NORTEL NETWORKS

Passport 7400, 15000, 20000

Dynamic Packet Routing System Guide

241-5701-425

Passport 7400, 15000, 20000

Dynamic Packet Routing System Guide

Publication: 241-5701-425
Document status: Standard
Document version: 5.2S1
Document date: November 2003

Printed in Canada

Publication history

November 2002

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

Contents

List of figures

List of tables

About this document

This guide describes the Dynamic Packet Routing System (DPRS).

Who should read this document and why

This guide is for persons who perform the following tasks for a Passport network with DPRS:

- planning
- engineering
- installing and configuring
- provisioning
- operating and maintaining
- troubleshooting

What you need to know

This guide assumes that you understand the Passport network architecture and the basics of network routing. You can learn more about the product by reading 241-5701-030 *Passport 7400, 15000, 20000 Overview*.

For basic information on routing systems, see 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*.

How this document is organized

The 241-5701-425 *Passport 7400, 15000, 20000 Dynamic Packet Routing System Guide*, contains the following information:

- "Overview" (page 19) presents an overview of the routing system.
- "DPRS routing" (page 27) explains how DPRS routes data traffic.
- "Traffic management" (page 53) describes the DPRS traffic management features.
- "Network planning and engineering" (page 75) provides the information you need to plan the DPRS network.
- "DPRS configuration" (page 107) describes how to provision DPRS.
- "DPRS monitoring" (page 125) provides information you can use in maintaining DPRS.

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 sections "Hierarchical addressing" (page 25), "RID subnets" (page 75), and "Guidelines for RID subnets" (page 78) 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 pathname refers to the full specification of a path starting from the root directory. Absolute pathnames always begin with the slash (/) symbol. A relative pathname takes the current directory as its starting point, and starts with any alphanumeric character (other than /).

Related documents

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

See the following Passport documents for information related to DPRS:

- 241-5701-600 *Passport 7400, 15000, 20000 Configuration Guide*
- 241-5701-050 *Passport 7400, 15000, 20000 Commands*
- 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*
- 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*
- 241-5701-420 *Passport 7400, 15000, 20000 Trunking Guide*
- 241-7401-110 *Passport 7400, DPN-100 Interworking Guide*
- 241-5701-060 *Passport 7400, 15000, 20000 Components*
- 241-5701-500 *Passport 6400, 7400, 15000, 20000 Alarms*

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
Overview

For an overview of the Dynamic Packet Routing System (DPRS) see the following sections:

For basic information on network routing, types of routing systems, and DPRS, see 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*.

What is DPRS?

DPRS is an efficient, connectionless routing system for delay-sensitive and high-throughput traffic. DPRS is ideally suited to carry data traffic such as frame relay. DPRS also handles the supported DPN-100 services on Passport 7400 series switches. Passport 15000 and 20000 do not use DPN-100 services. (For information about interworking with DPN-100, see 241-7401-110 *Passport 7400, DPN-100 Interworking Guide*.)

Why use DPRS?

DPRS provides the following benefits in network routing:

- dynamic routing—DPRS makes traffic forwarding decisions at each hop in a route, based on current topology information.

- automatic rerouting—If failures or topology changes occur, packets are instantaneously rerouted.
- scaling ability—The DPRS hierarchical method of addressing allows you to build large networks by grouping Passport nodes into RID subnets.
- quality of service routing—Applications with different requirements can use the routing class of service that meets their needs in the most cost-effective way. For example, delay-sensitive applications use the minimum delay class of service.
- traffic balancing—DPRS provides both multipath and multilink traffic management mechanisms for increased application throughput.
- congestion management—DPRS reduces congestion through mechanisms such as different discard priorities, reliability classes of service, congestion indicators, and overflow routing.
- advanced services—Call routing services such as call redirection and hunt groups provide further methods of ensuring an application's availability.

  For information about call redirection, see 241-5701-410 *Passport 7400, 15000, 20000 Call Redirection Server Guide*. For information about hunt groups, see 241-5701-415 *Passport 7400, 15000, 20000 Hunt Group Server Guide*.

How does DPRS work?

DPRS is a connectionless routing system. In this type of system, the basic routing process involves

- establishing a connection between two access service ports (for example, frame relay ports)
- packetizing the data for the connection
- forwarding the packets across the network from the source node to the destination node

Each packet has a header that contains the destination node address and the identifier of the access service process. DPRS uses the addressing information in the packet header to index into the forwarding tables that it maintains at each node. These tables specify how DPRS can send the packet on towards the destination.

Figure 1, “DPRS connectionless routing,” (page 22) shows DPRS connectionless routing over frame-cell trunks. DPRS makes a routing decision at each hop on the path from the source node to the destination node. Under normal, stable network conditions, all packets follow the same route. Under congested situations or when topology changes occur, DPRS can instantaneously reroute the packets. As the figure shows, DPRS routes the first six packets through nodes E and F to their destination at node G. When the link between A and E becomes congested, DPRS sends the next four packets through nodes D and F.

