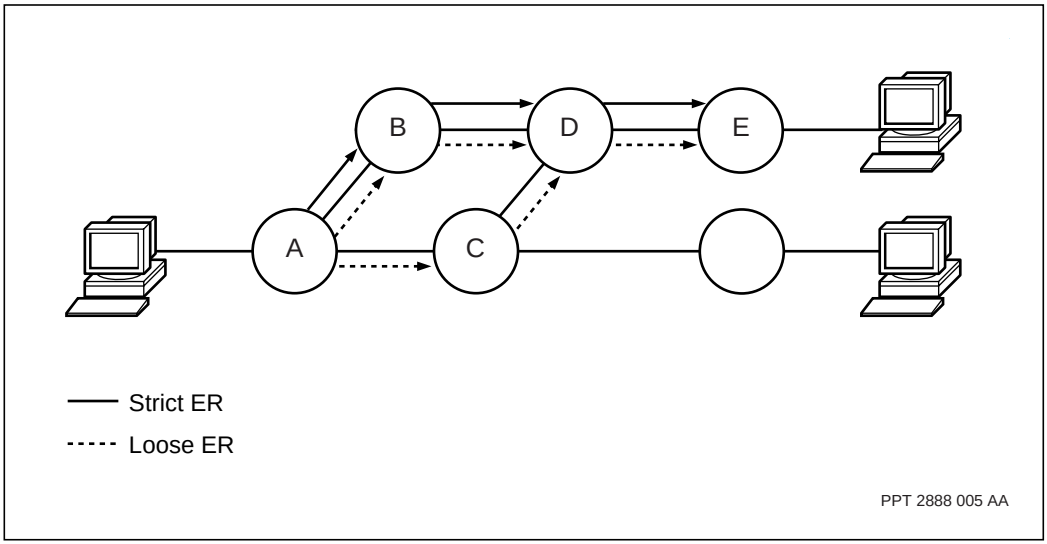

MPLS protocols

The MPLS protocols are used to set up, maintain, and tear down LSPs. The protocols allow LSP establishment by mapping network-layer routing information to data link-layer switched paths.

The basic signalling protocol in MPLS is the label distribution protocol (LDP). As its name indicates, LDP provides the method of distributing label binding information between MPLS nodes, or peers.

Constraint-based routing using LDP (CR-LDP) is an extension of the LDP protocol. CR-LDP allows forwarding on the basis of constraints such as explicit routes or traffic parameters. CR-LDP permits the traffic engineering that is essential to the management and predictability of large networks.

MPLS using ATM media

The supporting layer-2 medium for Passport MPLS is ATM. Passport MPLS interacts with ATM in the operating mode called ships-in-the-night. In this mode, MPLS and ATM control planes share hardware ports, but work independently.

The ships-in-the-night mode of operation allows a Passport node to function simultaneously as an ATM switch and an MPLS LSR. The MPLS control plane provides IP-based services, and the ATM control plane provides ATM-based services. The two control planes share memory, VPI/VCI space, processing capacity, and traffic management capabilities.

Chapter 5
Path-Oriented Routing System (PORS)

This section describes the Passport Path-Oriented Routing System (PORS) and includes the following topics:

For more detailed information on PORS and PORS services (BTDS, HTDS, Voice Transport, and AAL1 Circuit Emulation Service), see "PORS references" (page 19).

Role of PORS

PORS provides connection-oriented routing for a wide variety of traffic, including delay-variance-sensitive and bandwidth guaranteed traffic such as voice, video, facsimile, modem, bit transparent data, and constant bit rate data.

PORS dynamically routes calls across multiple topology regions while maintaining the majority of its unique capabilities such as bumping, rerouting, and dynamic path attribute (QoS) adjustment.

Trunks that span these region boundaries are referred to as inter-region Passport trunks. Calls may be comprised of one or more paths across a region or inter-region Passport trunk.

PORS routing features

This section describes the PORS routing features.

Bandwidth reservation

PORS enables the user to reserve bandwidth on a Passport trunk, based on peak, average, or any other traffic-use requirement.

> ***Note:*** DPRS and PORS share the bandwidth on a Passport trunk. The instantaneous unused bandwidth is available to these routing systems.

Call priority and bumping capability

PORS enables independent setup and holding priority for connections that can be used to prioritize important connections.

Packet efficiency

PORS headers are small. This small header and large data-to-header size ratio gives PORS a high throughput.

Fast packet forwarding

The transmission path is determined when the connection is set up and remains fixed for the duration of the connection (barring network changes). Each path is set up with logical channels on the Passport nodes of the path. Since some of the routing information is associated with the path, each packet has a simple routing algorithm. So, on a Passport network, PORS can take full advantage of frame and ATM-based links.

Packet order preservation

PORS does not retransmit packets. All packets follow the same path. Even if network changes cause some packets to be lost, the remaining packets arrive at the destination in the same sequence as they were transmitted.

Best effort delivery

PORS attempts to maintain uninterrupted service with automatic rerouting around any network changes that can otherwise cause packets to be lost.

Control over traffic pattern and route selection

PORS provides quality of service parameters such as cost metric, security levels, and emission priorities. With these parameters, the user can direct traffic over preferred routes.

Path optimization on a periodic basis

PORS can move an existing path to reduce the cost or delay metric, or to balance the load across link groups.

> ***Note:*** Path optimization and balancing are only available within a region path segment and not across inter-region Passport trunks.

PORS segmented optimization

PORS can move an existing path segment of a segmented PORS connection in order to reach a gateway or destination; this reduces the gateway cost or the path cost.

Congestion notification

During periods of traffic congestion, PORS notifies the service end points of this change in the network.

Interworking with ATM

PORS supports hardware forwarding with ATM services. Within a Passport network, PORS can dynamically establish VCCs and use hardware forwarding at ATM line speeds. This enables full utilization of high-speed ATM interfaces even when using small voice packets. It also allows routing of ATM services, such as AAL1 Circuit Emulation Service (CES) over PORS networks.

Interworking with MCS

PORS routing facilities are used to transmit frame relay service traffic in MCS connections. Each MCS connection is a PORS permanent logical connection (PLC) that benefits from the PORS value-added features, such as bandwidth reservation and path optimization.

Switched calls for voice service

PORS supports switched virtual circuits (SVCs) used by voice services.

Dynamic trunk speed change

The dynamic trunk speed change feature enables Passport 7400, 15000, 20000 trunking and routing to adapt to bandwidth changes without taking trunks out of service. For details, see 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide*.

Bandwidth load spreading

The intention of load spreading is to avoid link congestion caused by several inter-region Passport trunks originating from a single-node. Bandwidth load spreading is an option on inter-region Passport trunks that are provisioned with identical characteristics by the network operator. For details on bandwidth load spreading, see 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide*.

PORS architecture

This section describes the PORS architecture. The topics include the following supporting architecture layers:

- “Passport trunking system for PORS” (page 122)
- “PORS and base routing components” (page 123)
- “PORS virtual circuit (VC)” (page 125)
- “PORS logical channel (LCh)” (page 127)
- “PORS services” (page 128)
- “PORS traffic management” (page 128)

Figure “Passport networking architecture” (page 27) illustrates where PORS and supporting systems fit in to the networking architecture layers.

Passport trunking system for PORS

The Passport trunking system supports PORS by providing link information such as availability, speed (bandwidth), PORS reservable bandwidth, and remote node identification. PORS uses this link information for packet forwarding across Passport nodes and links.

For PORS to be active on a Passport trunk, the optional Passport 7400, 15000, 20000 trunk path administrator (PA) needs to be provisioned on the Passport trunk. The PA provides the infrastructure for PORS routing on a Passport trunk and can reside on physical or logical Passport trunks.

PORS uses the following link metrics:

- cost metrics based on a user definable cost
- delay metrics based on the link in the group with minimum measured delay

PORS uses both 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. For more information, see 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide* PORS efficiency chapter.

Each Passport node can support up to 128 PORS trunks.

For more information on Passport trunking, see “Passport trunking and base routing systems” (page 49).

PORS and base routing components

PORS packet forwarding across Passport nodes and links uses information from the PORS and base routing systems (see the figure “PORS and base routing architecture” (page 124)).

Figure 31
PORS and base routing architecture

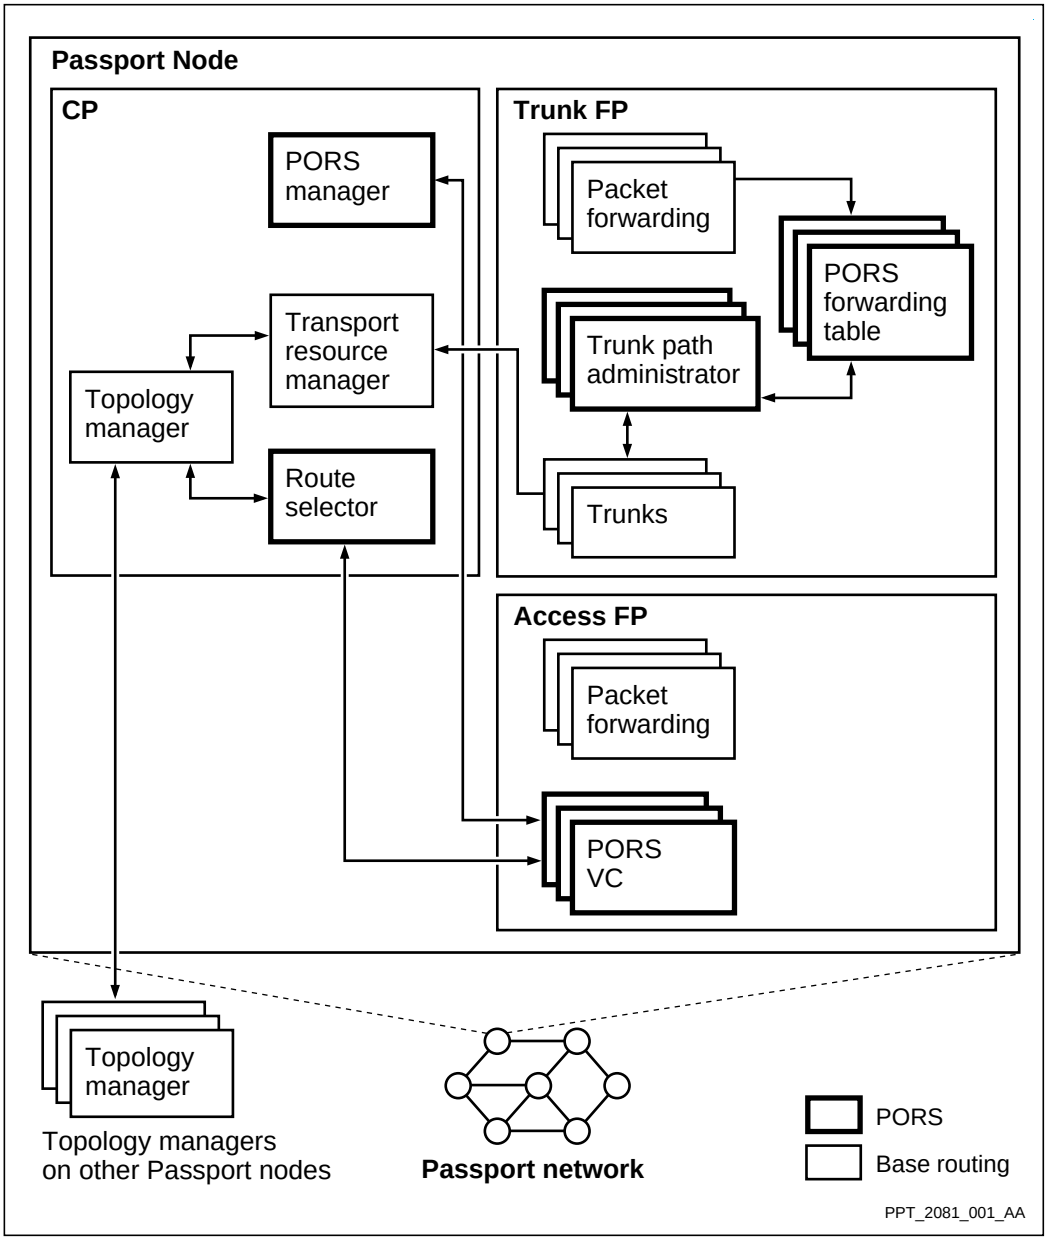

Transport resource manager (TRM) supporting PORS

Each link (Passport trunk) on a Passport node registers with the TRM, providing neighbor information, throughput, round trip delay, link status (up or down), and percent of bandwidth available to PORS.

Topology manager supporting PORS

PORS uses the Passport topology system, an open shortest path first (OSPF)-based link state routing protocol. This protocol is used to exchange topology information with topology managers on other Passport nodes and to create a topology region-wide view of the Passport topology.

PORS packet forwarding

The packet forwarding function is invoked when a packet arrives from a Passport trunk. This function looks up the path's logical channel number in the PORS forwarding table to determine the next hop (Passport trunk or path end point). Packet forwarding may discard packets based on the congestion level and the path's discard priority. Packet forwarding insures the packet is put on the queue indicated by the path's emission priority.

PORS manager

The PORS manager controls optimization for the module.

Passport 7400, 15000, 20000 trunk path administrator (PA) supporting PORS

The *Trunk PathAdmin* component is responsible for the following functions:

- establishing and maintaining logical channels
- managing bandwidth and bumping PORS connections

See "Passport trunking and base routing systems" (page 49) for more details.

PORS virtual circuit (VC)

A virtual circuit (VC) uses network resources such as the routing and Passport trunking systems to establish, maintain, and terminate logical connections across the Passport subnet for an application service. The VC provides reliable and ordered data delivery to PORS.

VC functions are described in "VC functions" (page 92).

PORS VC connections

The following two mechanisms for establishing connections are used:

- permanent logical connections (PLCs): also referred to as permanent virtual circuits (PVCs). A PLC is a permanent connection between two end points. Typically, a PLC is used to transparently carry access service protocols. An access service, such as voice, video, and HTDS, uses a PLC to exchange data with another access service. PLCs are meant to be permanent and (barring network changes) are broken only by the network administrator. A PLC provides an order-preserving, error-free, best-effort delivery path. See the figure "PLCs, LCs, and LCNs" (page 126) for more details on PLCs, LCs, and LCNs.

  PLCs also support a dynamic operator-controlled connection where an operator can set an override remote name on demand. This can be used for temporary bandwidth allocation such as video conferencing.

- switched logical connections (SLCs): also referred to as switched virtual circuits (SVCs). SLCs are similar to PLCs except that SLCs are set up only for the duration of a call. These connections are established through user signaling and can therefore be used to access a variety of destinations. The network establishes the connection on demand for the ATM services.

Figure 32
PLCs, LCs, and LCNs

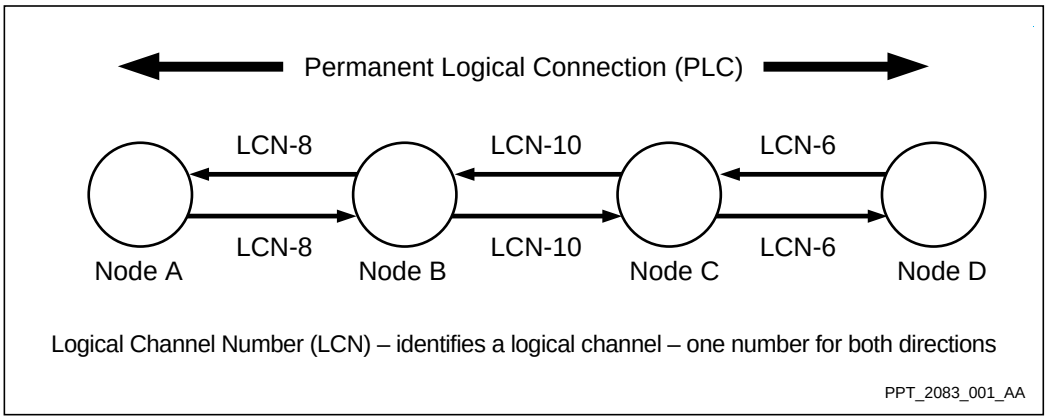

PORS VC types

PORS uses the PORS VC that has best effort delivery. See the table "PORS virtual circuits" (page 127) for details on the PORS VC characteristics.

Table 6
PORS virtual circuits

Characteristics	PORS
Guaranteed delivery of packets	No
Ordering of packets	Yes
Elimination of duplicate packets	N/A
Segmentation and reassembly	N/A
Flow control of packets	No
Congestion notification	Yes
Adaptation to network changes	Yes

Note: Because connection-oriented services cannot misorder or duplicate packets, the VC does not have to deal with this issue.

PORS logical channel (LCh)

A logical channel (LCh) is the instance of a path on a Passport trunk. LCh components exist only on the receiving side of a Passport trunk, as illustrated in the figure "Logical channel component for one Passport trunk" (page 128). The LCh components at each end of the Passport trunk use the same logical channel number (LCN) to the path. The LCh component owns the attributes of one path on that Passport trunk.

Figure 33
Logical channel component for one Passport trunk

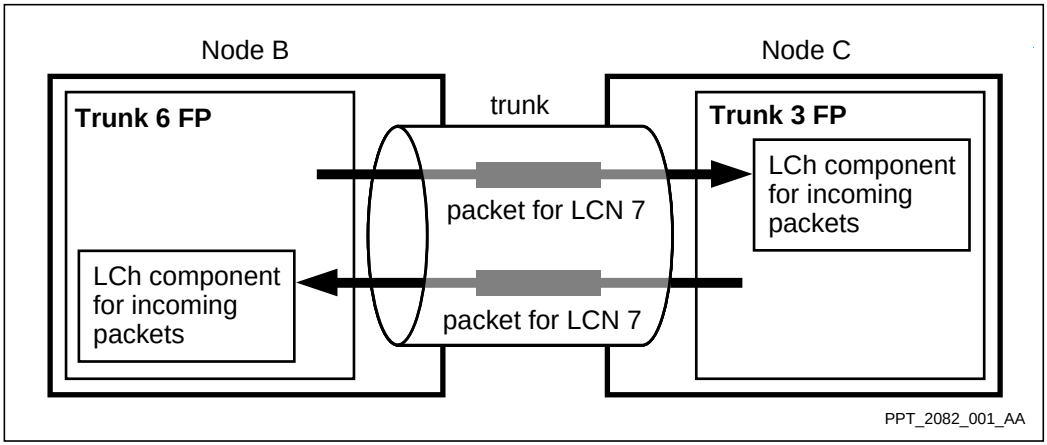

PORS services

PORS is the Passport routing system used to transport data traffic for connection-oriented services. The services are as follows:

- Bit Transparent Data Service (BTDS)
- HDLC Transparent Data Service (HTDS)
- Voice Transport
- AAL1 Circuit Emulation Service (CES)

Note: The services that use PORS also require network clock synchronization (NCS) to ensure the accurate transmission and reproduction of synchronous data such as voice or video.

PORS traffic management

PORS traffic management mechanisms operate primarily in the area of route selection. A path is expected to follow the same route for the lifetime of the logical connection. However, a new route can be established for the following reasons:

- “Path optimization” (page 129)
- “Path failure” (page 129)

Path optimization

An existing path can be moved to a more optimal path if the new path minimizes the original path's provisioned metric (cost or delay).

Path failure

An existing path can be interrupted if the network changes. For example, nodes can be removed or added during a call. PORS reroutes any failed paths if possible.

Setup priority, holding priority, and bumping

Each PORS connection has a separate setup and holding priority. This feature can be used to determine which calls remain active when bandwidth is reduced or paths are bumped.

For more information on traffic management, see the following documents:

- 241-5701-705 *Passport 7400, 15000, 20000 ATM Traffic Management Fundamentals*
- 241-5701-706 *Passport 7400, 15000, 20000 ATM Traffic Shaping and Policing*
- 241-5701-707 *Passport 7400, 15000, 20000 ATM Queuing and Scheduling*
- 241-5701-708 *Passport 7400, 15000, 20000 ATM CAC and Bandwidth Management*

PORS routing mechanisms

This section includes the following information:

PORS addressing system

Addresses uniquely identify network entities such as Passport nodes, telephones, and service ports. Addresses are stored in the header of each call packet and are used by the packet forwarding system to route the packet to its destination. The figure “PORS addressing system” (page 131) illustrates the PORS addressing system. A PORS address is defined by the following:

- “Provisioned name” (page 130)
- “Physical address” (page 130)

Provisioned name

Each service has a node name and a component name.

PORS uses the component name as the address of the source and destination end points on the network. The internal address plan identifier of a node and the component identifying the access service within the node are used as the address.

As an example, if a node has “EM/Ottawa” as the node name (EM for Passport enterprise module), and a voice service (VS) access port is identified as component name “VS/5”, the complete name of the end point is “EM/Ottawa/VS/5”.

Physical address

The node name maps to the nodeId and the component name maps to the processId. Alternatively, if the service uses DNA or NSAP addressing, PORS converts the address to a nodeId. PORS uses the nodeId to get a route.

NodeId values range from 1 to 4095. Nortel Networks engineering guidelines support up to 1000 nodeIds in a topology region.

Addressing is used for call setup only. After a connection is established, packet forwarding uses table lookups to route packets.

For more information on addressing, see “Passport addressing and call routing” (page 183).

Figure 34
PORS addressing system

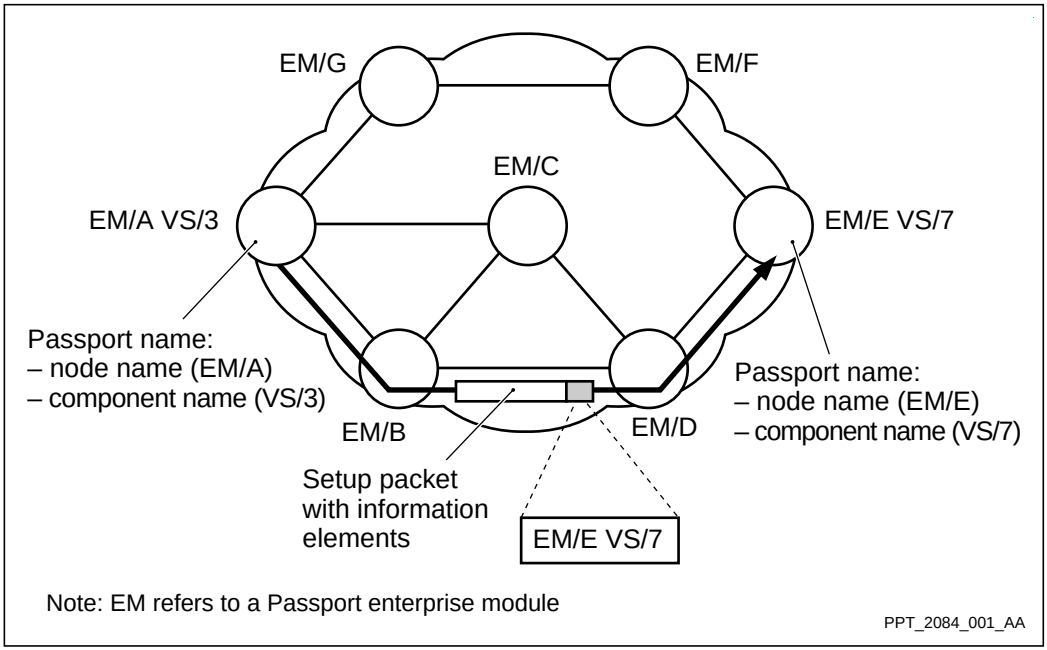

Physical network partitioning

A topology region is a group of interconnected Passport nodes that are partitioned as an autonomous section of the network. The nodes in a topology region know the topology of their own region only. Dynamic routing of PORS services that use NSAP addressing (CES, MCS, VTDS) is possible within and between topology regions, because inter-region Passport trunks for a given topology region, are provisioned with all of the NSAP addresses that they can reach in other topology regions. This allows nodes to identify which border node they can route to, in order to reach a node in another topology region. PORS services that do not use NSAP addressing can be dynamically routed between nodes within a topology region only. If you have VTDS services that will only be routing within a topology region, you can leave them under the current addressing scheme (do not need to migrate to NSAP addressing). For all services, calls between topology regions can be performed by establishing a manual path. For more details on routing within and between topology regions, see 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide*.

Passport clusters

Topology regions can be further partitioned with Passport clusters. Passport clusters are groups of Passport access switches, connected to a region's backbone. Passport clusters exchange only limited routing information with the backbone, and exchange no topology information with the backbone or other Passport clusters. Passport clusters are not included in the backbone's topology view. This results in an improvement in the engineering limitation on the total number of Passport nodes in a topology region. As nodeID reuse is not supported, the total number of Passport cluster nodes and backbone nodes that can be deployed in a topology region continues to be limited by existing design limits.

As with routing across topology regions, routing between a cluster and the backbone of a topology region is supported for PORS services (CES and VTDS using NSAP addressing).

For more information on Passport clusters and PORS, see 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide*.

Manual path connection

Manual paths are defined at the end points of a call. One or two paths can be defined for each connection. The first path to come up is used to route the call. The alternate path is used to reroute the call if an interruption occurs on the first path.

CAS interfaces for NSAP information

With PORS services that use NSAP addressing, a component administrative system (CAS) interface displays the addressing broadcast by the base routing topology system. The CAS interface displays both the node NSAP addresses and reachable NSAP addresses stored in the Base Routing Topology Database, learned by the local node from its neighbors. This functionality assists network operators with provisioning and troubleshooting anomalies.

Refer to 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide* for details on CAS interface functionality with NSAP addressing.

PORS packet header

PORS is designed for use by services that can tolerate the loss of some packets. Engineering can change the probability of loss for any PORS connection. PORS provides end-to-end reliable delivery of its control traffic and best-effort delivery of its data payload. External to PORS, service applications can add end-to-end guaranteed delivery of data.

PORS packets have a common header and a PORS header followed by the actual data. See the figure “PORS packet format” (page 133) for more details on the PORS packet format.

Figure 35
PORS packet format

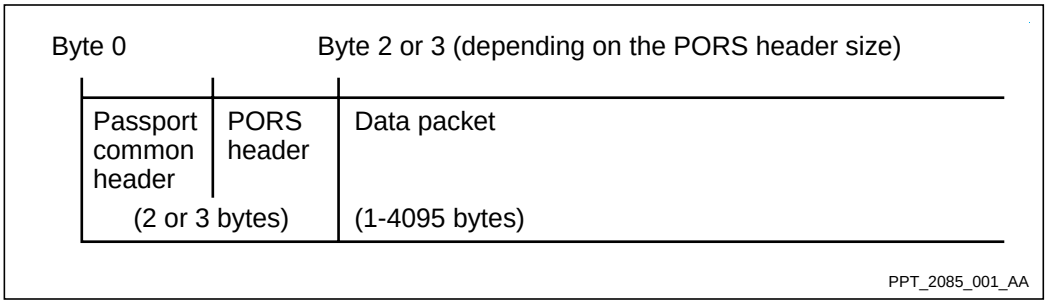

As a packet travels through the Passport subnet, the logical channel number (LCN) carried in the common header is used as an index into the forwarding table at each Passport node. The index determines which Passport trunk the arriving packet uses next.

Passport common header

A small Passport common header is appended to the front of frames and cells routed between Passport nodes. This header identifies that a PORS packet follows. The following header information is included:

- Routing mode identifies the packet type and its header format (for example, values of 0, 2, or 4 for PORS).

- Packet congestion indicators identify congestion encountered by the packet on its way to the destination node (forward congestion indication [FCI]). These indicators also identify congestion encountered by a packet issued by the source node (backward congestion indication [BCI]).
- Discard eligibility (DE) identifies the discard priority for the packet whenever discard decisions are made.
- Discard priority is compared to the link discard level in the forwarding tables. The discard priority determines the importance of forwarding the packet within the link group.

PORS header

The PORS header is a LCN that identifies the path to the next node. Routing information is stored at each hop of the route instead of being carried within the header. A three-byte header is used with ATM transport while a two-byte header is used unless the LCN exceeds 1023.

PORS call establishment

When a call is placed through the PORS, a path is created between the two Passport nodes located at the end points (source and destination) of the call.

The path is created as follows:

- A route that satisfies the transport characteristics is selected.
- The path and its associated switched logical connection (SLC) are instantiated.

After the path is created, the path is available for use by an access service, such as Voice Transport, Bit Transparent Data, and ATM services.

Selecting a PORS route

The PORS VC is the interface between the access service and the path. The PORS VC requests a route from the route selector (RS).

Several routes can exist between two Passport node end points. The RS selects the best available route from the topology database. See the figure "PORS and base routing architecture" (page 124) for more details. The RS

prunes the Passport trunks that do not meet the transport requirements. The RS determines the best route from among the remaining Passport trunks based on one of the following values:

- cost: assigned to Passport trunking facilities by a network administrator
- delay: calculated from the transmission delay of the Passport trunking facilities

Instantiating a PORS path

After the RS selects a route, the Passport node at the source end point builds or instantiates the path along the route. A call setup packet is sent along the Passport trunks that are required for the selected route, as illustrated in the figure "Instantiating a PORS route" (page 136).

The call setup packet is given to the Passport trunk PAs as it travels to the destination end point. The PAs add entries to the PORS forwarding tables on each Passport node along the route for the path being established.

For each Passport trunk of the selected route and for each direction of transmission, the PA:

- verifies that the required bandwidth is available on all Passport 7400, 15000, 20000 trunks
- reserves the required bandwidth
- creates the *LogicalChannel* (*LCh*) components
- assigns a logical channel number (LCN) to the path

The PA is responsible for all paths on a Passport 7400, 15000, 20000 trunk. The PA interfaces with the Passport 7400, 15000, 20000 trunk process on the function processor (FP) and reserves the bandwidth on the Passport 7400, 15000, 20000 trunk for the path. The PA then assigns the logical channel number (LCN) of the path on that Passport 7400, 15000, 20000 trunk and builds the forwarding table on the Passport 7400, 15000, 20000 trunk. The forwarding table is indexed by the LCh and indicates the next hop (Passport 7400, 15000, 20000 trunk or end point) of the route. The forwarding table includes path emission, discard priorities, and bandwidth requirements.

Figure 36
Instantiating a PORS route

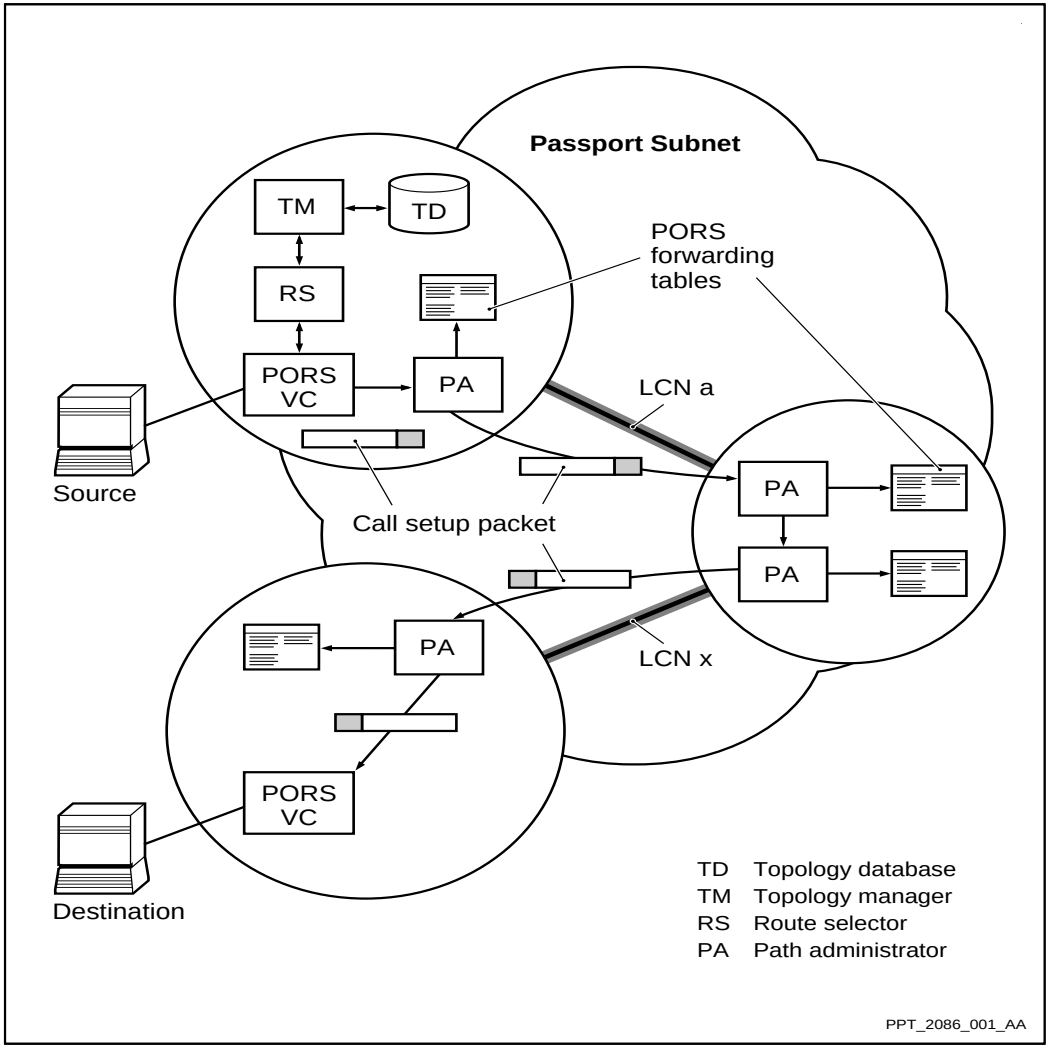

Path confirmation in the return direction

When the call setup packet reaches the Passport node at the destination end point, a call setup confirmation packet is returned to the Passport node at the source end point. Along the route, this packet gathers bandwidth for feedback to the topology system. When this packet arrives at the calling end, the path for data transfer is enabled.

Example of call establishment

The following sequence describes how a call is established by call routing in a PORS configuration (see the figure "Call routing and PORS — call establishment" (page 138)):

1 The source PORS VC requests a route from the route selector (RS).

2 The RS selects the best available route from the topology database, based on transport requirements, and cost and delay metrics.

3 The Passport 7400, 15000, 20000 trunk path administrator (PA):

4 reserves bandwidth on the Passport trunk that is chosen for the route

5 assigns a *LogicalChannel* (*LCh*) component and logical channel number (LCN) to the route

6 builds the forwarding table on the Passport trunk

7 The source sends a call setup packet along the Passport trunks that are required for the selected route.

8 At each Passport node along the route, the PAs update the forwarding tables.

9 The call setup packet reaches the destination.

10 The destination returns a call setup confirmation packet to the source. The call setup confirmation packet stops at all Passport nodes along the route.

11 The call setup confirmation packet reaches the source, and the path (a permanent or switched logical connection, depending on the application) is enabled.

Figure 37
Call routing and PORS — call establishment

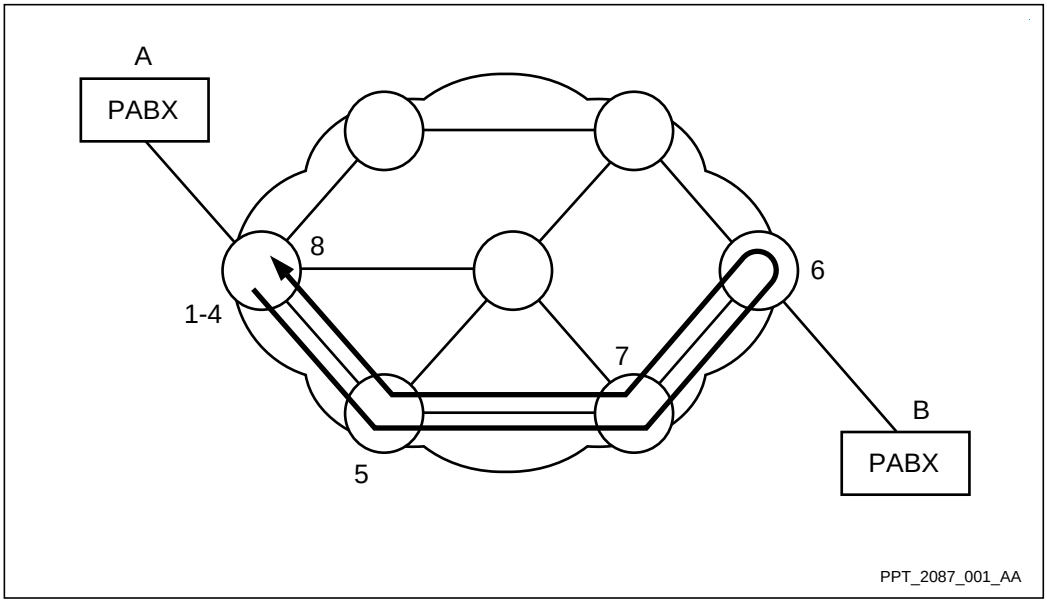

Path termination

Path termination is done in a progressive fashion. A path termination packet is sent from the end point that initiates the termination to the remote end point. At each Passport trunk along the path, the reserved bandwidth for both directions of transmission is released. The remote end point returns a path termination confirmation packet (this action deallocates the path LCN) to the initiating end point.

PORS data transfer

The packet forwarding function is invoked when a packet arrives from a Passport trunk to look up the path LCh in the forwarding table. This process determines the next hop (Passport trunk or path end point). The packet forwarding function can discard packets based on the congestion level and on the discard priority of the path. The packet forwarding function ensures that the packet is put on the queue indicated by the emission priority of the path.

Example of PORS data transfer

The transfer of data between Source A and Destination B involves the following steps (see the figure “PORS data transfer” (page 139) for an illustration of PORS data transfer):

1 PORS packet forwarding examines the packet and looks up the path LCh to determine the next hop. The packet can be discarded depending on the congestion level and discard priority of the path.

2 The packet is queued for forwarding, based on the emission priority of the path and forwarded to the next hop.

3 Steps 1 and 2 are repeated at each hop until the destination is reached.

4 At the destination, the packet is decapsulated (headers removed) and passed up to the corresponding higher layer protocol.

Figure 38
PORS data transfer

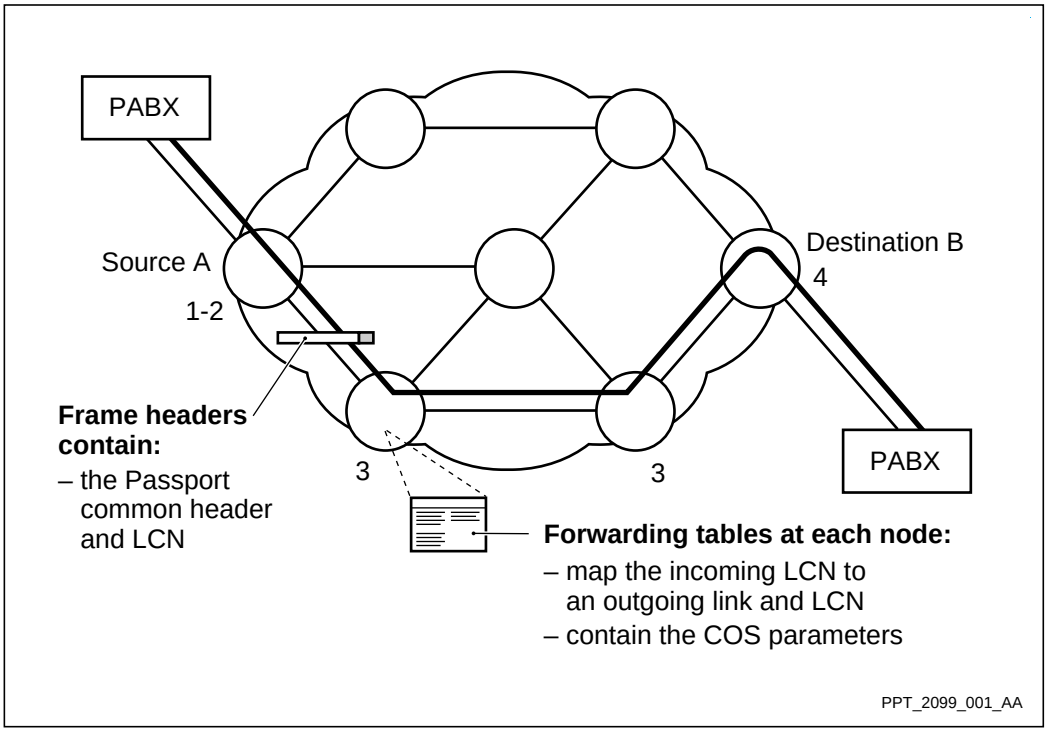

PORS route interruption

If network changes interrupt an existing path, teardown packets are propagated towards the end points from either side of the interruption, deallocating resources along the way. All the LCNs and the bandwidth reserved prior to the point of the network changes are deallocated. The source end point reports the reason for the interruption to the RS and requests a new route. The RS updates the topological database, thus learning from the interruption, and selects another route. The source end point starts the instantiation procedure again. If another route is not available, RS informs the source end point that no route is available.

PORS components

Each Passport 7400, 15000, 20000 node has a single *Routing* component that maintains a current view of the node and network topology. For PORS, the *Routing* component has subcomponent *Pors.*

Routing Pors is the component residing on all Passport nodes that carry PORS traffic. *Routing Pors* displays some attributes of the PORS connections on a node. *Routing Pors* also displays the attributes for setting optimization timing for PORS connections.

For more information on components, see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Chapter 6
ATM routing system

This chapter describes ATM networking on Passport, and includes the information on the following topics:

For more detailed information on ATM routing and ATM routing services, see "ATM routing references" (page 19).