For more information, see

- “Forwarding tables” (page 22)
- “Call establishment” (page 23)
- “Packet forwarding” (page 24)

Figure 1
DPRS connectionless routing

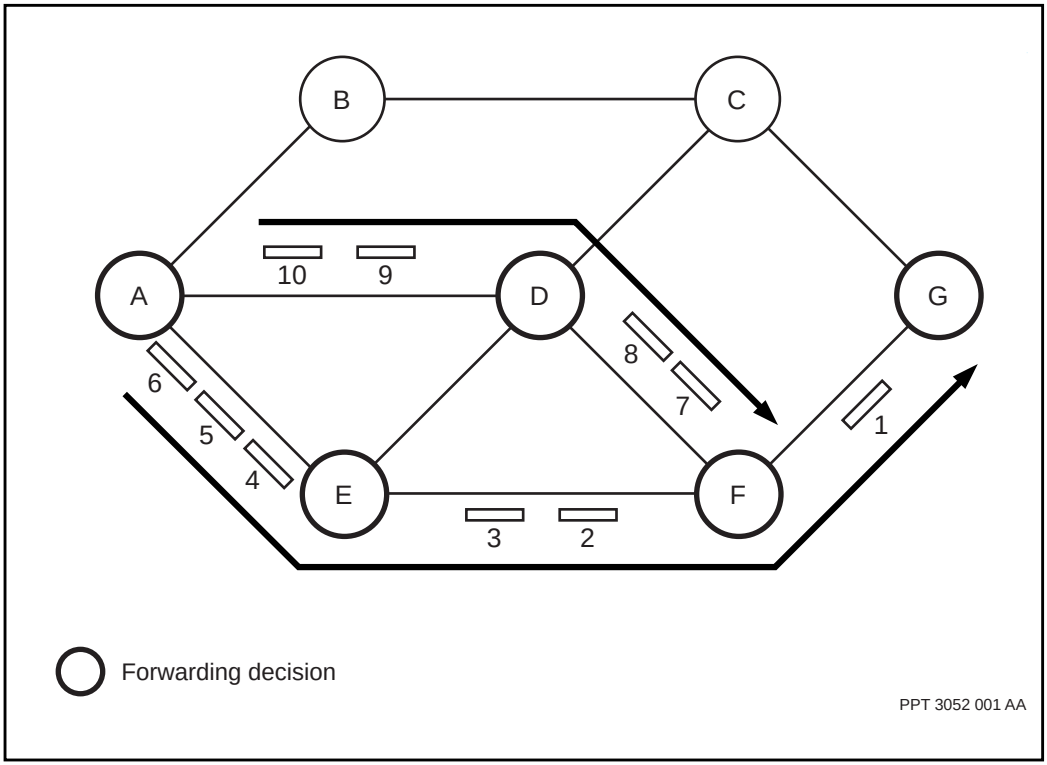

Forwarding tables

To be able to make the forwarding decisions at each hop, DPRS must have information about the current network topology and trunking characteristics. DPRS keeps this information on the control processors in tables known as routing tables. DPRS uses the information in these tables to create forwarding tables for each function processor.

DPRS creates the forwarding tables by obtaining

- topology information from the topology manager
- trunking information from the transport resource manager (TRM)
- DPRS-specific routing information from peer DPRS systems on other nodes in the network

DPRS uses this information to calculate the optimal route to each possible destination and stores the results in the forwarding tables. The forwarding tables map each possible destination address in the network to the appropriate next-hop outgoing link.

DPRS determines which links are preferred for delay-sensitive traffic, and which are best for applications that need greater throughput. For each destination address, the forwarding tables contain the next-hop link identifier for both the delay and throughput classes of service.

Call establishment

Before data packets can be routed, DPRS must set up a network connection between the two access service end points. This process is call establishment. In call establishment, DPRS works with the Passport call router (CR) to set up a virtual circuit (VC) for the call.

To begin call establishment, the originating access service sends a call request packet to the CR. The CR translates the first part of the data network address (DNA) provided by the service to determine the destination Passport node's address. This information allows the CR to deliver the call request packet to the CR at the correct destination node supporting that address. The destination CR translates the rest of the DNA to determine the destination access service. The CR sends the call request packet to that service. The service then accepts the call by returning a call accept packet to the originating node.

In Figure 2, "Establishing a call," (page 24), the call request packet is sent from the CR at node A to the CR at node B. The call accept packet returns to the originating node by the fastest possible route.

> ***Note:*** Networks that include DPN-100 nodes use a call server resource module (CSRM) for call establishment instead of a Passport call router. For more information on CSRMs, see 241-7401-110 *Passport 7400, DPN-100 Interworking Guide*.

Figure 2
Establishing a call

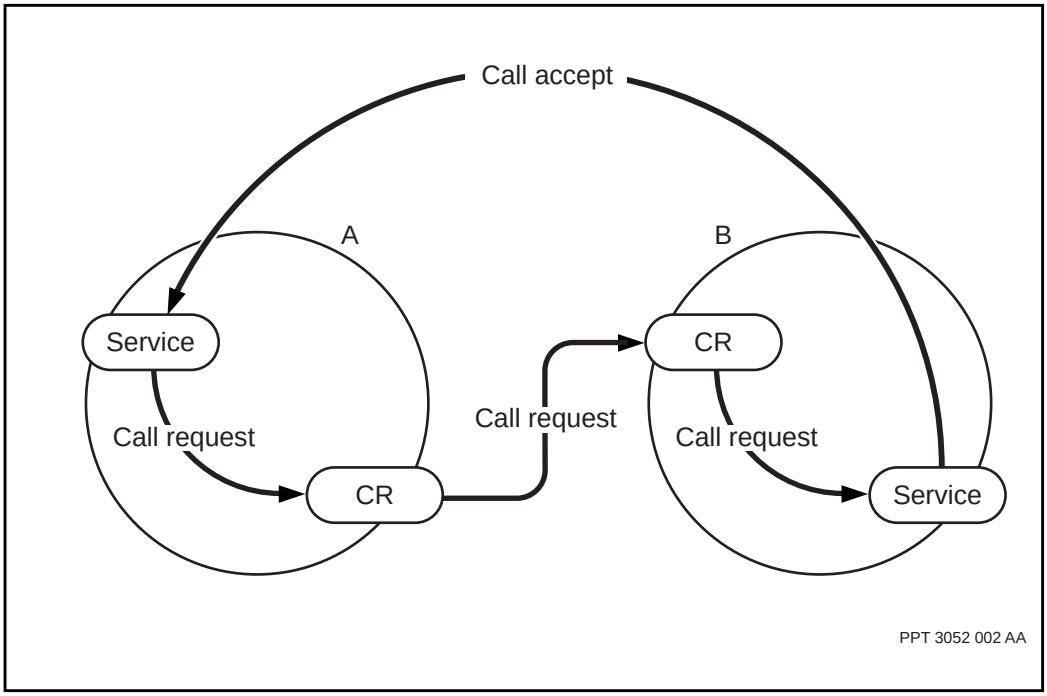

Packet forwarding

After the connection is established, DPRS can begin data transfer through the packet forwarding function. The packet forwarding function operates identically at each hop in the route.

When a packet arrives at a node, DPRS finds the destination address and the routing class of service (RCOS), either delay or throughput, in the packet header. DPRS uses the destination address as an index to the appropriate forwarding table (delay or throughput) to look up the next-hop link identifier in the packet's route.

Hierarchical addressing

Within the network, DPRS uses a form of hierarchical addressing known as routing identifier/module identifier (RID/MID). Figure 3, “DPRS addressing hierarchy,” (page 26) illustrates the three levels in the hierarchy:

- At the RID subnet level, Passport nodes are grouped and have a routing identifier (RID). Figure 3, “DPRS addressing hierarchy,” (page 26) shows three RID subnets: RIDs 10, 20 and 30. Passport networks support up to 126 RID subnets. In an interworked network, DPN-100 resource modules (RM) also have RIDs. In this case, the limit of 126 RIDs includes RMs.
- At the module level, each Passport node has a module identifier (MID). In Figure 3, “DPRS addressing hierarchy,” (page 26), the Passport nodes are identified with 4-digit MID values. Passport networks support up to 1909 MIDs in each RID subnet. Due to a hierarchical structure, individual MID values can be reused in different RID subnets. A Passport network, without Passport clusters, supports a maximum of 300 Passport nodes in a RID subnet, so only 300 of the 1909 possible MID values are used. Passport cluster nodes however, also have MIDs, and, in an interworked network, DPN-100 access modules (AMs) have MIDs. The total number of MIDs in a RID subnet then, including all Passport backbone nodes, Passport cluster nodes and AMs, cannot exceed 1909 MIDs.
- At the process level, each access service (for example, a frame relay PVC or SVC) has a module process identifier (MPID). Figure 3, “DPRS addressing hierarchy,” (page 26) shows MPID 66 assigned to a frame relay data link connection identifier (DLCI).

When DPRS routes packets between RID subnets, it examines only the RID field of the packet address. Traffic within a RID subnet routes according to the MID, and the MPID only routes packets within a module or node.

Figure 3
DPRS addressing hierarchy

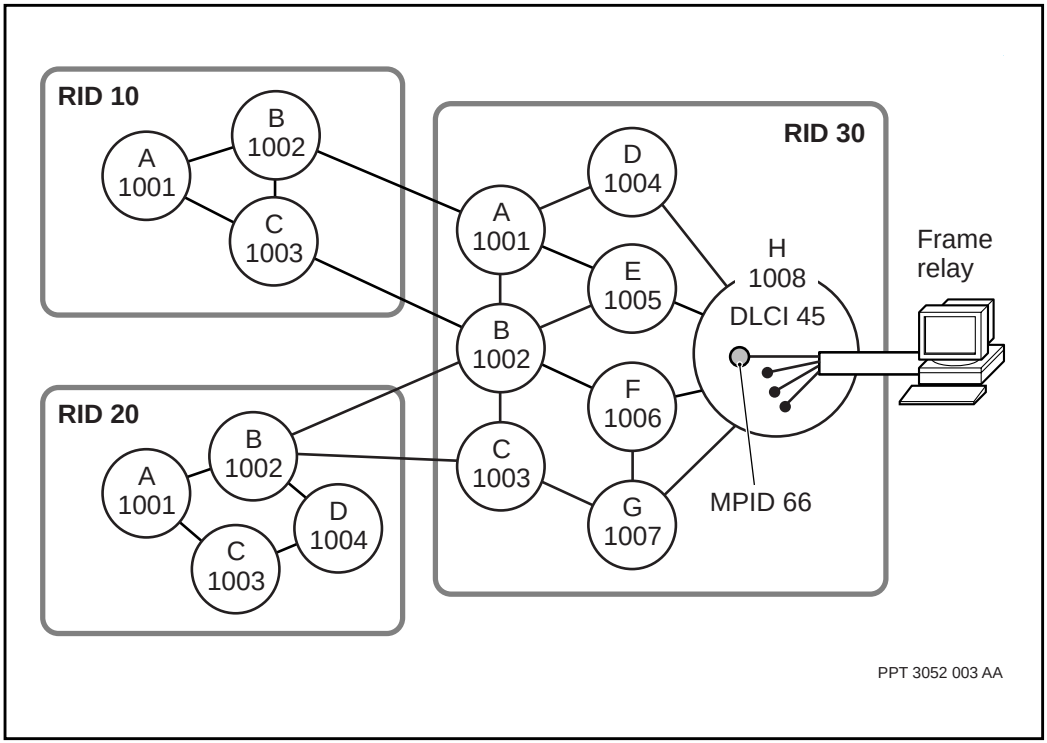

Chapter 2
DPRS routing

For a description of the Dynamic Packet Routing System (DPRS), see the following sections:

For basic information on network routing and types of routing systems, see 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*.