Role of ATM routing

ATM routing establishes dynamic routes between Passport nodes and other ATM equipment. ATM routing supports the following:

- ATM bearer service (ABS)
- AAL1 circuit emulation service (AAL1 CES)
- frame relay to ATM interworking service
- ATM multiprotocol encapsulation (MPE) service

Passport supports two types of routing for ATM:

- static routing, based on the ATM Forum user-to-network interface (UNI) specification Version 3.0 and Version 3.1, the interim inter-switch signaling protocol (IISP) specification Version 1.0, and the ATM inter-network interface (AINI) specification, *Version 1.0*
- dynamic routing (also known as dynamic route discovery) based on the ATM Forum private network-network interface (PNNI) specification Version 1.0

ATM routing features

ATM networking provides dynamic connection setup between Passport nodes, and lets Passport switches interwork with both Nortel Networks switches and switches from other ATM equipment vendors.

Passport ATM networking provides the addressing, signaling, and routing facilities to support switched virtual connections (SVCs), switched virtual paths (SVPs), soft permanent virtual connections (SPVCs), and soft permanent virtual paths (SPVP). These networking capabilities enable you to set up ATM connections in real time. These capabilities also provide a framework for future applications.

For more detailed information about Passport ATM routing features, see 241-5701-702 *Passport 7400, 15000, 20000 ATM Routing and Signaling Fundamentals*.

Standards-based routing

Standards-based routing permits a common networking evolution with the Vector and Concorde systems, and ensures multivendor operation across ATM networks.

Passport ATM networking is based on the following ATM Forum standards:

- *Integrated Local Management Interface (ILMI) Specification Version 4.0 (af-ilmi-0065.000),* ATM Forum Technical Committee, *1996*
- *Interim Inter-switch Signaling Protocol (IISP) Specification Version 1.0, (af-pnni-0026.000),* ATM Forum Technical Committee, 1994

- *ATM Inter-Network Interface (AINI) Specification, (af-cs-0125.000),* ATM Forum Technical Committee, 1999
- *Private Network-Network Interface (PNNI) Specification Version 1.0 (af-pnni-0055.000),* ATM Forum Technical Committee, 1996
- *PNNI V1.0 Errata and PICS (af-pnni-0081.000),* ATM Forum Technical Committee, 1997
- *Traffic Management Specification Version 4.0 (af-tm-0056.000),* ATM Forum Technical Committee, 1996
- *User-to-Network Interface Specification Version 3.0 (af-uni-0010.001),* ATM Forum Technical Committee, 1993
- *User-to-Network Interface Specification Version 3.1 (af-uni-0010.002),* ATM Forum Technical Committee, 1993
- *User-Network Interface Signaling Specification Version 4.0 (af-sig-0061.000),* ATM Forum Technical Committee, 1996

The previous listed standards enable Passport nodes to support UNI, IISP, and AINI interfaces as well as the basic Passport ATM interface. The PNNI standards enable Passport nodes to support the PNNIs required for dynamic networking. PNNI, UNI 4.0, and AINI ensure smooth transition as networking standards evolve. Passport supports all mandatory functionality subsets described in Annex G of the ATM Forum's PNNI Specification, Version 1.0.

The signaling procedures used to establish connections across an ATM interface comply with ATM Forum standards listed above. These signaling protocols are based on the following:

- ITU-T Recommendation Q.850, *Usage of Cause and Location in the Digital Subscriber Signalling System No. 1 and the Signalling System No. 7 ISDN user part*, 1993
- ITU-T Recommendation Q.2610, *Common Aspects of B-ISDN Application Protocols for Access Signaling and Network Signalling and Interworking*, 1995
- ITU-T Recommendation Q.2931, *B-ISDN User-Network Interface Layer 3 Specification for Basic Call/Bearer Control*, 1995

- ITU-T Recommendation Q.2971, *B-ISDN DSS2 UNI Layer 3 Specification for Point-to-Multipoint Call/Connection Control*, 1996

Supported connection types

For ATM signaling interfaces, the following ATM connections are supported:

- nailed-up permanent virtual connections (NPVC) and nailed-up permanent virtual paths (NPVP), which are manually provisioned on a hop-by-hop basis and remain in effect until they are manually removed through provisioning
- SPVCs and SPVPs, which are provisioned at the source end point and optionally provisioned at the destination end point.
- SVCs and SVPs which are not provisioned, but are established dynamically through signaling procedures.

Passport supports switched virtual paths (SVP) as part of the signaling requirement for SPVPs. SVPs are possible between relay points but not between relay points and end points.

For more information about connection types, see 241-5701-702 *Passport 7400, 15000, 20000 ATM Routing and Signaling Fundamentals*.

Interworking with MCS

ATM routing facilities transmit frame relay service traffic in MCS connections. Each MCS connection is an SPVC that takes advantage of ATM features such as dynamic call routing and automatic rerouting.

Hitless ATM services for Passport 15000 and 20000

For Passport 15000 and 20000, the following ATM services offer a hitless service:

- PVCs (VCCs, VPCs and VPT VCC (basic and standard) connections)
- point-to-point SVCs
- source and destination SPVCs
- point-to-point SVPs
- source and destination SPVPs

A service is hitless when the software that provides the service can run uninterrupted, even when the hardware providing the service changes.

The ATM hitless services can be offered on UNI, IISP, AINI and PNNI interfaces in non-associated signaling configurations.

Hot standby applications operate with a standby instance of the software that is fully synchronized with the active instance of the software. Hot standby applications and features use equipment sparing of optical and electrical FPs to incur a minimal interruption of cell forwarding and maintain any connections that are established.

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.

Hot standby applications and features can co-exist with other hot and warm applications and instances on an LP. Although you can create an LP that mixes cold standby features with hot standby or warm standby features, Nortel Networks does not recommend this action. A single cold standby application or feature in an LP changes all other applications and features into cold standby.

Hitless services minimize the interruption of cell forwarding only. During an FP switchover, all applications that were running on the FP can lose administrative data even if it is a hot standby or a warm standby application. Specifically, the following types of data can be lost:

- performance statistics, such as cell counts
- partial accounting records and any accounting records that reside in the memory of the FP before the FP switchover
- the OSI state (A service that is locked before the FP switchover becomes unlocked after the switchover.)

ATM routing architecture

This section describes the ATM routing architecture and includes information on the following topics:

- “Networking system for ATM routing” (page 146)

Networking system for ATM routing

Figure "Passport networking architecture" (page 27) illustrates where ATM routing systems fit in to the networking architecture layers. The illustration shows two areas of ATM routing: ATM networking and NPVCs. ATM routing does not need the support of the Passport trunking and base routing systems.

ATM interfaces

An ATM interface is the Passport software component that controls the connections on an ATM link. One ATM interface is located at each end of a physical ATM link. ATM networking supports the following interfaces:

- basic interface
- user-to-network interface (UNI)
- interim inter-switch signaling protocol (IISP) interface
- ATM inter-network interface (AINI)
- private network-to-network interface (PNNI)

Passport switches support ATM Address Screening, a security feature that screens incoming signaling connections based on their called (destination) or calling (source) addresses. Address screening allows you to provide security over a shared ATM network by configuring specific addresses to be accepted or rejected for each UNI, IISP, or AINI interface. The Nortel Networks technical publication (NTP), 241-5701-710 *Passport 7400, 15000, 20000 ATM Configuration Guide*, provides the details on how to use the ATM Address Screening feature.

Basic interface

A basic interface does not have signaling capabilities. As a result, a basic interface supports only NPVCs and NPVPs. A basic interface supports links between nodes, and between ATM terminal equipment and the ATM network.

User-to-network interface

The UNI interface connects ATM terminal equipment to the network. The signaling protocol on this type of interface is the ATM Forum UNI Version 3.0, 3.1, or 4.0. The UNI standard is based on ITU-T Q.2931.

UNIs support NPVCs, NPVPs, SPVCs, SVCs, SPVPs, and SVPs. The interface establishes the connections using the UNI signaling protocol. For dynamic connections, signaling provides the procedures for establishing, maintaining, and clearing connections across an ATM interface.

UNI interfaces also use integrated local management interface (ILMI) control procedures. ILMI signaling provides the bidirectional exchange of control information and dynamic address registration data across the interface.

Interim inter-switch signaling protocol interface

The IISP interface is used between network nodes. The IISP signaling protocol on these interfaces provides interconnection between Passport switches, other Nortel Networks switches, and switches from other vendors.

IISP interfaces also support all three connection types. The figure “UNI, IISP, and AINI interfaces” (page 148) shows a typical ATM network with UNI interfaces at the edge of the network, and with IISP, UNI, and AINI interfaces connecting the Passport nodes between them.

Figure 39
UNI, IISP, and AINI interfaces

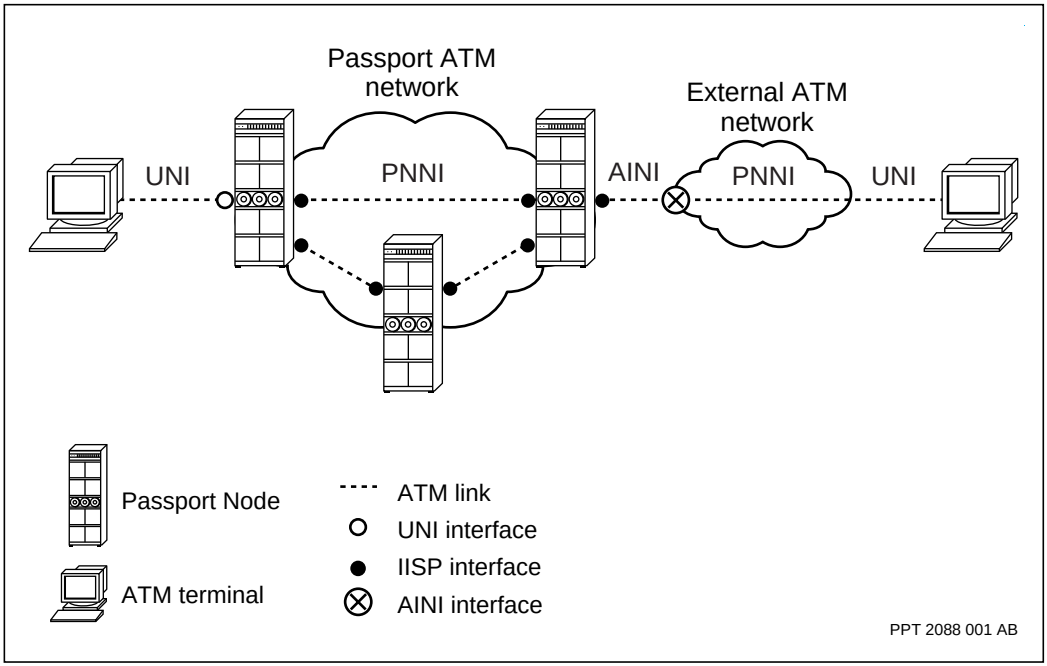

ATM inter-network interface

Passport is compliant with the ATM Forum standard version 1.0 AINI specification. The ATM inter-network interface (AINI) is used between network nodes. The AINI signaling protocol on these interfaces provides interconnection between Passport switches, other Nortel Networks switches, and switches from other vendors.

AINI interfaces also support all three connection types. The figure “UNI, IISP, and AINI interfaces” (page 148) shows a typical ATM network with UNI interfaces at the edge of the network, and with IISP, UNI, and AINI interfaces connecting the Passport nodes between them.

Private network-to-network interface

The PNNI also supports connections between network nodes, including other Passport switches, Nortel Networks switches, and switches from other vendors. The signaling protocol is the ATM Forum PNNI Version 1.0. The Passport PNNI implementation supports the following:

- Nailed-up connections and functionality. Nailed-up ATM bearer connections and logical Passport trunks can traverse PNNI interfaces.
- Standards-based SPVCs between Passport and any ATM Forum PNNI Version 1.0 compliant vendor equipment that implements SPVCs and SPVPs.
- Point-to-point SVCs and SVPs and point-to-multipoint SVCs
- UNI/ILMI scope mapping to PNNI
- All mandatory PNNI functionality subsets as described in Annex G of the ATM Forum PNNI Specification Version 1.0.

Figure 40
PNNI interfaces in a hybrid network scenario

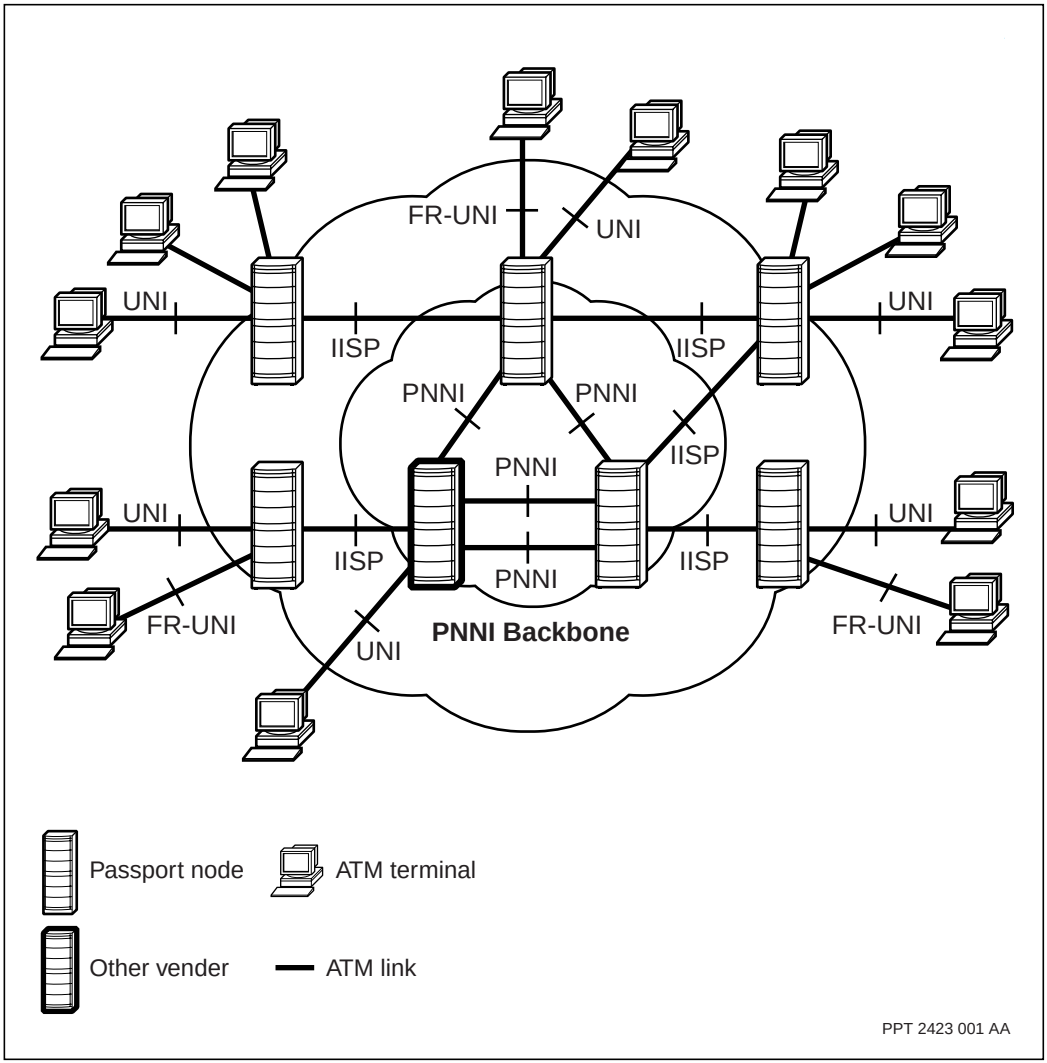

PNNI provides dynamic routing and signaling. PNNI-based switching systems monitor network topology and available resources. As a result, calls automatically route around points of congestion and failure.

The figure “PNNI interfaces in a hybrid network scenario” (page 150) shows a typical hybrid ATM network with UNI and IISP interfaces at the edge of the network, and PNNI links connecting the Passport nodes in the core.

Virtual interfaces

Passport nodes can support multiple virtual interfaces on one physical port by tunnelling SVCs through permanent virtual paths (PVP). A virtual interface may be a UNI, IISP, AINI, or PNNI interface. Each port requires one physical path to the core network. Through configuration, you associate each virtual interface with a PVP, and each virtual interface and has its own signaling channel for setting up SVCs within the PVP. In this way, Passport nodes can provide the required SVC tunnelling through a PVP network.

Note: ILMI is not available for virtual UNIs.

ATM virtual connections

Passport supports several types of virtual connections:

- “Nailed-up permanent virtual connections and paths” (page 154)
- “Soft permanent virtual connections and paths” (page 157)
- “Switched virtual connections and paths” (page 159)

Passport establishes these connections using the virtual channel connections (VCCs) and virtual path connections (VPCs) in ATM links. A VCC is an association established between two ATM layer users that communicate through a virtual channel. A VPC is a similar association using a virtual path. VPCs are supported only for NPVCs and NPVPs.

ATM components

Each VCC or VPC that is associated with an ATM port is represented by a *Vcc* or *Vpc* component. These components are subcomponents of their associated *ATMInterface* component (*AtmIf*) which represents the ATM interface. One *AtmIf* component represents the ATM layer functionality for each ATM port. The *AtmIf* component has a provisioned link to the port component that embodies the physical layer function. The figure “ATM components” (page 152) shows these relationships.

Figure 41
ATM components

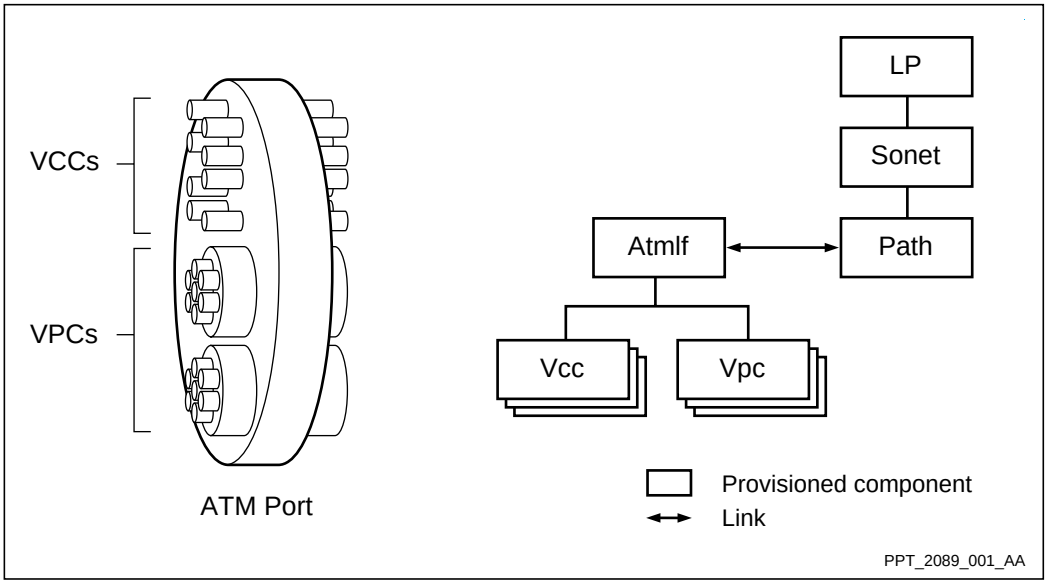

The UNI, IISP, AINI, and PNNI interfaces are represented by *Uni*, *Iisp*, *Aini*, and *Pnni* components under their associated *AtmIf* component. You provision one of these subcomponents for each *AtmIf* to define the type of interface. If you do not provision one of these subcomponents, the interface is a basic ATM interface. The *Uni*, *Iisp*, *Aini* and *Pnni* components include subcomponents that define addressing and signaling, ILMI functions for UNIs, and route control connections for PNNIs. The figure "ATM interface components" (page 153) shows the relationships between these components.