Routing process overview

Figure 4, "DPRS routing process," (page 29) shows an overview of the DPRS routing process. The first step in the process is to obtain the information for building the routing tables. The transport resource manager (TRM) provides link information to the topology manager. This information includes information about the links, such as availability, bandwidth, and round trip delay, as well as the identification of the remote nodes at the other ends of the links.

The topology manager exchanges topology information with other Passport nodes to create a RID subnet-wide view of the Passport topology. The topology manager uses this information to determine the best next-hop delay and throughput link groups to reach each node in the subnet. The topology manager passes a summary of this information to DPRS.

DPRS combines the topology information with DPRS information broadcast from other nodes to develop the forwarding tables. Then, DPRS creates the forwarding tables on each node as it activates, and updates them whenever there is a change in the network topology that affects the best paths chosen.

There are three sets of forwarding tables: RID, MID, and logical MID (LMID). DPRS creates a set of the three tables on each functional processor (FP). Figure 4, “DPRS routing process,” (page 29) provides a graphic representation of the tables.

Figure 4
DPRS routing process

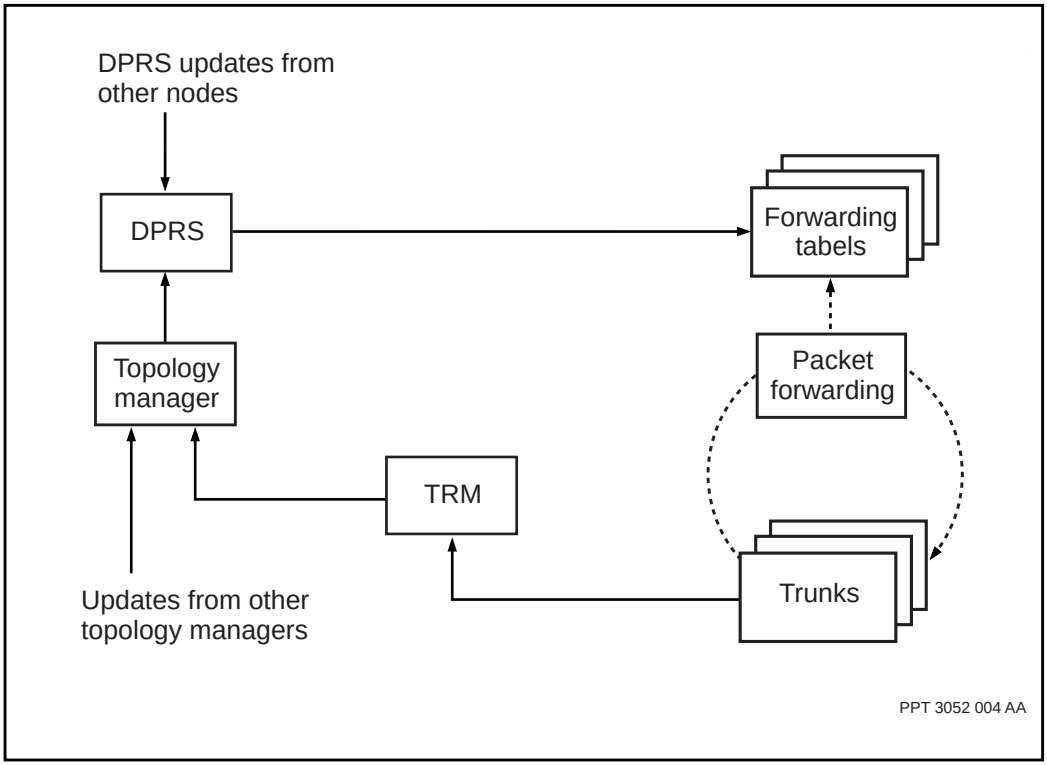

For each set of tables, there are two separate tables: one containing the best paths for the delay routing class of service (RCOS) and one containing the best throughput paths. These tables contain up to two next-hop link group identifiers for each destination RID, MID, or LMID address.

The RID tables contain the link group information for each possible destination RID in the Passport network. The MID tables contain the information for every destination MID in the RID subnet. The LMID tables contain an entry for each call router in the network.

After the forwarding tables are available, the DPRS forwarding system can use them to route the packets. At each hop in the route, the forwarding system examines the packet header to determine its destination. DPRS packet forwarding uses the packet header information and the forwarding tables to determine the next-hop link group for the packet.

During this process, packet forwarding first looks at the RID in the packet header. If it does not match the current RID subnet's RID, DPRS accesses the RID forwarding tables to determine the link group to forward the packet towards that RID.

When the RID matches that of the current subnet, packet forwarding checks the MID field in the packet header. If the MID does not match the current node's MID, DPRS uses the MID forwarding tables to determine the next-hop link group for that MID. When the MID matches that of the current node, DPRS uses the MPID field to determine which access service on the node is the destination for the packet.

To accommodate call establishment, the MID field in the packet header may contain a logical MID (LMID). LMIDs are used to identify the call routers used during call establishment. If the MID field in the packet header contains an LMID, packet forwarding checks the LMID forwarding table to determine how to forward the packet to the CR.

Figure 5
Forwarding tables

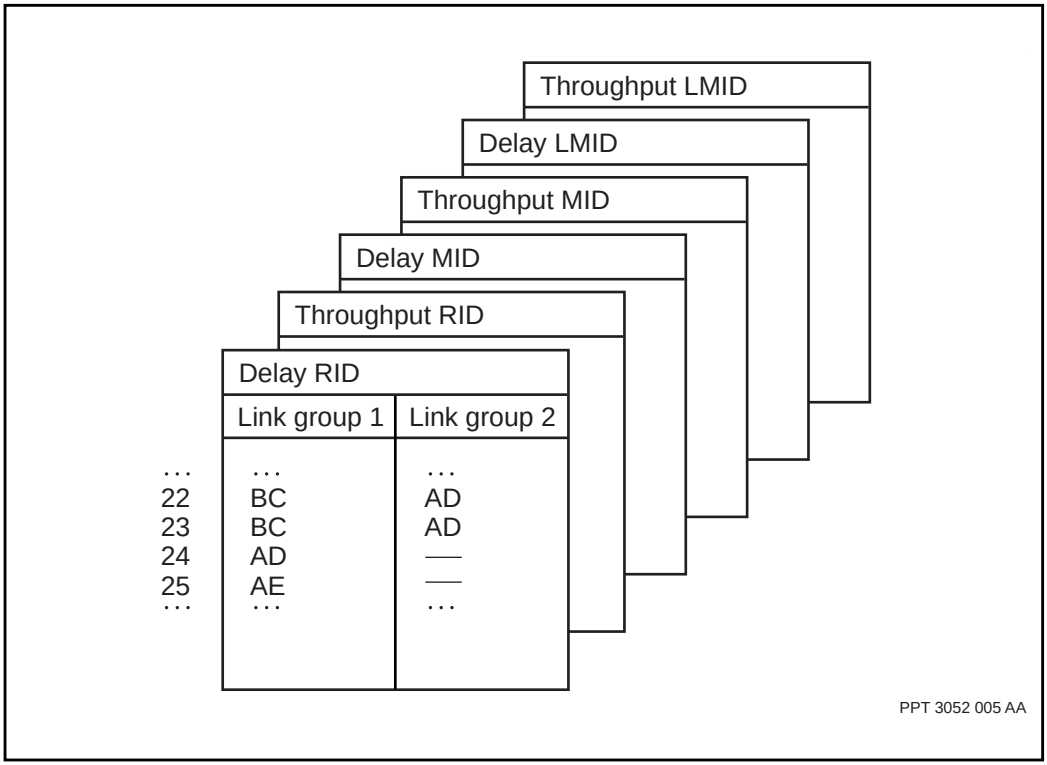

Spared DPRS on Passport 15000 and 20000

DPRS can be provisioned on a spare LP to be a warm standby feature for Passport 15000 or 20000. Hot and warm standby applications and features use sparing to reduce service down time and increase service reliability.

A warm standby application or feature can also operate together with a hot standby application or feature on the same FP without affecting the ability of the hot standby application or feature to provide hitless services during an equipment switchover.

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

Although DPRS is a warm standby feature, DPRS routing is interrupted during an equipment switchover. DPRS routing becomes available after:

- the switchover to the standby FP is complete
- the ATM networking and signaling capabilities are re-established

Creation of forwarding tables

The process of creating forwarding tables includes the following steps:

1 The topology manager determines up to two best paths to each node in the subnet. See “Topology metrics” (page 32).

2 DPRS combines the topology information with MID broadcast information to create the MID forwarding tables. See “MID metrics” (page 35).

3 DPRS combines the topology information with RID broadcasts to create the RID forwarding tables. See “RID metrics” (page 37).

4 DPRS combines the topology information with call services information to create the LMID forwarding tables. See “LMID metrics” (page 40).

5 When there are more than two best paths for a route, DPRS uses determinism rules to select between them. See “Routing determinism” (page 97).

Topology metrics

The topology manager uses the link information from the TRM and the topology information from other nodes in the subnet to compute optimal paths through the network using the shortest path first algorithm. This process determines the best path to each node in the subnet. The best path is the path with the minimum metric.

Delay and throughput metrics

Metrics are used to determine the best route through a Passport network. In effect, metrics are numbers used to represent the routing cost of a link group. They are calculated for every link group based on the characteristics of the links in the group. The lower the delay or the higher the throughput, the lower is the corresponding metric. The best route through the network is the one for which the sum of all the link group metrics between the source and destination is the minimum.

There are two types of metrics:

- The throughput metric is a value based on the reported speed of all the links in the link group. The reported throughput of the link is either the measured bandwidth of the link, the remaining bandwidth after the path-oriented routing system (PORS) reserved bandwidth is subtracted, or the override speed if it is provisioned. If both a PORS reserved bandwidth and an override speed are provisioned, the override value has precedence.
- The delay metric is a value based on the reported average delay of the links in the link group. The reported delay is the round trip delay for a 128-byte packet, which is measured at staging time. If a PORS reserved bandwidth or an override value for the round trip delay is provisioned on the link, either of these takes precedence over the reported delay. If both these attributes are configured, the override value has precedence.

The topology database

Because it has a view from each node in the RID subnet, the topology manager can compare the metric values on different links and build the best path to each node. The best path information includes the node identifier, two next-hop link group identifiers, and the metric for the path through each link group.

"Topology metrics" (page 35) shows a sample network of seven linked Passport nodes in a RID subnet. For simplicity, each link in the example is assigned a delay metric of 1000. The following table shows part of the topology information provided to DPRS on node G.

Table 1
Topology information

Node ID	Link group 1	Link group 2	Metric for link group 1	Metric for link group 2
C	EG	DG	2000	2000
E	EG	—	1000	—

According to this information, there are two possible equal-metric paths from node G to node C, one through links EG and CE, and one through links DG and CD. There is only one best path to node E, the path on link group EG at a metric of 1000. (The path on links DG and DE does not qualify as a best path, because the metric is 2000, twice as large as the EG link metric.)