Figure 42
ATM interface components

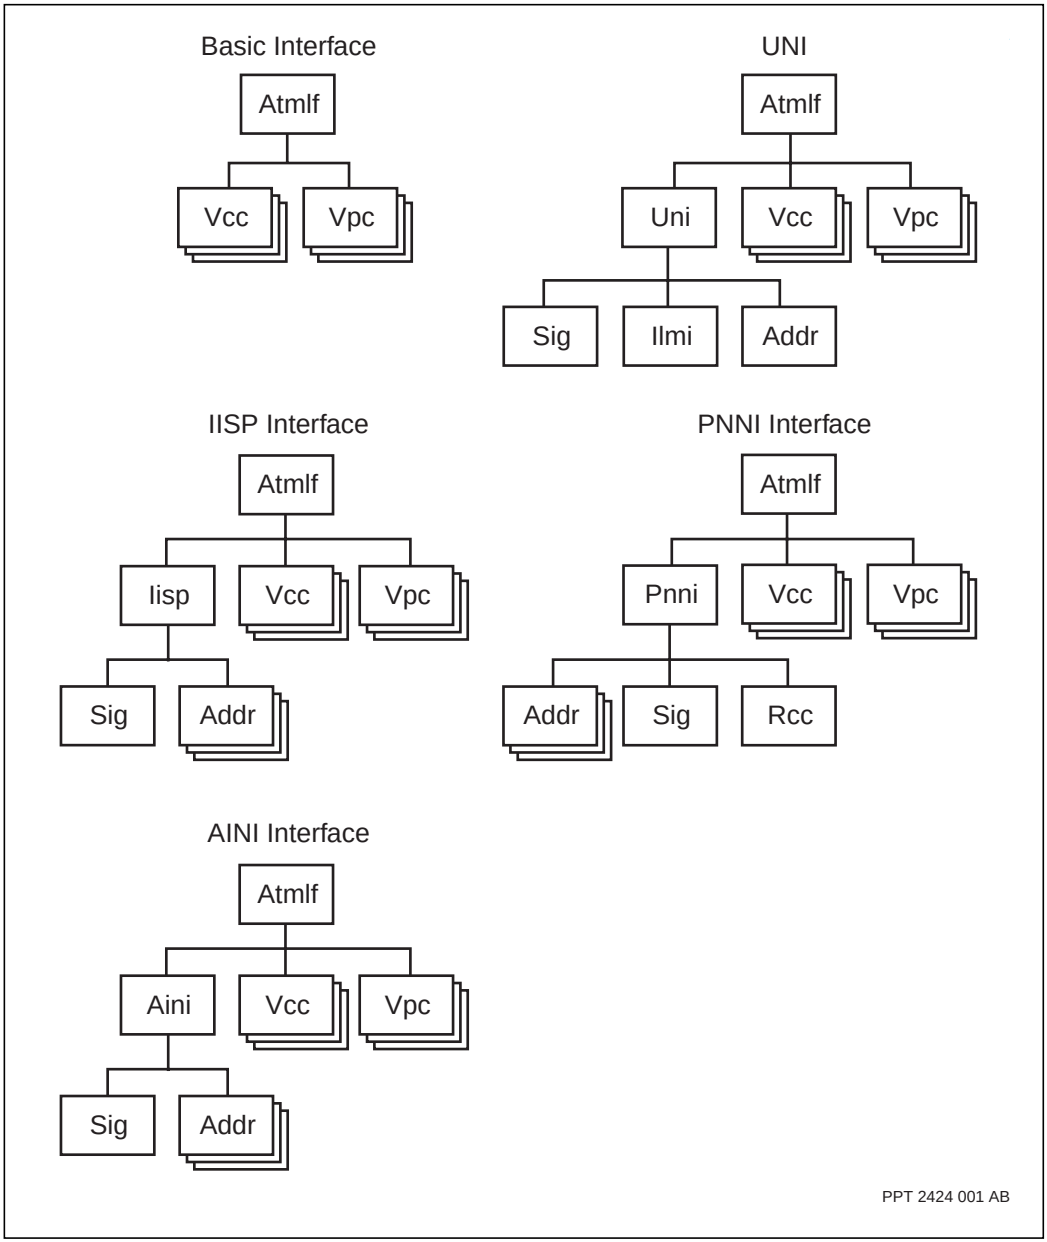

Nailed-up permanent virtual connections and paths

NPVCs and NPVPs support the ATM bearer service (ABS). You implement NPVCs and NPVPs in areas of the network that need consistent, high-volume traffic. If a facility or node failure occurs along the route, the connection goes down. Automatic rerouting is not possible. The figures “Example of an NPVC” (page 155) and “Example of an NPVP” (page 156) shows examples.

Figure 43
Example of an NPVC

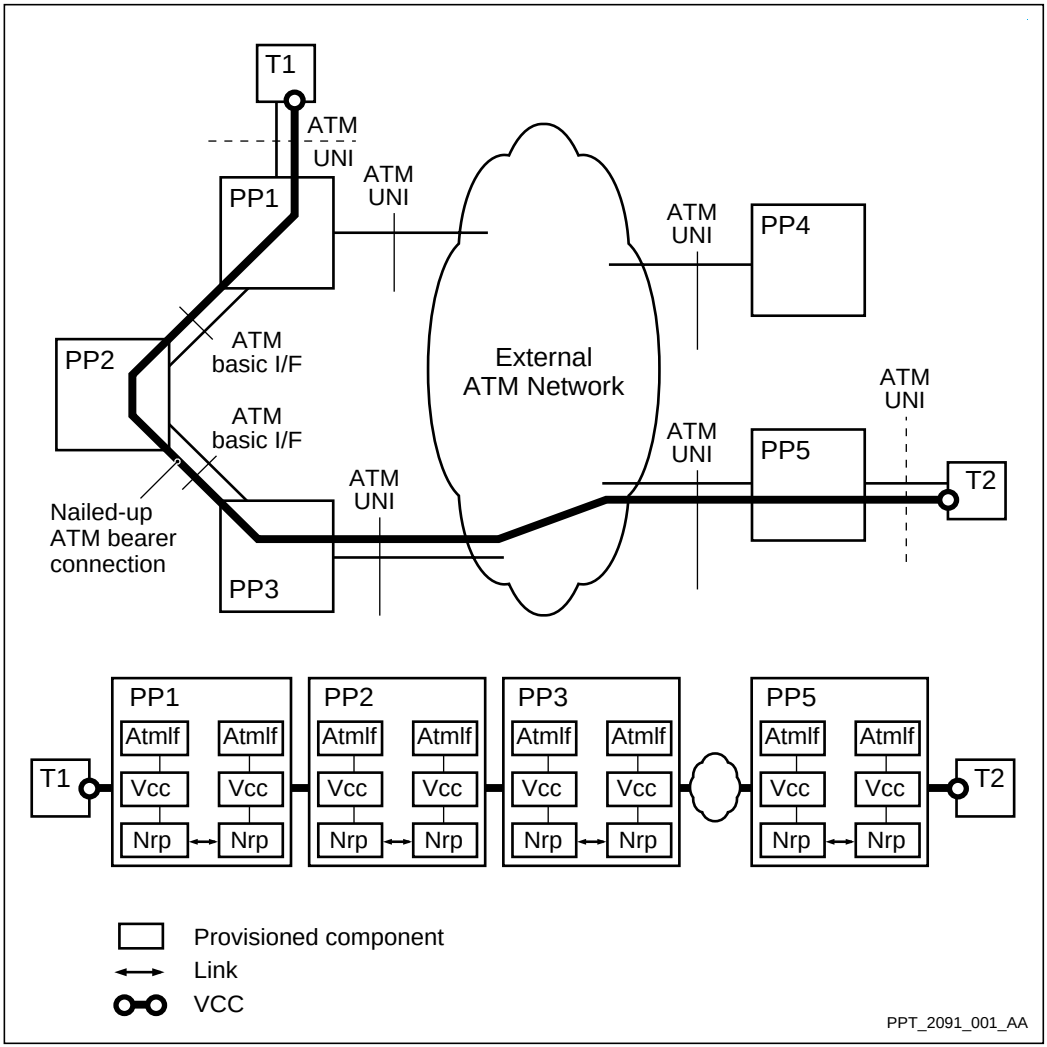

Figure 44
Example of an NPVP

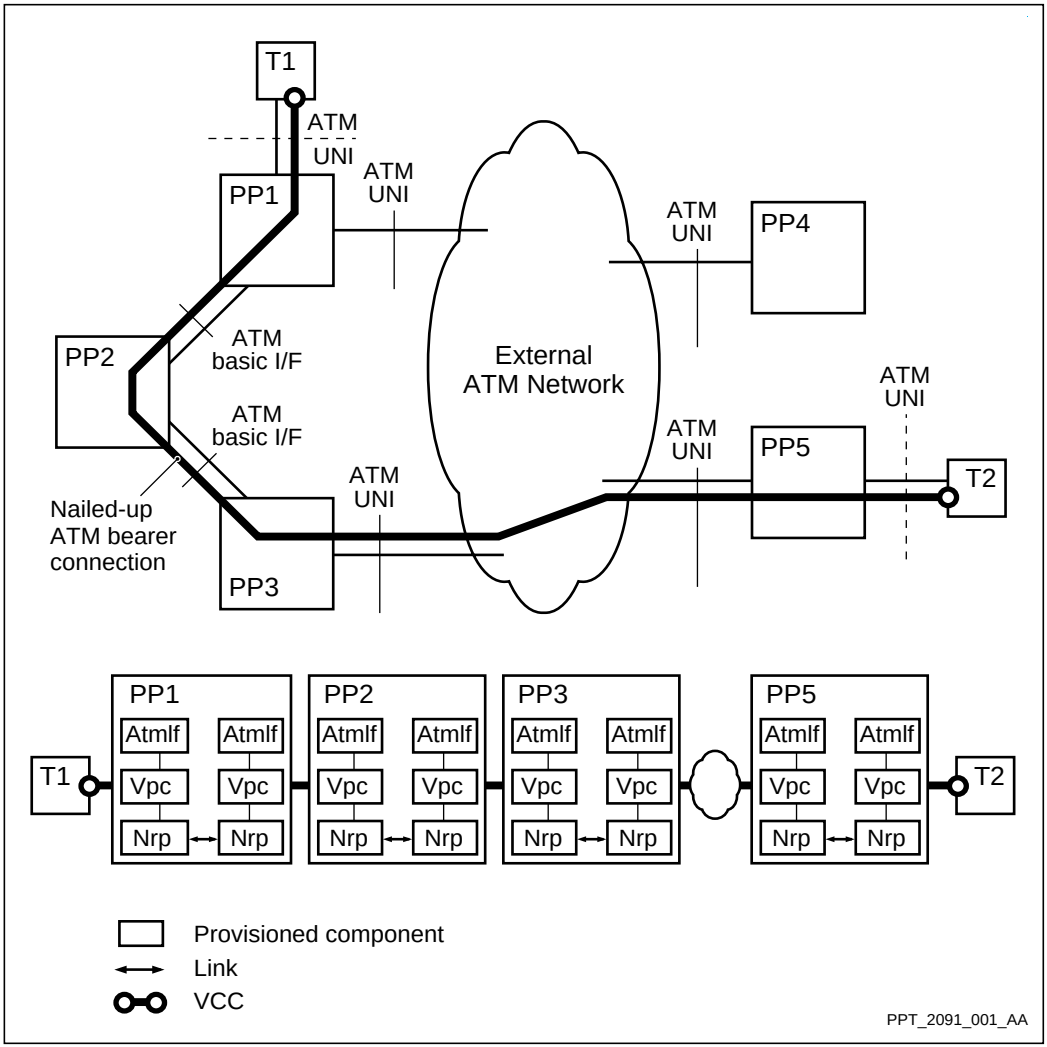

You manually provision the NPVCs and NPVPs, which remain in effect until you manually deprovision them. To establish the connection, you associate each pair of *Vcc* or *Vpc* components for each adjacent node in the connection using *NailedUpRelayPoint* (*Nrp*) components. The figures “Example of an

NPVC" (page 155) and "Example of an NPVP" (page 156) show how to establish an NPVC and an NPVP across a Passport network by linking *Nrp* components across four nodes.

Because NPVCs and NPVPs do not have a dynamic element, they are not strictly considered part of ATM networking. (The term "ATM networking" implies dynamic operation.) However, Passport supports NPVCs and NPVPs on all ATM interfaces to accommodate high traffic demands that require dedicated connections.

Soft permanent virtual connections and paths

The network establishes SPVCs and SPVPs in real time using signaling procedures. In an SPVC or SPVP, route selection and connection establishment is automatic. Further, rerouting is automatic in the case of a facility or node failure.

Through UNI and IISP, Passport supports proprietary SPVCs that are specific to the Nortel Networks product line. Through the PNNI and AINI protocols, Passport supports standards-based SPVCs and SPVPs that can terminate on any ATM interface, regardless of the equipment vendor.

Provisioning for SPVCs and SPVPs is simpler than for NPVCs and NPVPs since you need provision only the source end point. To establish the connection, the network creates dynamic *Vcc* and *Vpc* components on demand at run time. (The *Vcc* and *Vpc* components are either provisioned or dynamic.)

The figure "Example of an SPVC" (page 158) shows an example of a *VCC-based* SPVC linking three Passport nodes. In the first node, a provisioned *Src* component defines the source point. To establish the connection from that point onward, the network dynamically links two *RelayPoint* (*Rp*) components in each Passport node. A dynamic *Dest* component defines the destination point. This example shows UNIs, an IISP interface, and a PNNI in a single connection. The figure "Example of an SPVP" (page 159) shows a similar example for VPC-based SPVPs.

Figure 45
Example of an SPVC

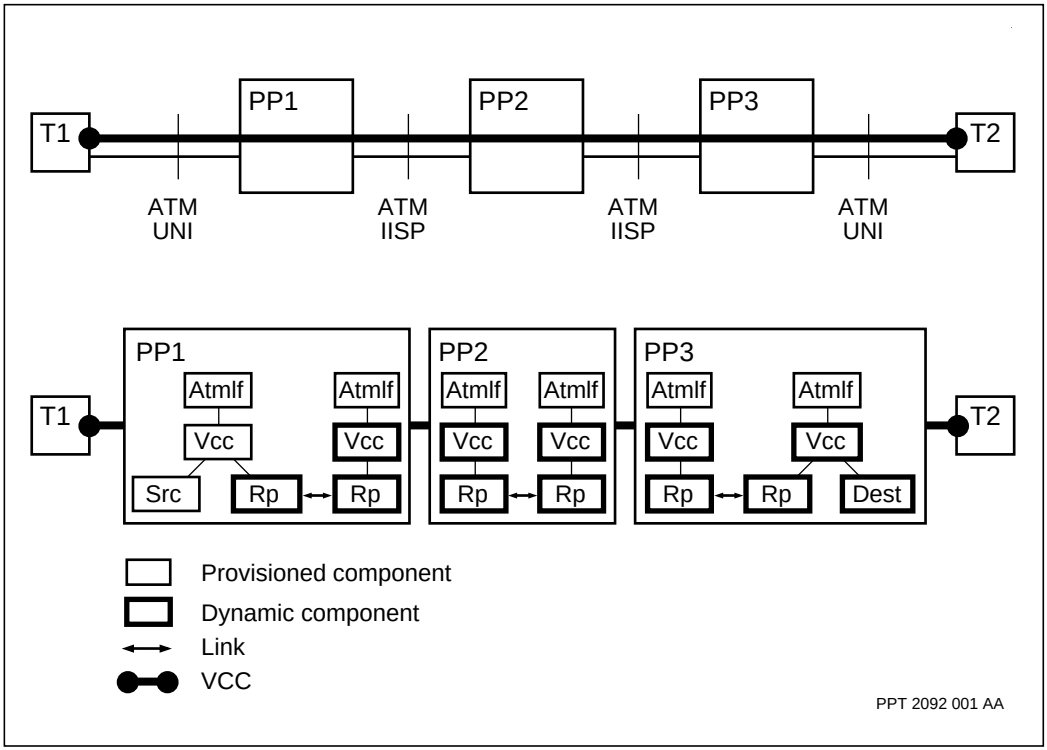

Figure 46
Example of an SPVP

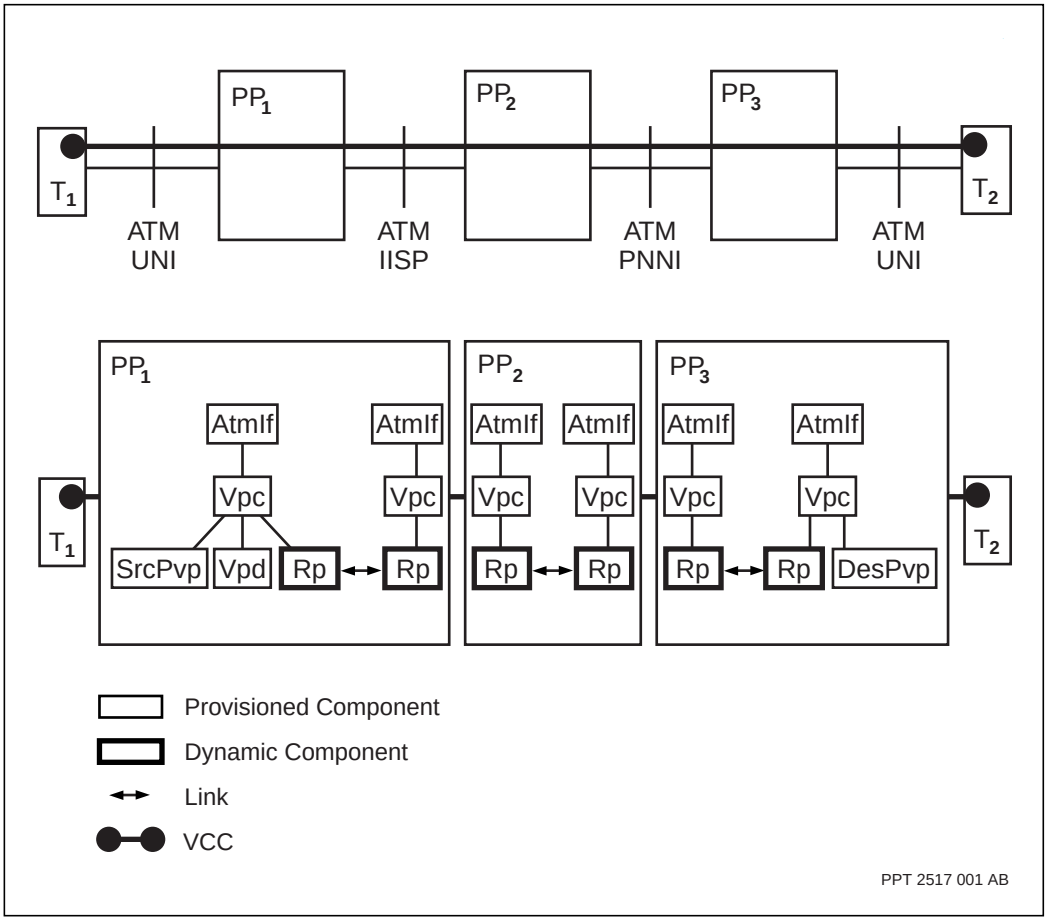

Switched virtual connections and paths

The network sets up and takes down SVCs and SVPs on demand. For connections over UNIs, point-to-point SVCs and SVPs; and point-to-multipoint SVCs using static routing are possible. For SVCs over PNNIs, both point-to-point and point-to-multipoint QoS-based dynamic connections are possible.

The figure “Example of an SVC” (page 160) shows three Passport nodes in series to form an SVC under static routing. A similar configuration using SVPs is shown in the figure “Example of an SVP” (page 161). The network dynamically links two *Rp* components in each Passport node. A provisioned source point is not needed for an SVC. This example shows UNIs, and IISP interface, and a PNNI in a single connection.