The topology manager combines the metric information broadcast from each node, and at the end of the calculation process, has the two best paths for each destination node for each RCOS. The topology manager passes the next-hop link group of each route and the metric to DPRS. DPRS packet forwarding determines which of the two paths to use depending on the method of traffic balancing provisioned. See “Link group selection” (page 47).

By default, the topology manager selects two equal minimum-cost paths to each node. If there is no minimum-metric path equal to the first one selected, the topology manager selects only one path. However, if the variance option is turned on, the topology manager uses a slightly different algorithm, which allows it to select a second path that has a larger metric than the first path. For more information on variance, see “Variance” (page 61).

Figure 6
Topology metrics

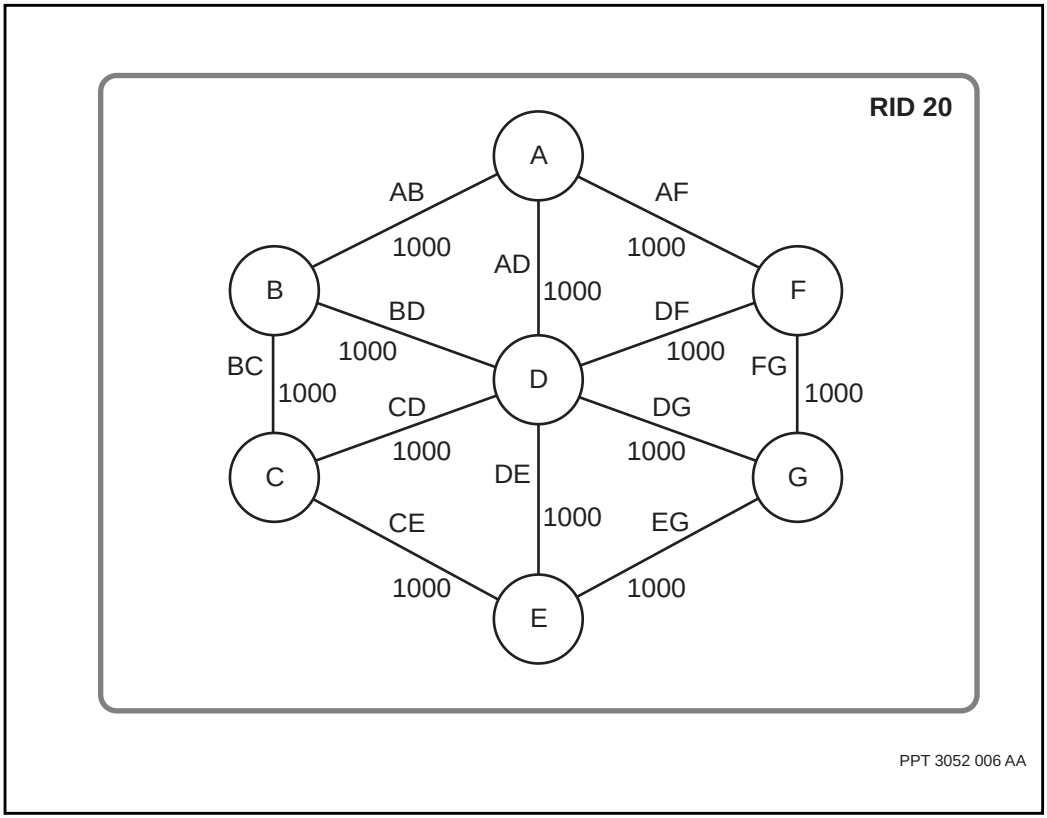

MID metrics

To create the MID forwarding tables, DPRS combines the topology database information with the MID value broadcast by each node. For example, Figure 6, “Topology metrics,” (page 35) shows a sample network with each link having a delay metric of 1000 and MID values for each node. The figure also shows part of the MID forwarding table DPRS creates at node G. The MID values provide an index to the next-hop link groups to those nodes.