Figure 47
Example of an SVC

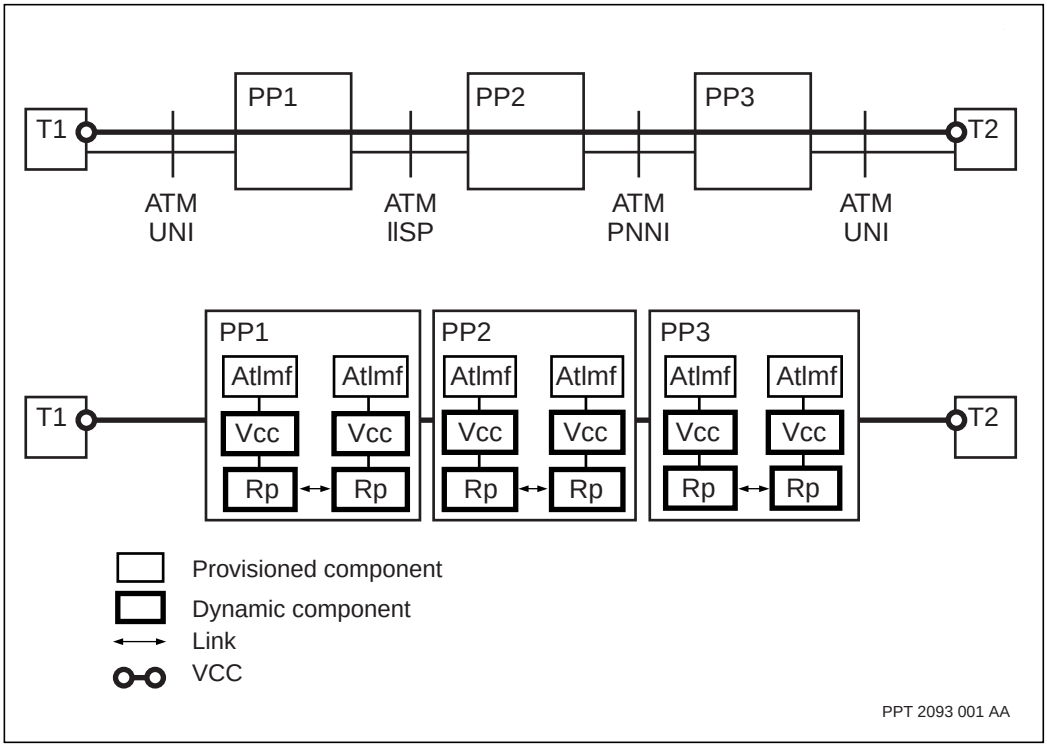

Figure 48
Example of an SVP

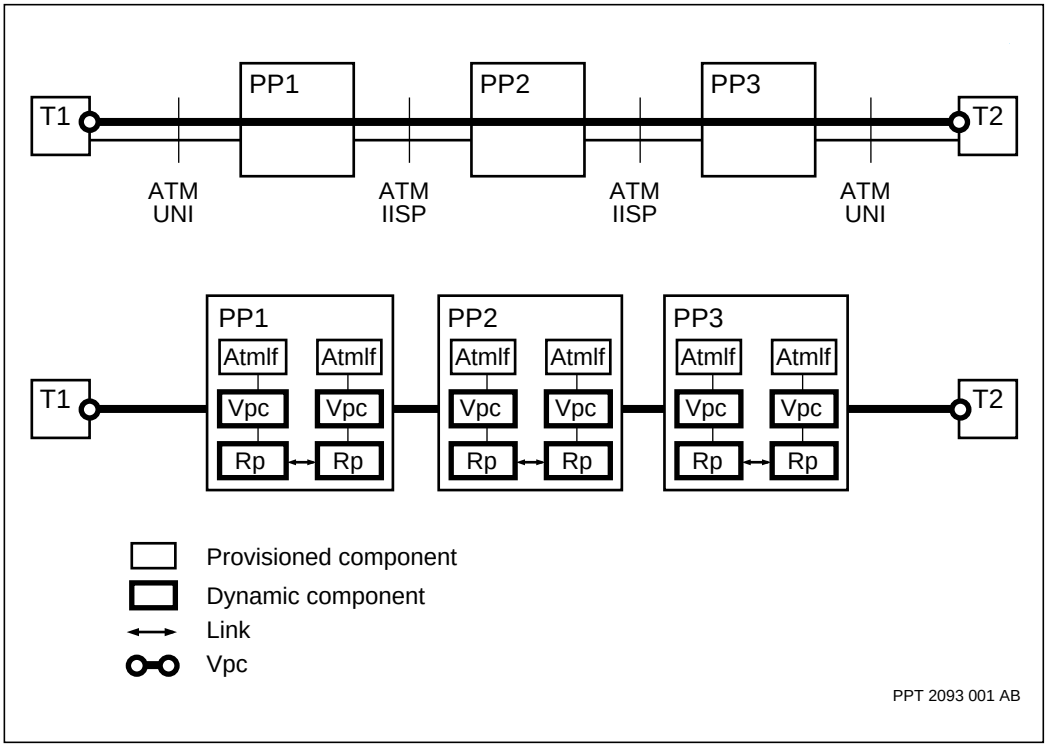

Dynamic *EndPoint* (*Ep*) components are created for the signaling and ILMI VCCs. The figure "Signaling channel VCC (SVC under static routing)" (page 162) shows an example of the use of *Ep* components in establishing a signaling channel.

The figure "Signaling channel and RCC VCCs (SVC under dynamic routing)" (page 163) shows an example of the use of dynamic *Ep* components in establishing the PNNI routing control channel (RCC), signaling channel, and ILMI channel for an SVC under PNNI.

Figure 49
Signaling channel VCC (SVC under static routing)

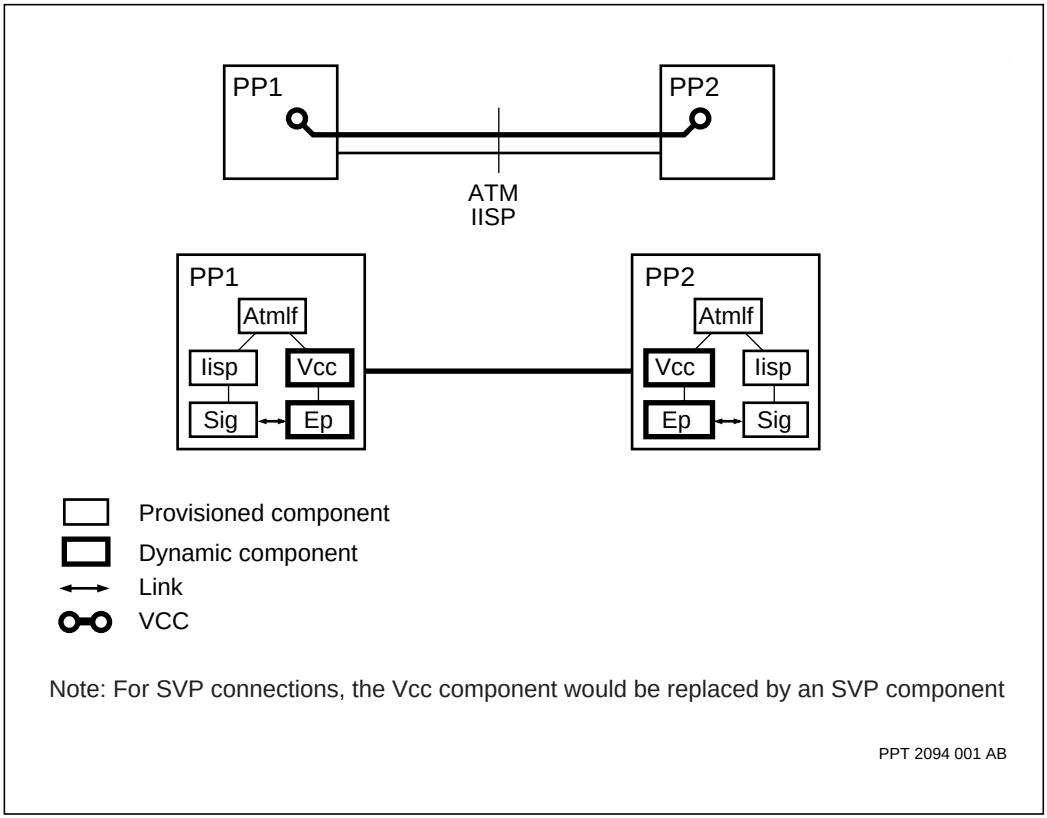

Figure 50
Signaling channel and RCC VCCs (SVC under dynamic routing)

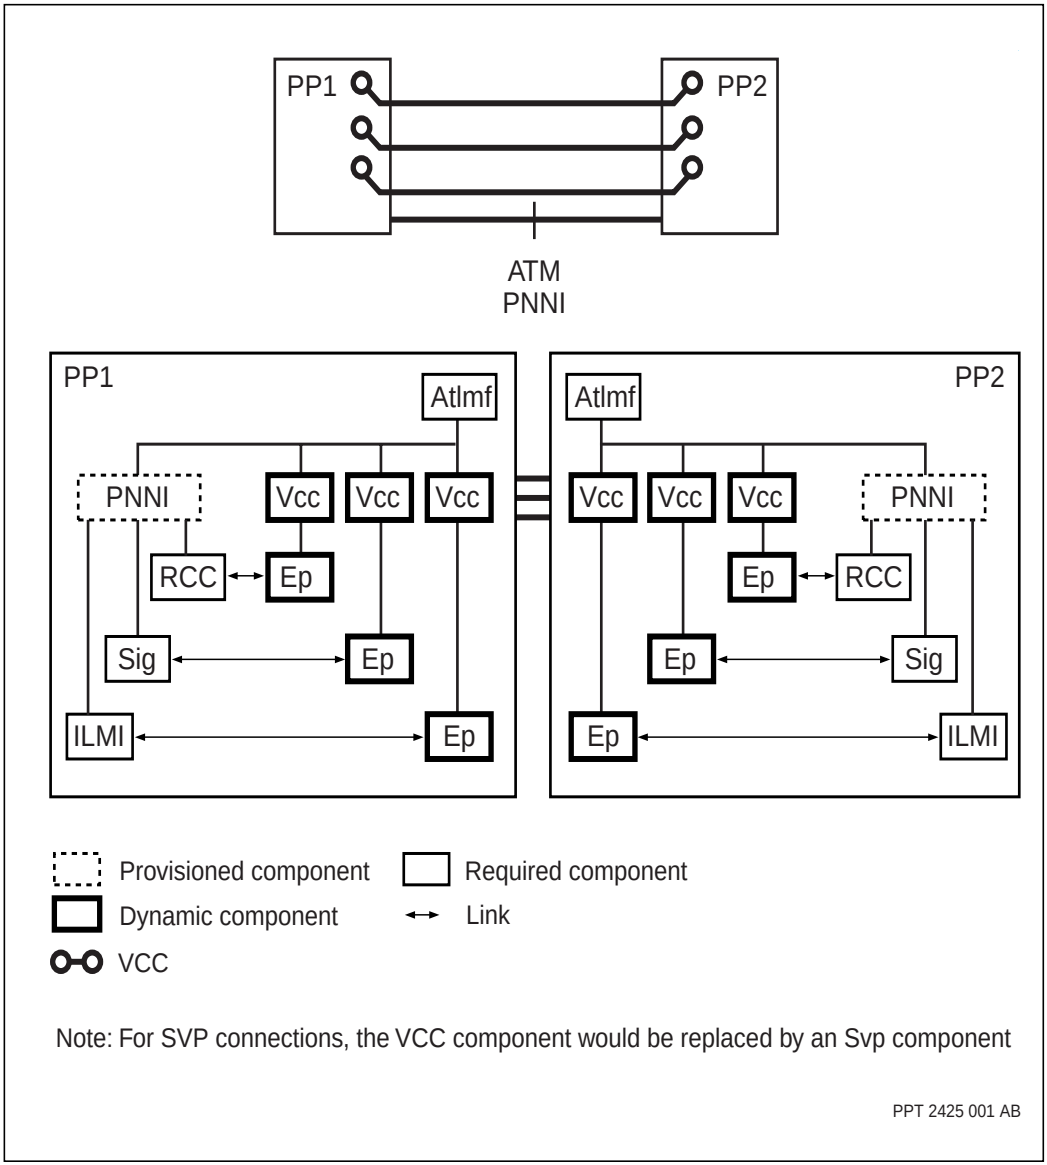

Dynamic Vcc and Vpc components

The capacity to create dynamic *Vcc* and *Vpc* components is essential to supporting SPVCs, SPVPs, SVCs and SVPs. The nodes, according the defined protocol for an interface, create dynamic components at run time as needed (as opposed to provisioned components that you provision and are permanent until you deprovision them). The *Vcc* and *Vpc* components are either provisioned or dynamic. It is provisioned when part of a nailed-up connection and dynamic when part of a switched connection.

When created as part of a switched connection, a *Vcc* component is either a relay-point or an end-point. This definition is created through the *RelayPoint* and *EndPoint* dynamic subcomponents under the *Vcc* component. When created as part of a switched connection, a *Vpc* component is a relay-point only. This definition is created through the *RelayPoint* dynamic subcomponent under the *Vpc* component. These subcomponents are described in the following paragraphs.

Dynamic relay points

Dynamic relay points are used to set up switched VCCs and VPCs. For each Passport node along a connection route, two dynamic *Vcc* or *Vpc* components are linked together through their *RelayPoint* subcomponents. *RelayPoint* components are used to establish SPVCs, SPVPs, and SVCs in the same way as *NailedUpRelayPoint* components are used to establish NPVCs and NPVPs. In a switched connection, however, the *RelayPoint* component is dynamic. The figure “Switched VCC showing relay points” (page 165) shows how a switched VCC connection is established across a Passport network by linking dynamic *RelayPoint* components together.

Figure 51
Switched VCC showing relay points

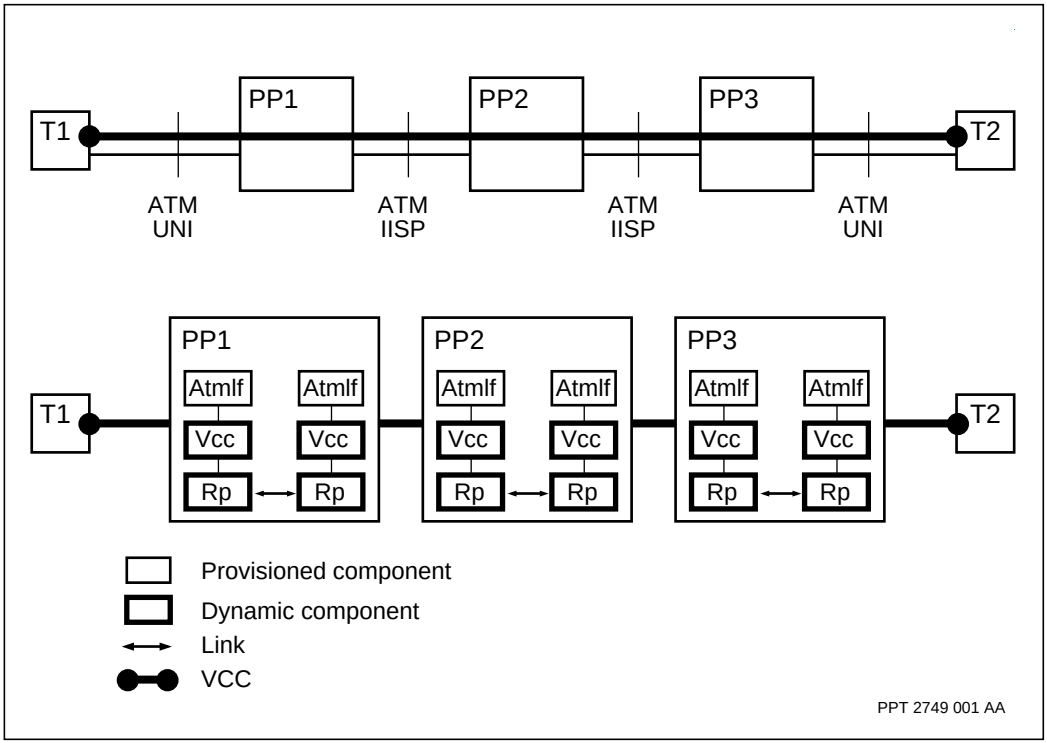

Dynamic end-points

There are three applications that need to communicate between nodes using a dynamic VCC:

- signaling channel
- ILMI channel
- RCC

For each application, a dynamic *Vcc* component with a dynamic *endPoint* (*Ep*) subcomponent is created at each end of the ATM connection. End-points are used because these applications generate/terminate ATM traffic on the

VCC. The figure “Signaling channel VCC” (page 166) illustrates this configuration by showing how a signaling channel is established across an IISP interface.

Figure 52
Signaling channel VCC

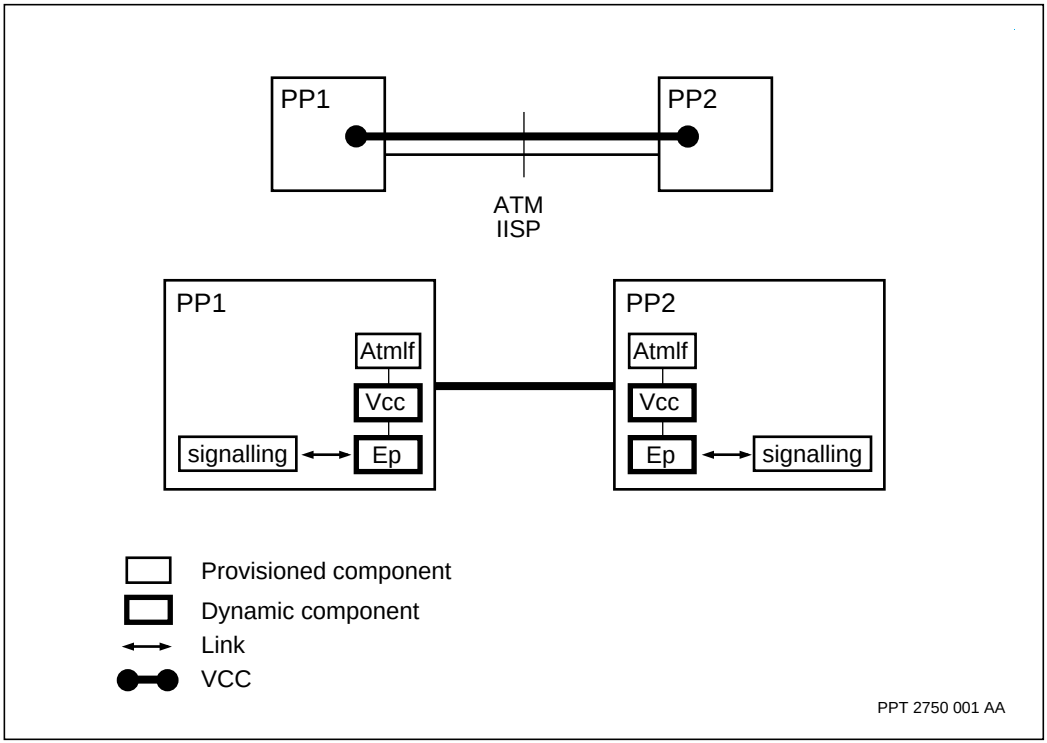

EndPoint components are used in the same way as *NailedUpEndPoint* components are used in nailed-up connections: they are used for applications that generate/terminate ATM traffic from within a Passport node. The difference is that the *EndPoint* component is dynamic.

ATM traffic management features

An ATM source needs to send the connection traffic characteristics to the ATM network before the traffic can be transported efficiently. The traffic description information is either provisioned or signaled at call establishment time. This information includes characteristics such as ATM service category and cell rate parameters.

Depending on the traffic description, different traffic management strategies are selected. These strategies include

- connection admission control (CAC): to permit or reject connection setup
- traffic shaping: to control traffic in the transmit direction
- usage parameter control (UPC): to monitor and control traffic in the receive direction
- buffer, queue, and congestion/discard management: to handle congestion situations as they arise

Traffic management strategies for connections over PNNIs include the following:

- resource availability information group (RAIG), which defines topology metrics and attributes for each PNNI-based link in the network, and which the network uses to select the best path for a connection
- two additional QoS information elements (QoS-IE) in the SETUP and CONNECT messages, which supplement the QoS parameter IE defined for UNI Versions 3.0 and 3.1
- two optional traffic descriptor information elements (TD-IE) in the SETUP message
- traffic management for the PNNI routing control channel (RCC)

For more information on traffic management, see the following documents:

- 241-5701-705 *Passport 7400, 15000, 20000 ATM Traffic Management Fundamentals*
- 241-5701-706 *Passport 7400, 15000, 20000 ATM Traffic Shaping and Policing*

- 241-5701-707 *Passport 7400, 15000, 20000 ATM Queuing and Scheduling*
- 241-5701-708 *Passport 7400, 15000, 20000 ATM CAC and Bandwidth Management*

ATM routing mechanisms

This section describes ATM routing mechanisms and includes the following topics:

ATM addressing system

To establish SPVCs, SPVPs, SVPs, and SVCs across the network, the ATM source and destination points need to be uniquely identified with ATM addresses.

Each UNI, IISP, AINI and PNNI interface has one default address and can have other associated addresses. UNIs can have static (provisioned) or dynamic addresses that are registered through ILMI control procedures. IISP and AINI interfaces have static addresses.

Both static and dynamic addressing methods use the OSI Network Service Access Point (NSAP) format. Passport ATM routing supports all three NSAP address formats: Data Country Code (DCC), International Code Designator (ICD), and E.164.

NSAP addressing plan format

NSAP addresses contain 40 hexadecimal digits divided into two parts.

- The first 26 digits form the network prefix.
- The last 14 digits form the end system identifier (ESI).

The figure "NSAP address format" (page 169) shows the NSAP address format.

Figure 53
NSAP address format

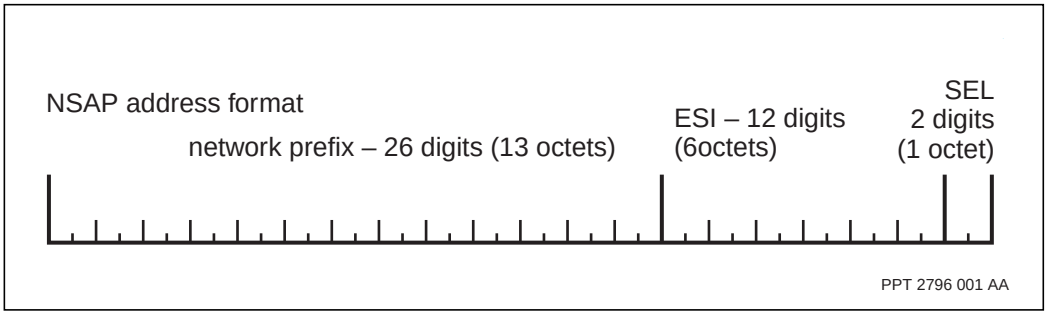

Default addresses

Each UNI, IISP, AINI and PNNI interface has a built-in address that is automatically registered in the routing tables of the node as its default address. You can use this address as the destination address for an SPVC or SPVP, or as the called party number for an SVC call.

The figure "Default address format" (page 170) shows the default address format. The first part of the default address is the node, or network, prefix. (One network prefix is provisioned for each Passport node.) The ESI portion of the address includes the instance value of the *AtmIf* component for that interface. The default address is displayed as an *Address* subcomponent under the *Uni*, *Iisp*, *Aini*, or *Pnni* component.

Figure 54
Default address format

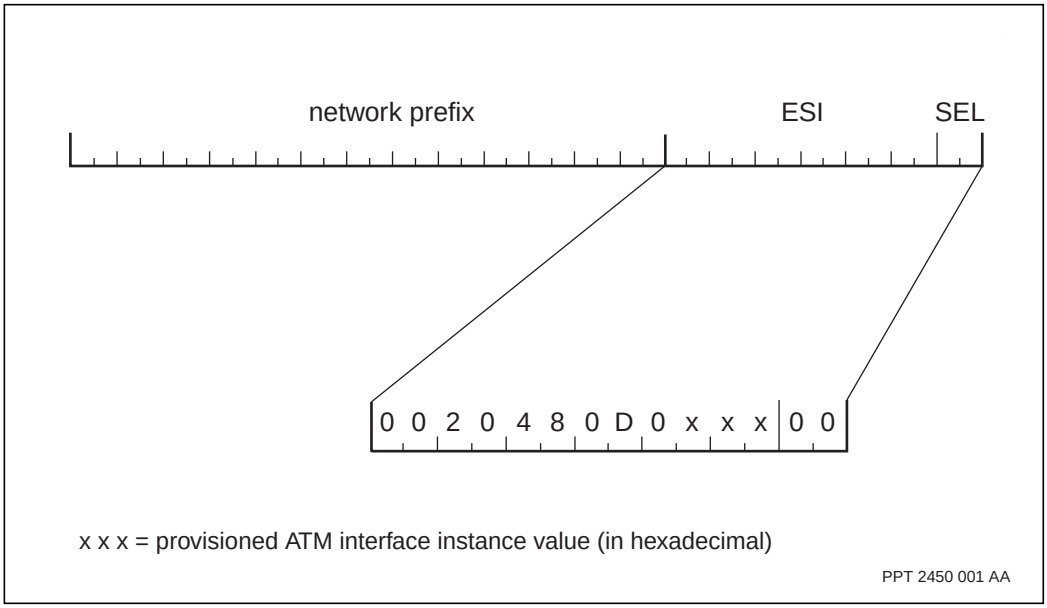

Static addresses

Static addressing is used on IISP and AINI interfaces to assign static routes. This type of addressing can also be used on UNI interfaces when dynamic addressing is not supported or when additional exception addresses need to be assigned. In static addressing, a list of addresses is assigned to an interface through manual provisioning. This list represents all the addresses that can be reached through the interface.

A provisioned static address can specify less than the whole 40-digit NSAP address. In this case, the address refers to all addresses that start with the specified digits. For example, if an IISP is provisioned with the static address 4751, it means that all addresses beginning with 4751 are reached through that interface.

Dynamic addresses

The NSAP address of an ATM device on the user side of a UNI interface can be dynamically registered with the network side of the interface at run time. Addresses are automatically registered when the interface comes up and remain registered as long as the interface remains up. New addresses can be added and existing ones removed at any time.

Dynamic address registration uses ILMI control procedures between the network and user sides of the interface. The network side provides the network prefix value. The user side responds by registering an NSAP address for the ESI value it supports, by appending the ESI value to the network prefix.

Native E.164 addresses

Although the Passport ATM networking system uses NSAP addresses within the network, it can also handle native E.164 addresses. When an incoming call setup protocol data unit (PDU) uses a native E.164 address as a destination address, it is converted to an NSAP-encapsulated address before being processed. The address field is restored to its native E.164 form as the call setup PDU leaves the Passport switch on the egress link.

Through provisioning, you can also specify that all outgoing call setup PDUs on an interface need to have NSAP or native E.164 destination addresses. In that case, the conversion is performed accordingly.

Group addressing

A group address for ATM anycast capability uses the same format as an individual ATM endpoint address. The authority and format identifiers (AFI) for group address adhere to standards defined by the ATM Forum in *User-Network Interface Signalling Specification Version 4.0* (af-sig-0061.000). See 241-5701-702 *Passport 7400, 15000, 20000 ATM Routing and Signaling Fundamentals* for more information on group addressing and anycast point-to-point connections.

ATM signaling

In ATM networking, signaling allows Passport switches to dynamically set up connections at run time. The UNI, IISP, AINI, and PNNI interfaces have a signaling channel that carries signaling protocol messages to control the setup and tear-down of connections. Basic ATM interfaces do not support signaling.

Signaling protocol

UNIs use the ATM Forum UNI signaling protocol (Version 3.0 or 3.1). IISP interfaces use the ATM Forum IISP signaling protocol version 1.0 (which is based on UNI signaling Version 3.0 or 3.1). The AINI signaling protocol is based on PNNI1.0 signaling. PNNI uses the ATM Forum PNNI signaling protocol version 1.0. The UNI signaling protocol is based on ITU-T Q.2931, which is the standard SVC signaling protocol adopted by the ATM Forum for signaling on UNI and IISP interfaces. Passport supports signaling for both point-to-point and point-to-multipoint connections.

Standard UNI and IISP signaling does not transport the end-to-end information that the network needs to set up an SPVC or SPVP. Therefore, a proprietary IE is added to the call setup PDU sent across each IISP signaling channel. This IE is proprietary, and is defined in the ATM Forum specification for PNNI signaling. The IE specifies the Virtual Path Identifier/Virtual Channel Identifier (VPI/VCI) of the destination point for the SPVC or SPVP.

Signaling VCC

Each UNI, IISP, AINI, and PNNI interface has a signaling channel, or VCC, that carries signaling PDUs across the interface. The PDUs are used to establish, maintain, and clear switched connections across the interface. The default signaling channel is VCC 0.5, but you can change the value to a different VCC through configuration.

Signaling process

Calls are routed on a hop-by-hop basis for nodes under UNI/IISP/AINI and through source routing for nodes under PNNI. For example, the Passport switch receives a call setup request across a UNI, IISP, AINI or PNNI interface. It then switches the request out on the UNI, IISP, AINI or PNNI interface that supports the address specified in the call request. Eventually, the call request leaves the last Passport node through a UNI, IISP, AINI or PNNI interface.

Signaling version interworking

Passport supports signaling version interworking functions that allow multiple signaling versions to coexist within a Passport ATM network. Interworking provides a way for the network to transport signaling PDUs generated on any interface over any other interface.

Passport provides the following signaling version interworking functions:

- between UNI Version 3.1 and PNNI Version 1.0
- between UNI Version 3.0 and UNI Version 3.1
- between UNI Version 3.0 and PNNI Version 1.0 (using a combination of the first two functions)
- between UNI Version 3.0 and UNI Version 4.0
- between UNI Version 3.1 and UNI Version 4.0
- between UNI Version 4.0 and PNNI Version 1.0
- between AINI Version 1.0 and PNNI Version 1.0

Interworking functions involving UNI 3.1 include the IISP 1.0 protocol.

ATM static routing

Call setup is a dynamic process. ATM networking over UNI, IISP, and AINI interfaces uses a scheme based on static call routing tables, rather than the dynamic exchange of routing information between switches.

Static call routing tables

Each switch in the network maintains a call routing table on the control processor (CP). This table contains all the ATM addresses associated with each UNI, IISP, and AINI interface on the switch.

The call routing table maps all the static and dynamic addresses to their corresponding UNI, IISP, or AINI interface. Default addresses are automatically added to the table, and static addresses are added as they are provisioned. Dynamic addresses are automatically placed in the table during the registration process. When an interface goes down, its associated addresses are removed from the table.

Address matching

An incoming call setup request is received over the signaling channel of a UNI, IISP, or AINI interface. The call routing table is then scanned for the specified called address. From the table, the switch selects the best matching address. This address is the table entry with the most hexadecimal digits identical to those of the called address. This process is called a maximal address match.

The call setup request is then forwarded to the next-hop UNI, IISP, or AINI interface associated with the best-match address. If no match exists, the call is cleared back to the previous node.

The figure “Address matching process” (page 175) shows a simplified example of the matching process. In this example, a call setup request with the following address enters the switch at UNI interface A:

```
473763111111111111222222222233333333333
```

Three IISP interfaces with mapped addresses are in the call routing table:

- IISP A handles all addresses starting with 473
- IISP B handles all addresses starting with 473763
- IISP C handles all addresses starting with 4737632211

In this case, the best-match process selects address 473763. The call is routed to interface IISP B.

Figure 55
Address matching process

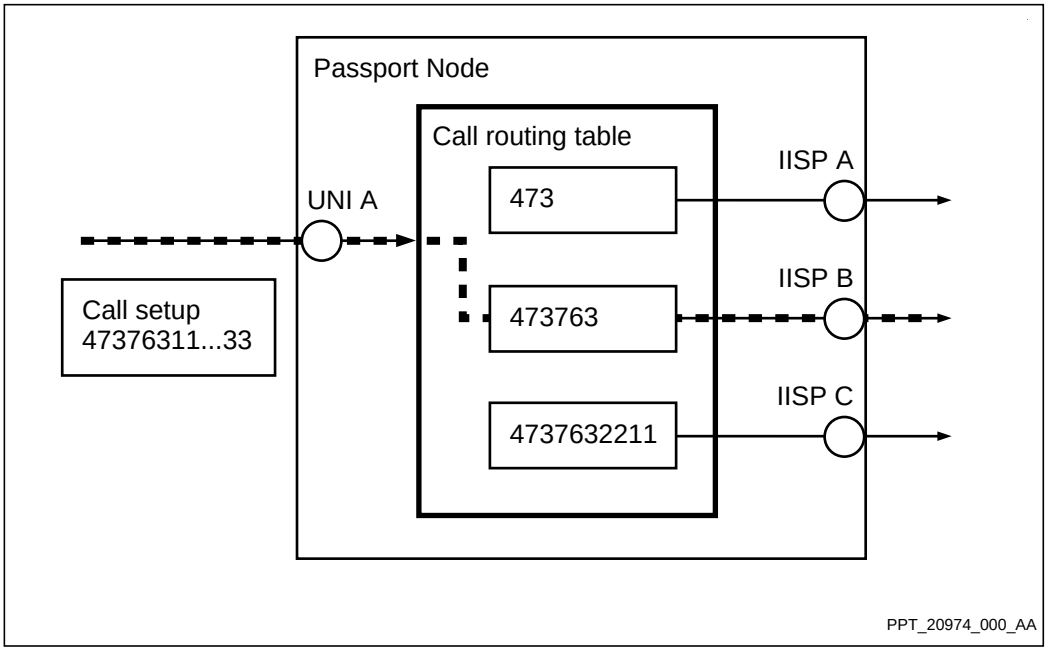

When a maximal address match produces a tie (when more than one interface handles the best-match address), the selected interface is chosen in a round-robin fashion. That is, the next time the same best-match address tie occurs, another interface is selected. This improves load sharing when multiple interfaces lead to the same routes, and also helps in not always re-selecting a bad route that leads to a failed network element.

When setting up a connection, the call routing function on a node is designed such that, when a maximal address match produces a tie, it tries all routes. Each address has two lists of ATM interfaces. The primary list is tried first and then, if necessary, the alternate list is tried. The starting point in the primary list changes for every new call setup request in round-robin fashion. Once the primary list is exhausted, only then is the alternate list tried using the same round-robin rule.

Route failure under UNI, IISP, and AINI

When all the possible paths from a node in the route fail, the call setup request is sent back to the previous node. At this node, all other paths are attempted until a successful one is found. This process, which is a PNNI concept, is known as a crank-back mechanism.

The figure "Crankback mechanism (static routing)" (page 176) shows an example of the crank-back mechanism for UNI-, IISP-, and AINI-based routing. In this example, User A originates a call for User B. Although the call setup fails at node 6 and node 5, the connection is still successfully routed through other nodes to the destination.

Figure 56
Crankback mechanism (static routing)

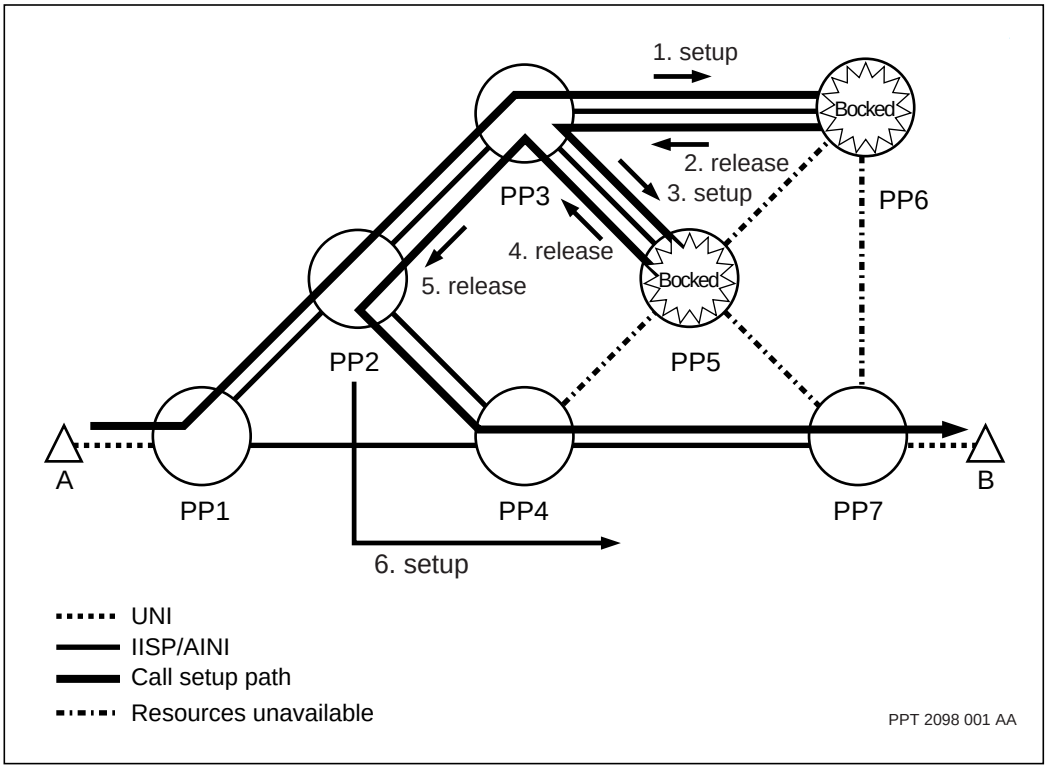

ATM dynamic routing

Nodes in a PNNI network exchange topology and link state information on an ongoing basis. In this way, nodes maintain an up-to-date view of the state of the network. This approach differs from static routing protocols like IISP, which require manual provisioning for reachable addresses across interfaces.

For more information on ATM dynamic routing, see 241-5701-702 *Passport 7400, 15000, 20000 ATM Routing and Signaling Fundamentals*.

Topology and link state maps

In a PNNI call set-up, the Passport uses the topology and resource information acquired from its neighbors to find the best route for the call. This can include the exchange of topology and logical link information in an hierarchical PNNI network.

Passport determines a route to the node that is advertising the longest prefix match for the destination address. This route satisfies the QoS requirements specified in the call set-up request. The route selection process also takes into consideration the user-specified optimization criteria (for example, giving preference to routes with minimal administrative weight).

Address matching

When an incoming call set-up request enters a node through UNI or IISP interface, Passport performs a longest prefix match of the destination address with all addresses supported on UNIs and IISP interfaces and with all of the addresses supported by PNNI nodes. It is possible to have several choices with the same best prefix match. Passport resolves the choice depending on the interface type and the quality of the address match. As a rule, Passport selects a UNI or IISP interface first, and then a PNNI.

Route setup failure under PNNI

If a node blocks the call set-up request before the request reaches its destination, the call cranks back to either the source node or to the entry border node that generated the DTL for that level or for a numerically higher level. This mechanism is not a crankback as in IISP-based networks. Instead, the node receiving the cranked-backed call reissues the call setup with a new calculated route. If call setup times out, the source node releases the call setup.

The figure “Crank-back mechanism (dynamic routing under PNNI)” (page 179) shows an example of the crankback mechanism for PNNI-based routing in a hierarchical network. For more information on crankback, see 241-5701-702 *Passport 7400, 15000, 20000 ATM Routing and Signaling Fundamentals*.

Figure 57
Crank-back mechanism (dynamic routing under PNNI)

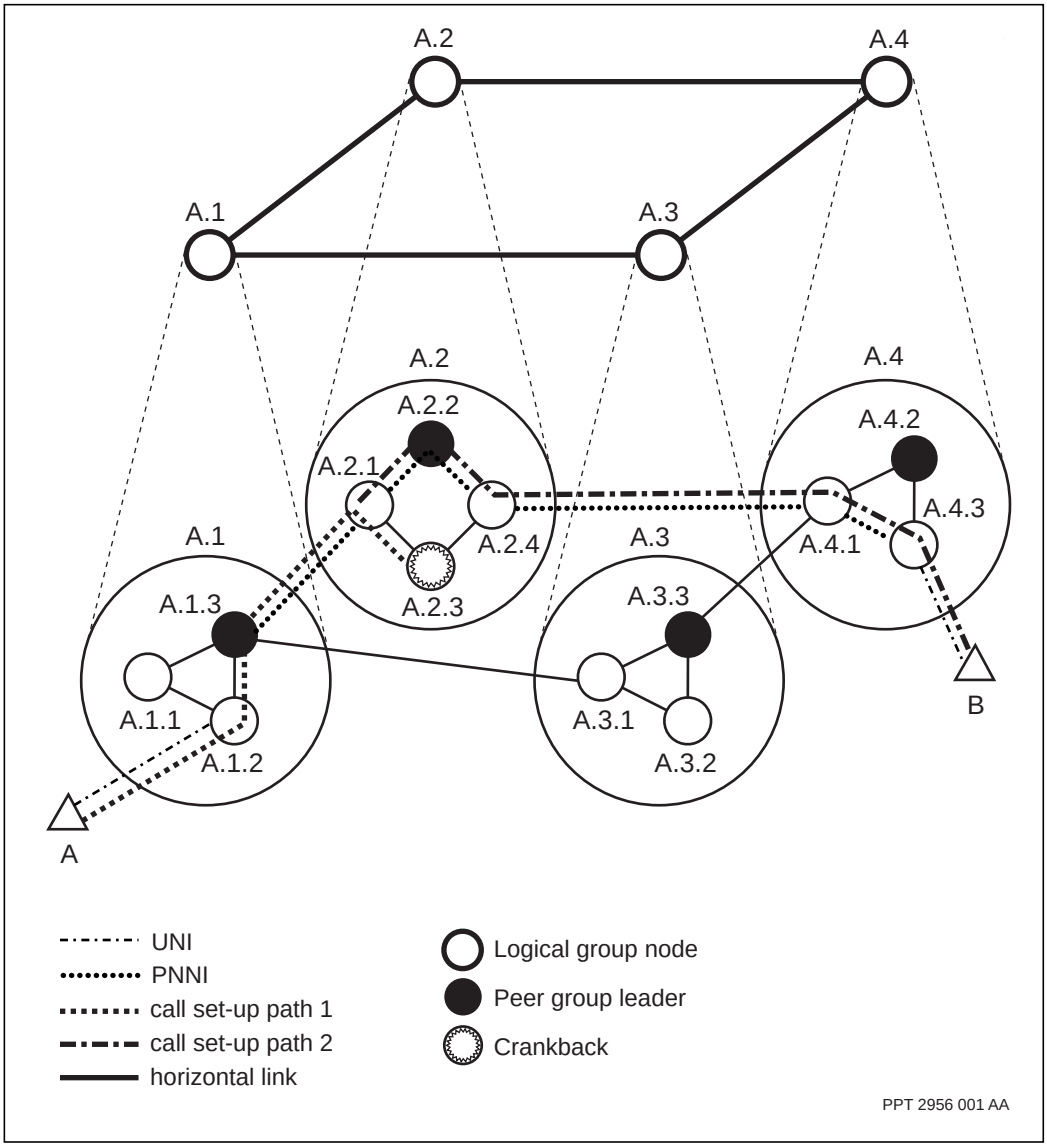

EBR connection recovery and path optimization under PNNI

Edge-based rerouting (EBR) capabilities operate in a PNNI network automatically to recover and optimize existing point-to-point SVC, SVP, SPVC, or SPVP connections. EBR implements two types of rerouting procedures: connection recovery and path optimization. When network failures occur, EBR allows connection recovery by reestablishing the connection without intervention from the end systems. In the connection recovery process, EBR finds an alternative route for a connection that would otherwise be cleared back to the originator.

EBR also operates by moving active connections to more optimal PNNI paths in conditions when no failure has occurred. The path optimization aspect of EBR preserves the QoS of a point-to-point connection and provides more efficient use of network resources.

ATM routing components

The following paragraphs describe the major components in the ATM routing system. For more information on components, see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

AtmIf

You provision an *AtmIf* component for each ATM interface. The *AtmIf* component manages the set of ATM connections that terminate at or traverse the interface.