Figure 7
MID metrics

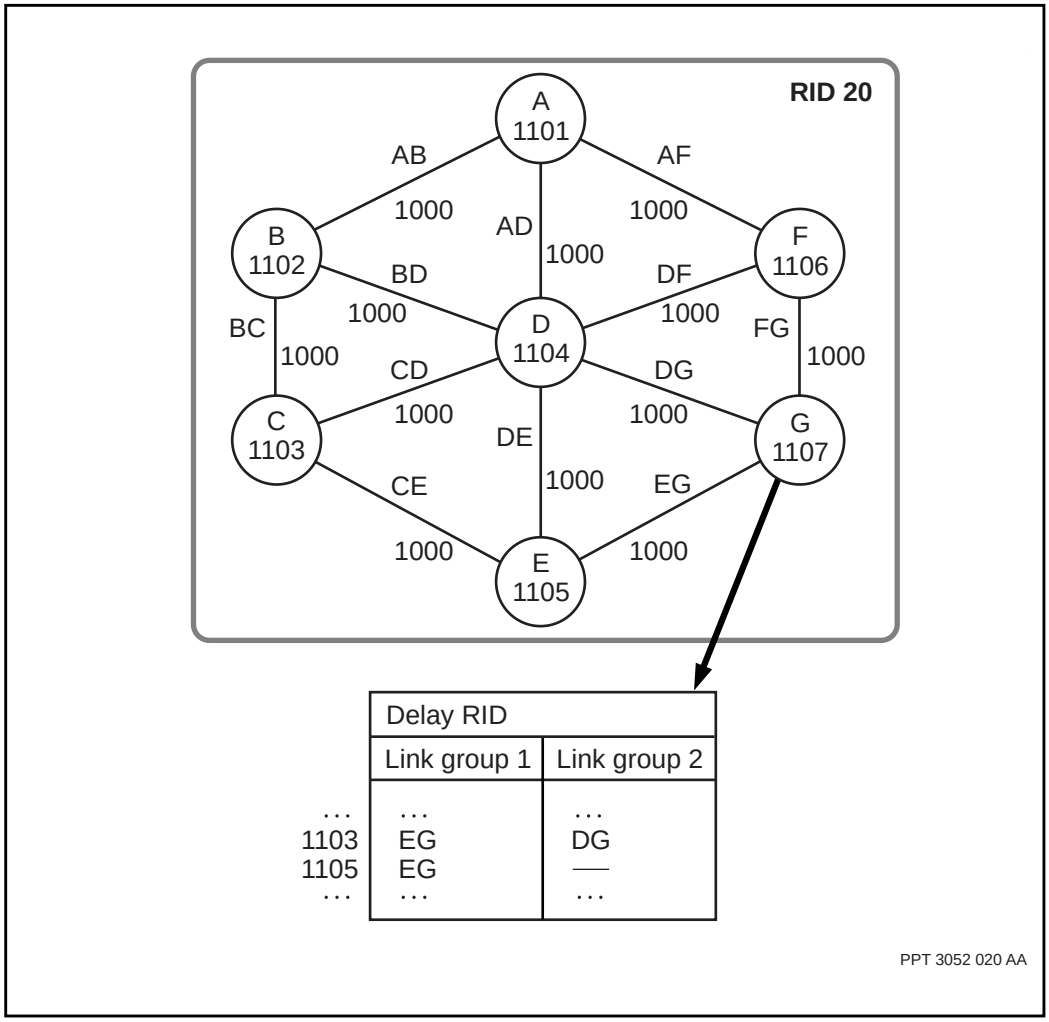

RID metrics

To create the RID forwarding tables, DPRS combines the topology database information with DPRS RID information from every connected node in the other RID subnets in the network. The RID routing information contains the delay and throughput metrics for every RID that can be reached through the neighbor subnet nodes.

Figure 8, “RID metrics,” (page 39) shows a sample network of three linked RID subnets. The following table shows part of the RID information shared by node C with other nodes in RID 20.

Table 2
RID information from node C

RID	Delay metric	Throughput metric
10	10	25
30	30	30
40	30	60

Node C can advertise the metrics to RID 10 and RID 30 because it has direct links to nodes in those subnets. Node C can also advertise the metrics to RID 40, which it receives from node H in RID 10. The metrics to reach RID 40 include the metric for the link to RID 10 (10 for delay and 25 for throughput), the metric to cross RID 10 (10 for delay and 10 for throughput), and the metric for the link to RID 40 (10 for delay and 25 for throughput).

The following table shows the RID information broadcast from node E to the other nodes in RID 20.

Table 3
RID information from node E

RID	Delay metric	Throughput metric
30	30	60

DPRS at node F receives these two broadcasts and combines this information with the topology information created for the subnet to form the RID forwarding tables. Figure 8, “RID metrics,” (page 39) shows part of the RID delay forwarding table at node F. This table shows that there is only one best path to RID 10, but two equal best paths to RID 30.

This example simplifies the metric calculations. In fact, the RID metrics (also known as DPN metrics) are calculated on a different scale from the MID metrics. Due to the difference in scale, DPRS converts the RID metrics to the topology manager’s metric scale before calculating the best paths to put in the forwarding tables. DPRS converts the throughput metric using the formula

$14240 \times \text{RidMetric}^2$

DPRS converts the delay metric using the formula

$\text{RidMetric} \times 6000$

By default, DPRS selects two equal minimum-metric paths. When the variance option is on, you can provision DPRS to select the best minimum metric path and the next-best path when there is only one best path. For more information about variance, see “Variance” (page 61).