AtmIf Uni

This component represents an ATM UNI. If you provision this component under an *AtmIf* component, you cannot provision the *Iisp*, *Aini*, or *Pnni* components under the same *AtmIf*.

AtmIf Iisp

This component represents an ATM IISP interface. If you provision this component under an *AtmIf* component, you cannot provision the *Uni*, *Aini*, or *Pnni* components under the same *AtmIf*.

AtmIf Aini

This component represents an ATM Aini interface. If you provision this component under an *AtmIf* component, you cannot provision the *Uni*, *Iisp*, or *Pnni* components under the same *AtmIf*.

AtmIf Pnni

This component represents an ATM PNNI interface. If you provision this component under an *AtmIf* component, you cannot provision either of the *Uni* or *Iisp* components under the same *AtmIf*. To implement ATM networking, you must provision one of the *Uni*, *Iisp*, *Aini*, or *Pnni* components under the *AtmIf* component. If you do not provision one of these components, Passport defines the interface as basic.

AtmRouting

This component contains the provisionable and operational attributes related to ATM PNNI routing. If the *AtmCallRouter* component is provisioned on the node, the *AtmRouting* component is added in its place. For more information on the *AtmCallRouter* component, see 241-5701-060 *Passport 7400, 15000, 20000 Components* and 241-5701-710 *Passport 7400, 15000, 20000 ATM Configuration Guide*.

Chapter 7
Passport addressing and call routing

This section provides further background on addressing plus additional call routing examples to complement the Passport routing system chapters. The topics covered are as follows:

For more detailed information on addressing and call routing, see the individual routing system chapters, and "Call routing references" (page 20).

Role of addressing

An address uniquely identifies a network entity. This identifier may be unique within or outside a network, depending on the addressing plan. The network entity may be a Passport node, DPN-100 module, work station, telephone, service port, or other network device.

An important aspect of the addressing function is to identify and resolve any differences between address formats. Routers and servers translate external network addresses to Passport internal addresses. This address resolution allows connections to be established and packets routed.

Addressing features

The Passport system has several addressing features.

Compliance to external addressing plans

The Passport system supports addresses complying to the following standard addressing plans and formats:

- E.164
- X.121
- Internet Protocol (IP)
- Media Access Control (MAC)
- Network Service Access Point (NSAP)

Note: Data network address (DNA) is a Passport term for an address that identifies a device in a Passport network. DNAs conform to X.121 and E.164 numbering systems.

Support of Passport internal identifiers

The Passport system uses the following internal identifiers, which are used to optimize call establishment and packet routing:

- RID
- MID
- node identifier (nodeId)
- node name

See the figure “External addressing plans and Passport internal identifiers” (page 185) for the addressing plan usage for several network entities.

Figure 58
External addressing plans and Passport internal identifiers

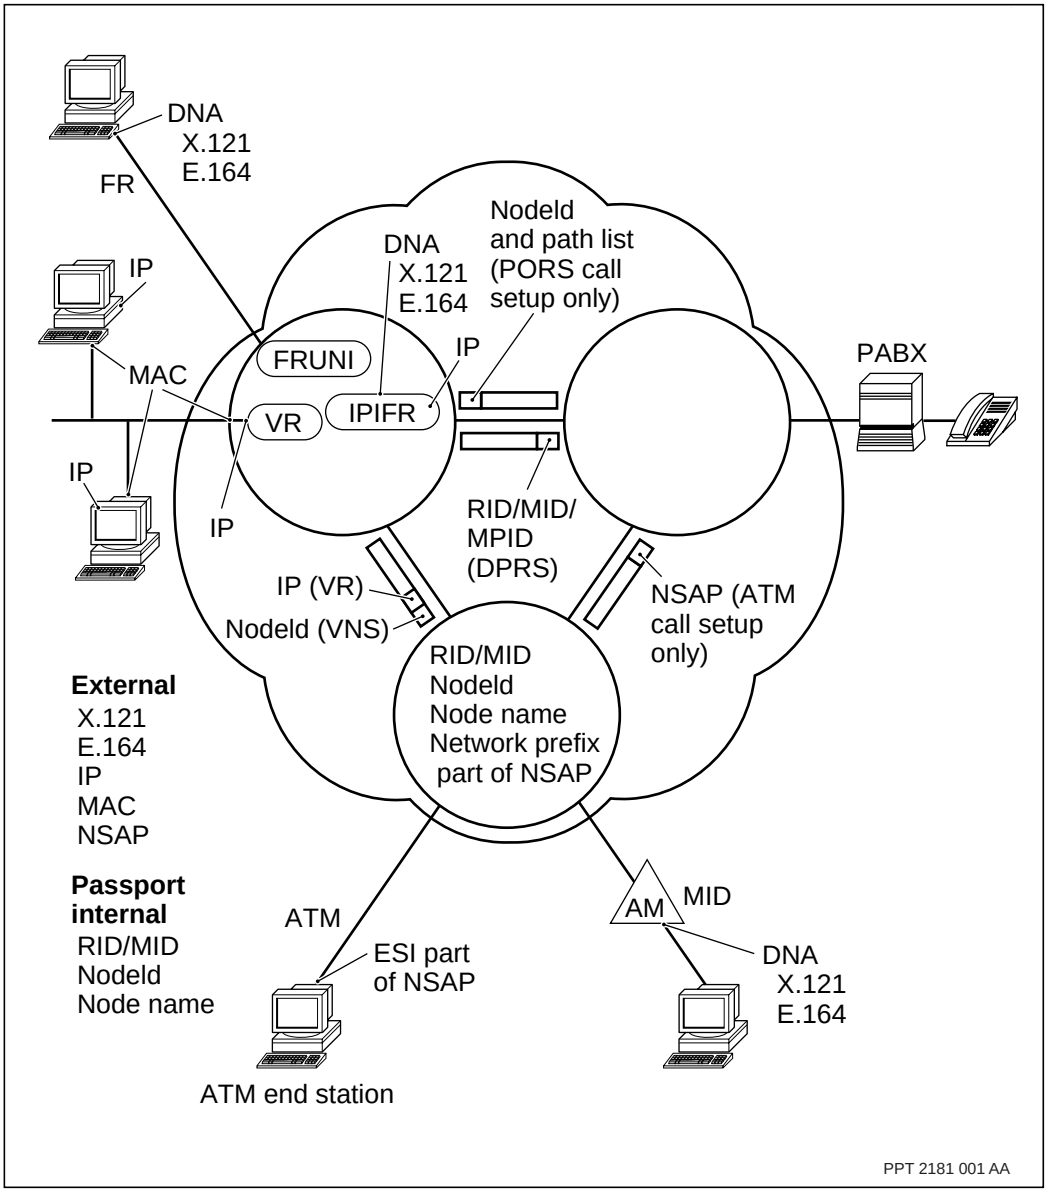

Addressing mechanisms

This section describes the following addressing mechanisms:

- "External addressing plan formats" (page 186)
- "Passport internal identifiers" (page 191)

External addressing plan formats

An addressing plan, also known as an addressing scheme or numbering plan, defines a format for identifying an endpoint. The Passport system supports the most common external addressing plans recognized by ITU-T and IETF.

These external addressing plan addresses represent

- entities on a switch that are reachable by an application
- custom access lines (for example frame relay)

Passport services using the Dynamic Packet Routing System (DPRS) can accept the X.121 and E.164 addressing format for both international and national addresses. IP services use IP and MAC addressing formats. PORS services use X.121, E.164, and NSAP addresses. ATM services accept E.164 and NSAP addresses. See the table "External address information" (page 186).

Table 7
External address information

Name	Length	Example	Origin	Use
X.121	14 digits maximum	30211500510001	provisioned by network operator	DPRS, PORS
E.164	15 digits maximum	0116137633500	provisioned by network operator	DPRS, PORS, ATM
IP	32 bits	47.208.133.76	provisioned by network operator	IP

(Sheet 1 of 2)

Table 7 (continued)
External address information

Name	Length	Example	Origin	Use
MAC	48 bits	00-20-1B-AB-CD-00	assigned by hardware vendor	IP
NSAP	40 hex digits	4737631111111111112222 222222233333333333	provisioned by network operator	PORS, ATM
(Sheet 2 of 2)				

X.121 addressing plan format

X.121 is defined as the ITU-T standard describing an addressing scheme used in frame relay and X.25 networks or for public switched data networks (PSDN). X.121 addresses are sometimes called IDNs (International Data Numbers). See the figure "X.121 addressing plan format" (page 187) for the X.121 address format.

Figure 59
X.121 addressing plan format

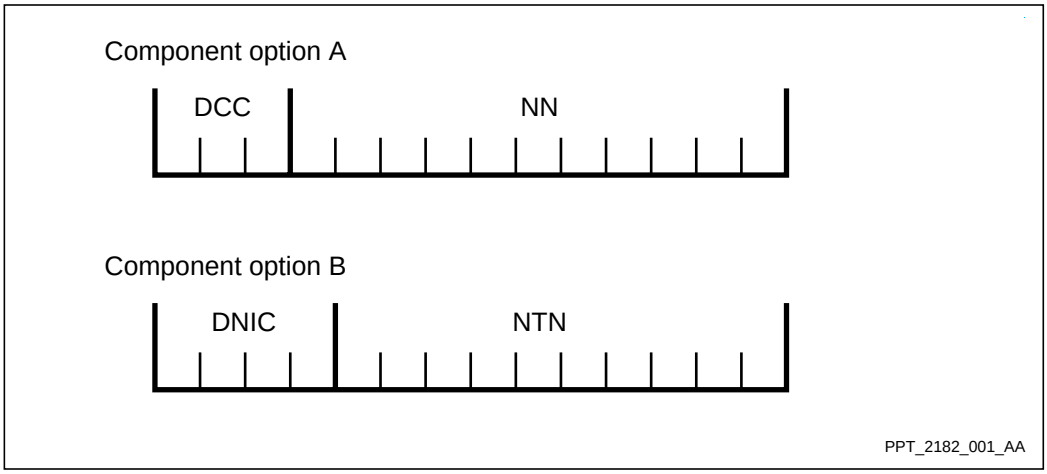

Passport implements the X.121 formats in accordance with ITU-T standards.

The full X.121 international address must be less than or equal to 14 digits in length. X.121 addresses are globally unique.

An international address consists of one of the following components:

- a data country code (DCC) of exactly three digits and a nationally defined national number (NN) of up to 11 digits
- a data network identification code (DNIC) of exactly four digits (a 3-digit DCC followed by a 1-digit network digit), and a nationally defined network terminal number (NTN) of up to ten digits. DNICs are assigned by the ITU-T standards central authority.

The national address is comprised of either the NN or the NTN.

Refer to the X.121 standard for more details.

E.164 addressing plan format

The telephone system, also known as the public switched telephone networks (PSTN), uses the E.164 addressing plan. E.164 is the ITU-T recommended numbering plan for ISDN. E.164 addresses are globally unique.

All of the Passport access services can accept the E.164 addressing format for both international and national addresses. See The figure “E.164 address format” (page 188) for the E.164 address format.

Figure 60
E.164 address format

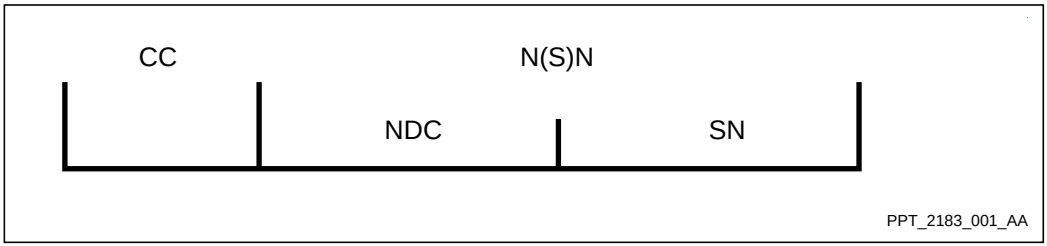

Passport implements the E.164 formats in accordance with ITU-T standards.

The full E.164 international address must be less than or equal to 15 digits in length, including the escape digit.

An international address consists of the following variable length components:

- a country code (CC) to select the destination country
- a national destination code (NDC), or area code
- a subscriber number (SN)
- the national (significant) number (N(S)N) comprises the NDC and SN

Refer to the E.164 standard for more details.

IP addressing plan format

The network administrator assigns a globally unique IP address to all workstations and router interfaces in an IP network. This addressing plan supports multiple routing protocols such as OSPF and RIP. See the figure "IP address format" (page 189) for the IP address format.

Figure 61
IP address format

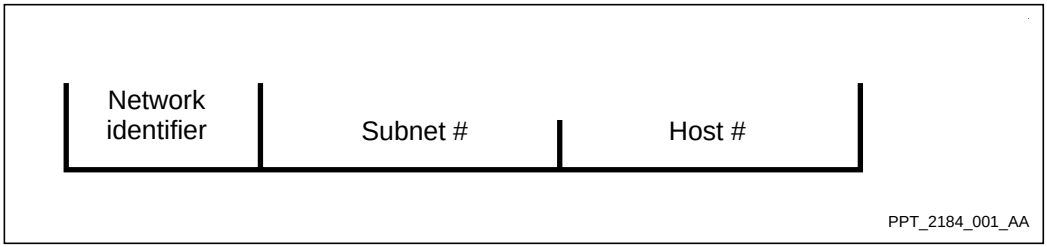

Passport nodes implement the IP version 4 address formats in accordance with IETF standards.

The IP address is 32 bits, or 4 bytes, in length. The bytes are separated by a period and represented by decimal numbers.

An IP address consists of the following components:

- A network identifier, also known as net Id and network field, is assigned by a central authority.

- A subnet number represents one or more networks, and is assigned by the network operator. The subnet-number length depends on the size of the represented network.
- A host number represents the station, node, or other device, and is assigned by the network operator.

IP addresses use the concepts of classes and masks to create subnetworks. These subnets provide flexibility to the operator in designing the network. For information about IP address classes and subnets using masks, see the addressing information in 241-5701-805 *Passport 7400, 15000, 20000 Understanding IP*.

Media Access Control (MAC) addressing plan format

MAC addresses, also known as data link layer addresses, are used in delivery of data to a destination on a local area network. Every IP device interface has a MAC address which is globally unique. This address is the device's physical address.

A MAC address consists of 48 bits, represented as a hexadecimal string, and used between end systems. The figure "MAC address format" (page 190) shows the MAC address format.

Figure 62
MAC address format

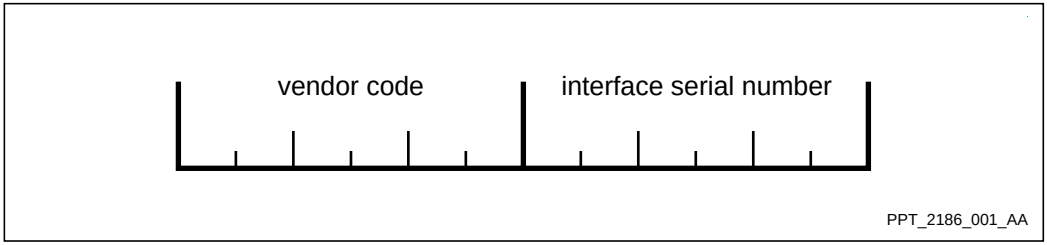

Each Passport node is assigned a unique set of MAC addresses consisting of a base address and a range of addresses. The addresses are generated by an algorithm. For each logical processor (LP) a set of 64 MAC addresses is assigned. Each Passport node has 1024 MAC addresses (64 addresses for each of the 16 LPs).

MAC addresses are assigned to all the LPs on a shelf, regardless of the LP type. Ethernet LAN components have MAC addressing on their interfaces. The MAC addresses cannot be changed through provisioning.

NSAP addressing plan format

OSI Network Service Access Point (NSAP) addressing plan exists for PORS and ATM traffic. The figure "NSAP address format" (page 191) shows the NSAP address format.

Figure 63
NSAP address format

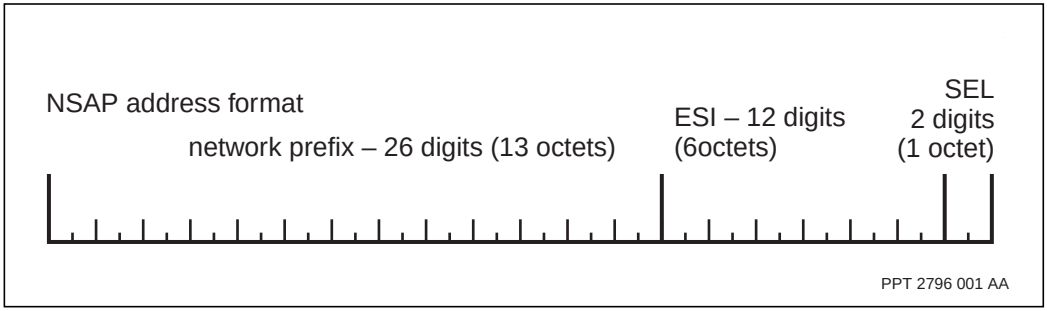

Passport ATM routing supports all three NSAP address formats: Data Country Code (DCC), International Code Designator (ICD), and E.164.

The NSAP address format consists of three main parts:

- a 26 hexadecimal digit prefix, known as the network prefix
- a 12 hexadecimal digit identifier, known as the end system identifier (ESI)
- a 2 hexadecimal digit selector (SEL)

Refer to the ATM Forum UNI standard, section 5.1.3 Addressing for more details.

Passport internal identifiers

The Passport system supports external addressing plans by translating them when necessary to proprietary or internal identifiers for Passport network use.

Passport networks and nodes have several identifiers (see the table “Passport internal address information” (page 192)). The identifiers are provisioned by the network operator. Each identifier is used by the routing system, and in some cases for network management too.

Table 8
Passport internal address information

Level	Name	Length	Example	Use
Network	RID	7 bits	15	DPRS for routing
Node	MID	11 bits	275	DPRS for routing
	nodeId	32 bits	1244	PORS for call setup; and network management
	node name	12 character ASCII string	EM/Toronto	mnemonic for a node. Synonymous with nodeId within a region. PORS for call setup; network management.

Note: Network management is outside the scope of this NTP.

RID/MID

DPRS uses the hierarchical routing identifier/module identifier (RID/MID) addressing system. For details on the RID/MID Passport internal addressing plan see “DPRS addressing system” (page 94).

NodeId

The node identifier (nodeId) is an internal concept within base routing. NodeIds are stored in the topology database and used by the topology process to select the best routes to each nodeId for the DPRS and PORS routing systems. These routes are placed in the routing system’s forwarding tables. For details on the nodeId Passport internal addressing plan see “PORS addressing system” (page 130).

Node name

PORS BTDS, HTDS, and voice service use the node name addressing system for call set up. Node name is synonymous with nodeId and is used to determine nodeId. For details on the node name Passport internal addressing plan see “PORS addressing system” (page 130).

Role of call routing

At a high level, all Passport routing systems have the same concept of address resolution. IP makes reference to address resolution only. Call routing is a term associated with DPRS, PORS, and the ATM Routing System.

Call establishment sets up the connection between the end points of two access services across a network to enable data transfer. Call routing is one phase of call establishment. Call routing involves address resolution and the routing of call request or call setup packets between the end points. The address resolution aspect of call routing translates external address formats to Passport internal identifiers.

Some other aspects of call establishment are creating the call request and call accept packets, routing the call accept packet, call setup failure, and rerouting. See the individual routing chapters and routing NTPs for further information on these topics.

Passport supports call servers and call routers

DPRS uses a call router and PORS uses a call server. A call router receives the whole call request packet. First the call router performs an initial address lookup to begin translating the DNA to a physical network identifier. Next the call router updates the information in the call request packet. The call router now forwards the packet to the next entity (node or process) to resolve the Passport internal address further.

A call server accepts the external address, translates it to an identifier, then returns the information to the requestor. Within the Passport system the call server is rarely accessed again for the duration of that call. The call setup packet with the Passport identifier and route information is forwarded on to the destination end point.

ATM calls are routed on a hop-by-hop basis, based on the ATM Forum standard. At each hop the called address (NSAP) is matched against the list of NSAP prefixes covered by each ATM link. The call is forwarded out the ATM link that has the longest prefix match.

Call routing features

A Passport network supports call routing-type features for all of the Passport routing systems:

- Dynamic Packet Routing System (DPRS)

  DPRS provides call routing options such as throughput or delay RCOS, call forwarding, call redirection, trace, and hunt groups.

- Path-Oriented Routing System (PORS)

  PORS provides many call routing transport characteristics consisting of cost or delay, setup priority, bump preference, required bandwidth, security, traffic type, and Passport trunk type.

- IP routing service

  IP provides address resolution only.

- ATM Routing System

  ATM provides many call routing QoS levels consisting of constant bit rate (CBR), real time and non-real time variable bit rate (RT-VBR, NRT-VBR) and unspecified bit rate (UBR) connections. Passport ATM also provides a hop-by-hop crankback mechanism.

DPRS call routing

DPRS call routing refers to address resolution and call request packet routing for DPRS services. With DPRS, the endpoints are identified by an X.121 or E.164 DNA. The address resolution involves mapping the DNA assigned in a Passport network to its corresponding RID, MID, MPID for call establishment by either a CSRM or a Passport call router. See the figure “External addressing plans and Passport internal identifiers” (page 185) for the DPRS addressing plan usage.

Figure 64
External addressing plans and Passport internal identifiers—DPRS focus

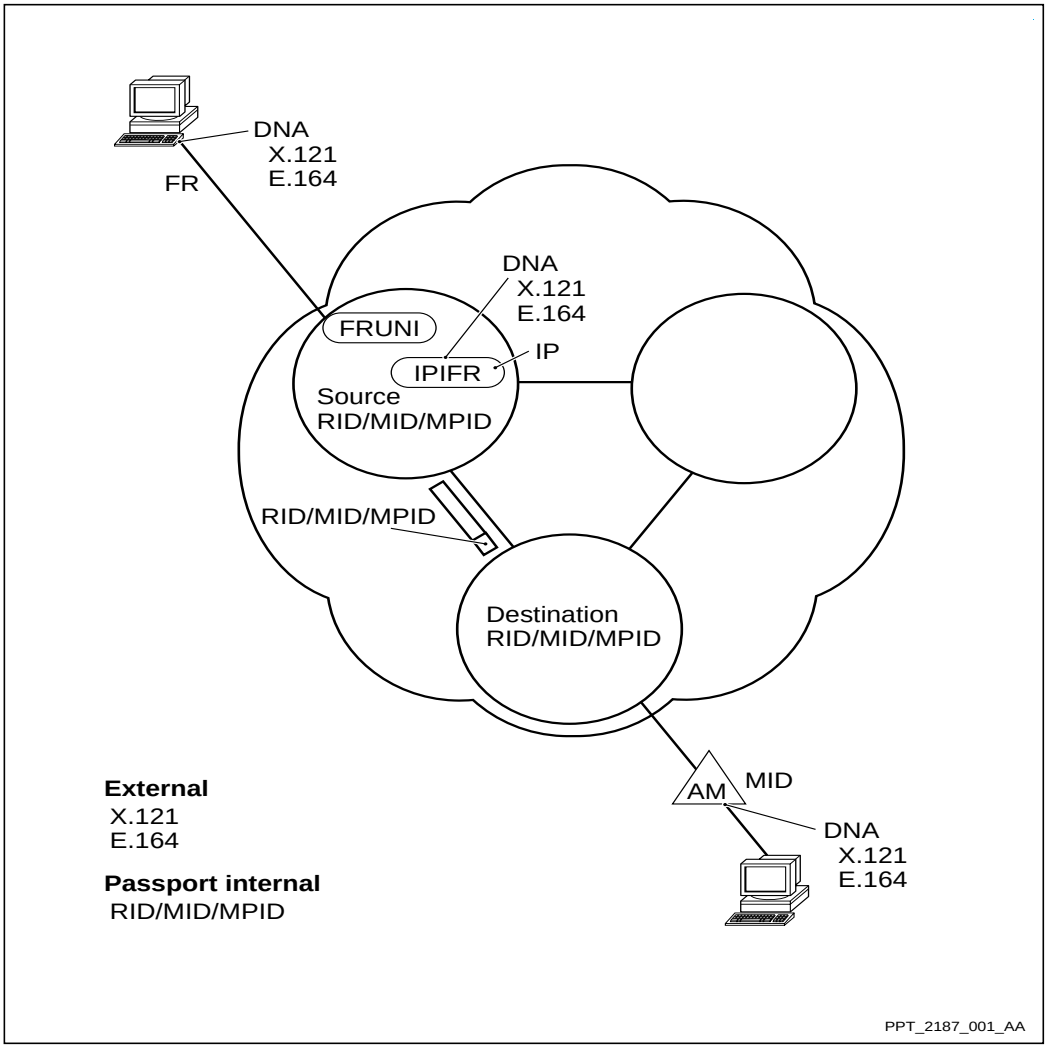

CSRM used in a mixed Passport and DPN-100 network

The Call Server Resource Module (CSRM) provides call routing services for both the Passport nodes and DPN-100 modules connected to a Passport network. The CSRM supports advanced networking services like call forwarding and call redirection.

In a network with CSRMs

- a source call router (SCR) identifies the destination RID
- a destination call router (DCR) identifies the destination MID

DPRS call establishment—using CSRMs in a mixed Passport and DPN-100 network

The call routing process using CSRMs is described in the DPRS chapter (see the figure "Call routing and DPRS—call establishment" (page 105)).

Passport call router used in a Passport-only network

A Passport network that does not include DPN-100 modules uses call routers. Each Passport node may be provisioned with a call router to handle the call routing functions. This functionality eliminates the need for a CSRM, thereby simplifying the management of network devices.

A call router uses the DNA to determine the RID and MID of the destination DNA.The RID/MID address information is placed in the call packet header for forwarding the call to the destination address end point.

Optionally, a Passport-only network can include call redirection servers. Call redirection servers (CRS) direct failed call attempts to alternative destinations. A CRS has a database of primary addresses mapped to alternative addresses or RID/MID locations. When a destination cannot be reached, the CRS redirects the call to an alternative location. Fore more information on call redirection servers, see 241-5701-410 *Passport 7400, 15000, 20000 Call Redirection Server Guide*.

DPRS call establishment—Passport call router, multiple RID subnet

The following sequence describes DPRS call establishment when using Passport call routers in a multiple RID subnet (see the figure "DPRS call establishment—with Passport call routers in a multiple RID subnet" (page 198)):

1 The Passport virtual circuit (VC) process creates a call request packet which includes the data network address (DNA), RID (15), MID (3) and MPID (x) for the source and the DNA for the destination. The packet is forwarded to the node's Call Router (CR) to determine the destination address.

2 In multiple RID subnets the CR uses its provisioned prefix-DNA to map the DNA of the destination to the RID. The RID (16) of the destination is placed in the call request packet which is then forwarded to the closest Passport in the destination RID subnet.

3 The call request packet is delivered to the CR of the closest Passport in the destination RID subnet. The CR uses its provisioned prefix-DNA to map the DNA to the destination MID (6). The call request packet is then sent to the destination MID (6).

4 At the destination Passport MID (6), the local call routing system first determines which process supports the DNA, then forwards the call request packet to that process, for example frame relay VC. That process learns the source RID/MID/MPID from the call request packet. The process places the RID (16), MID (6) and MPID (z) of the destination in a call accept packet and sends it back to the source, by the shortest available route.

5 Data transfer can begin because both ends know each others RID, MID and MPID addresses.

Note: The call establishment steps for a single RID subnet are reduced from the multiple RID subnet. In step 2 and 3, the CR uses its provisioned prefix-DNA to map the DNA of the destination to the RID. Next the CR uses its provisioned prefix-DNA to map the DNA of the destination to the MID. The RID and MID of the destination are placed in the call request packet which is then sent to the destination MID. The process then follows from step 4.

Figure 65
DPRS call establishment—with Passport call routers in a multiple RID subnet

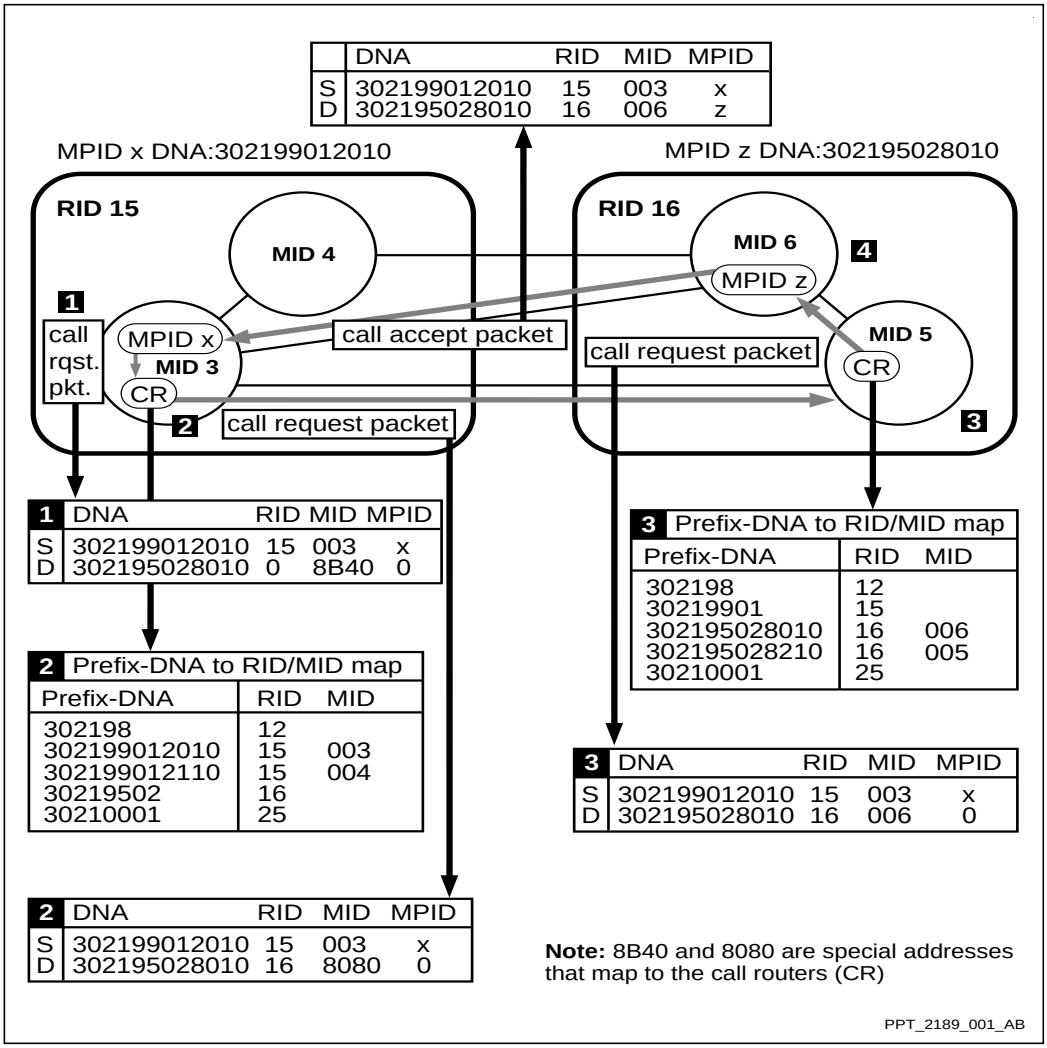

IP address resolution

The Passport system uses the standard IP Address Resolution Protocol (ARP) for its IP address resolution. See 241-5701-805 *Passport 7400, 15000, 20000 Understanding IP* for more information.

PORS call routing

PORS call routing refers to address resolution and call setup packet (also referred to as path setup packet) routing for the PORS services. PORS accepts identifiers in the form of node names and nodeIds. The endpoints are identified by a node name and component name for voice and transparent services. The addresses are translated to node identifiers (nodeId) and process identifiers (PID), respectively, for call establishment. For ATM services, PORS uses NSAP addressing, which it converts to PORS addressing plan identifiers. See the figure “External addressing plans and Passport internal identifiers—PORS focus” (page 200) for the PORS addressing plan usage.

Figure 66
External addressing plans and Passport internal identifiers—PORS focus

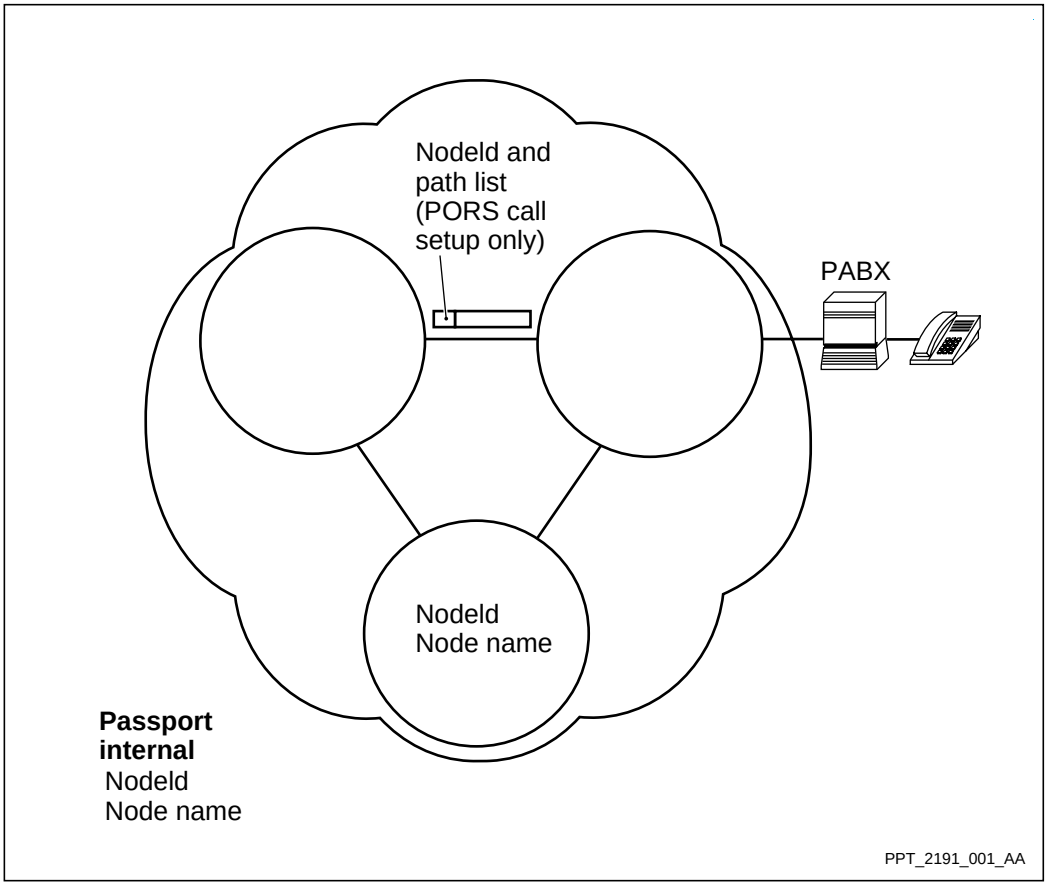

When a call is placed via PORS, a route and a logical connection (either permanent or switched, depending upon the application) are created between the source and destination Passport nodes. The address is no longer used by PORS which now uses the established path to forward packets.

In addition, PORS is capable of routing on DNA and NSAP-based addresses when the originating services uses one of these methods of addressing. For example, the CES service uses NSAP-based addressing. When it routes on DNA or NSAP addresses, PORS translates the address into a Passport node identifier.

To route on a DNA address, PORS uses the DPRS call establishment infrastructure to complete the call, so the same provisioning procedures are needed as for DPRS. (See “DPRS call routing” (page 194).)

To route on an NSAP address, PORS uses ATM routing. (With ATM, the endpoints are identified by an OSI Network Service Access Point (NSAP) address. Each UNI, IISP interface, and AINI has one default address, and may have other associated addresses. UNI interfaces can have static (provisioned) or dynamic addresses. IISP interfaces and AINIs can only have static addresses.) To support this method of addressing, the network operator must provision a network identifier for each node. (See the section on provisioning NSAP addressing with PORS in 241-5701-435 *Passport 7400, 15000, 20000 Path-Oriented Routing System Guide*. See the section on provisioning ATM routing in 241-5701-710 *Passport 7400, 15000, 20000 ATM Configuration Guide*.)

PORS call establishment—voice transport call example

The following sequence describes PORS call establishment for permanent virtual circuits (PVCs) when performing a Passport voice transport call (see the figure “PORS call establishment — voice example” (page 202)):

1 The PORS VC requests a route from the Route Selector (RS) between source node name (EM/A) and destination node name (EM/C).

2 The RS performs a mapping of the destination node name (EM/C) to the destination nodeId (1236). The RS then selects a route to that nodeId using the topology database.

3 The source sends a call setup packet along the Passport trunks specified in the selected route. The packet contains the source and destination node and component names, the nodeIds, as well as the route information. At each Passport node along the route, the PAs set up the forwarding tables.

4 When the call setup packet reaches the destination, the PORS FDCR maps the component name (VS/3) to the PID (7).

5 The destination returns a call setup confirmation packet to the source. The call setup confirmation packet stops at all Passport nodes along the route.

6 The call setup confirmation packet reaches the source and the route (a permanent or switched logical connection, depending upon the application) is enabled.

Note: PVC and SVC call establishment are very similar.

Figure 67
PORS call establishment — voice example

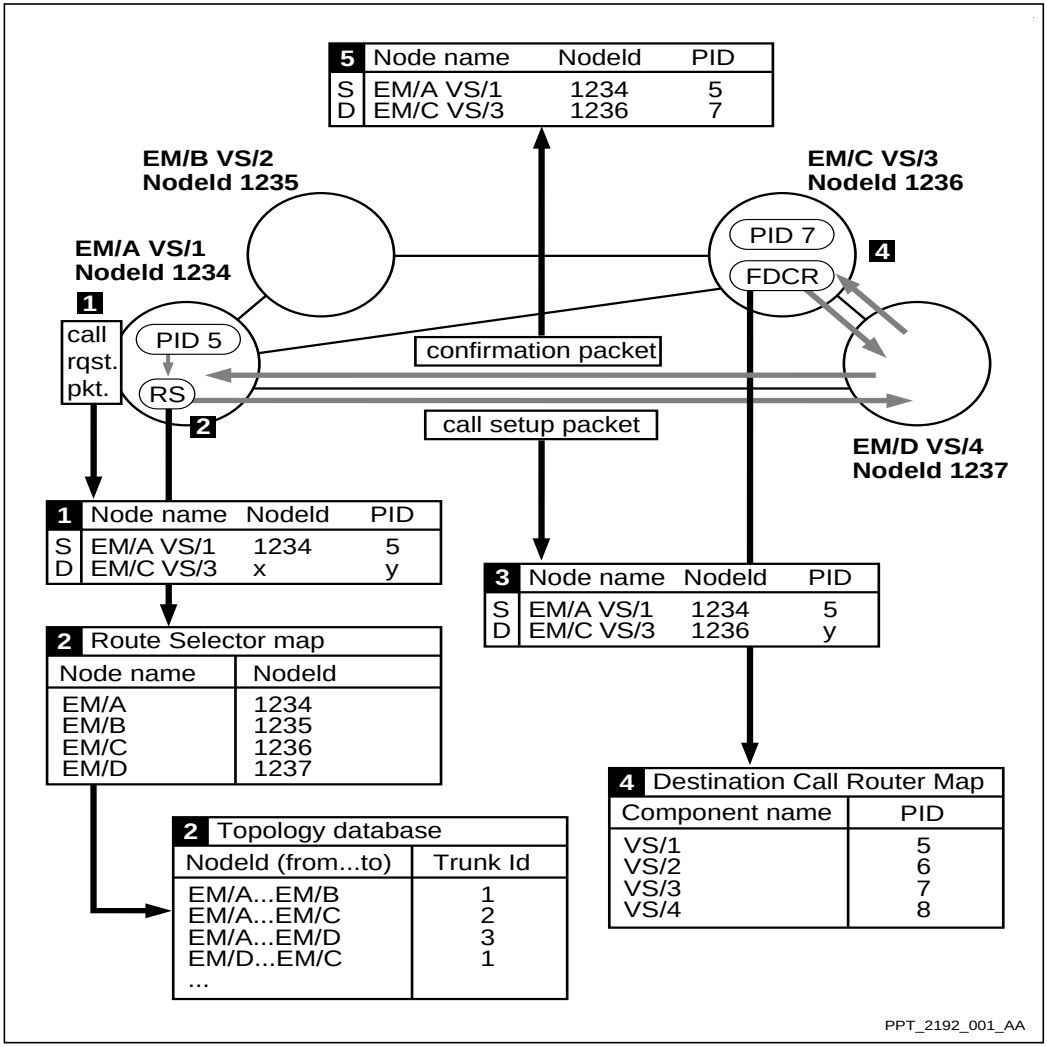

ATM call routing

ATM call routing refers to address matching and routing call setup request packets for ATM services.

For networks under UNI 3.x, IISP 1.0, and AINI nodes route calls on a hop-by-hop basis. The Passport nodes receive call setup requests over the signaling channel of a UNI, IISP interface, or AINI. The call router searches the call routing table on the control processor for the UNI, IISP interface, or AINI address that supports the called address. The call router selects the best address match. The call setup request is then forwarded to the next-hop UNI, IISP interface, or AINI. The ATM call routing process is described in detail in “ATM static routing” (page 173).

For networks under PNNI 1.0, nodes use source routing. Before undertaking transport of cells, the node establishes a connection that satisfies QoS requirements. Nodes in a PNNI network are organized into groups and advertise node and link state parameters. This arrangement is the basis for source routing.

Passport 7400, 15000, 20000

Networking Overview

Release 5.2

Publication: 241-5701-400
Document status: Standard
Document version: 5.2S1
Document date: November 2003
Printed in Canada

NORTEL NETWORKS