Figure 8
RID metrics

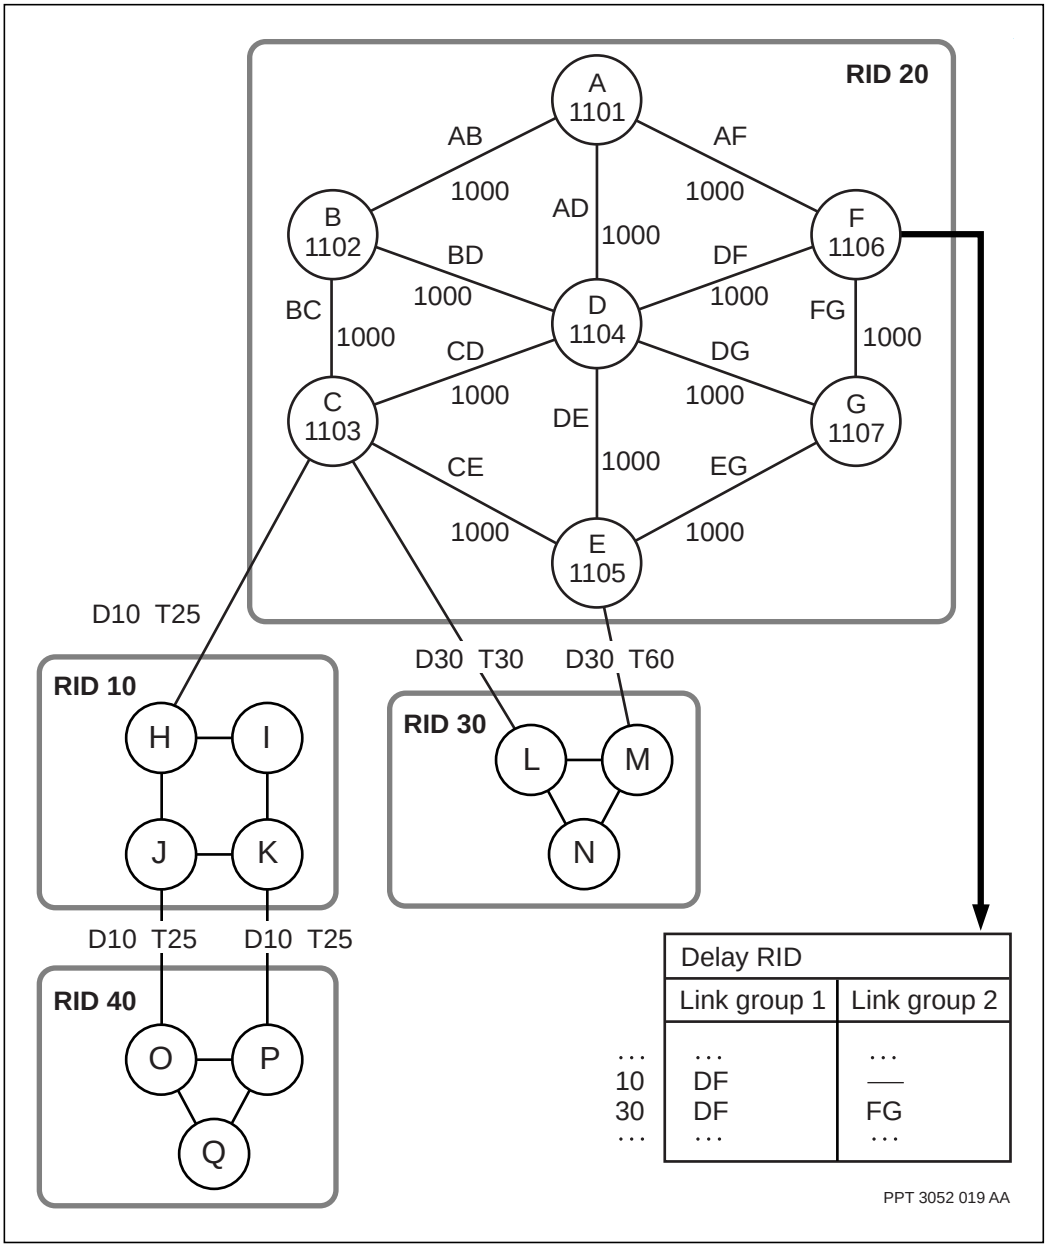

LMID metrics

Within the RID subnet, Passport nodes generate LMID broadcasts. These LMID broadcasts contain the source node's nodeId and the call services (such as call routing or call redirection) available on the node.

In the case of services like call redirection that are available on only a few Passport nodes in the subnet, DPRS compares the best paths in the topology database to the nodes supporting the LMID. This comparison determines the best next-hop link groups to the appropriate node (or nodes) to put in the LMID forwarding table. If there are two equal paths to a call service, DPRS puts both paths in the forwarding table and shares traffic between the two paths.

In the case of Passport call routing services, which are on every Passport node, DPRS puts the PID of the local call router directly into the LMID forwarding table.

Routing determinism

When the metric calculation process produces more than two equal minimum-metric paths for a route, DPRS uses routing determinism algorithms to determine which paths to use for the forwarding tables. For more information on routing determinism, see "Routing determinism" (page 97).

Call establishment

The process of establishing the connection differs slightly depending on whether the call destination is in the same RID subnet as the source, or in another RID subnet.

For more information, see the following sections:

- "Call establishment inside a RID subnet" (page 40)
- "Call establishment outside the RID subnet" (page 43)

Call establishment inside a RID subnet

Figure 9, "Call routing in the subnet," (page 42) shows the process of establishing a call in a single RID subnet.

Note: For simplicity, the DNIC values are not shown at the beginning of the sample DNAs.

The figure shows the three phases in the process:

1 The VC process of the access service creates the call request packet. The packet contains:

— the DNA (1021001), RID (51), MID (1001), and MPID (x) of the source access service

— the DNA (1030101) of the destination access service

— a destination RID of 0 and the LMID (8B40) of the call router supported by the source node

The VC process sends the call request packet to DPRS packet forwarding, which recognizes the RID of 0 as indicating the current subnet. Packet forwarding determines that the LMID is a call router address. It looks up the LMID in the LMID forwarding table, and sends the call request packet to the node's call router. For more information on how packet forwarding works, see "Packet forwarding" (page 45).

2 The CR translates the DNA in the call request packet to the destination MID using the prefix-DNA mapping tables. These tables contain the provisioned values of prefix-DNAs and their corresponding MIDs or RIDs. The CR finds that the prefix-DNA 10301 is the best (longest) match for the destination DNA. This DNA is supported by a MID in the local subnet.

The CR puts the destination MID (1004) and RID (0) values in the call request packet with the logical MPID (LMPID). The LMPID (1) is the address of the Passport final destination call router (FDCR) that handles DNAs supported on the destination node. The CR then invokes DPRS packet forwarding again to send the packet to the destination node.

3 At the destination node, the FDCR translates the rest of the DNA to determine the MPID (y) of the destination access service.The FDCR sends the call request to the destination access service, which creates the destination VC. The VC process exchanges the source and destination fields in the call request packet to create the call accept packet. Packet forwarding returns the call accept packet to the source access service.

Figure 9
Call routing in the subnet

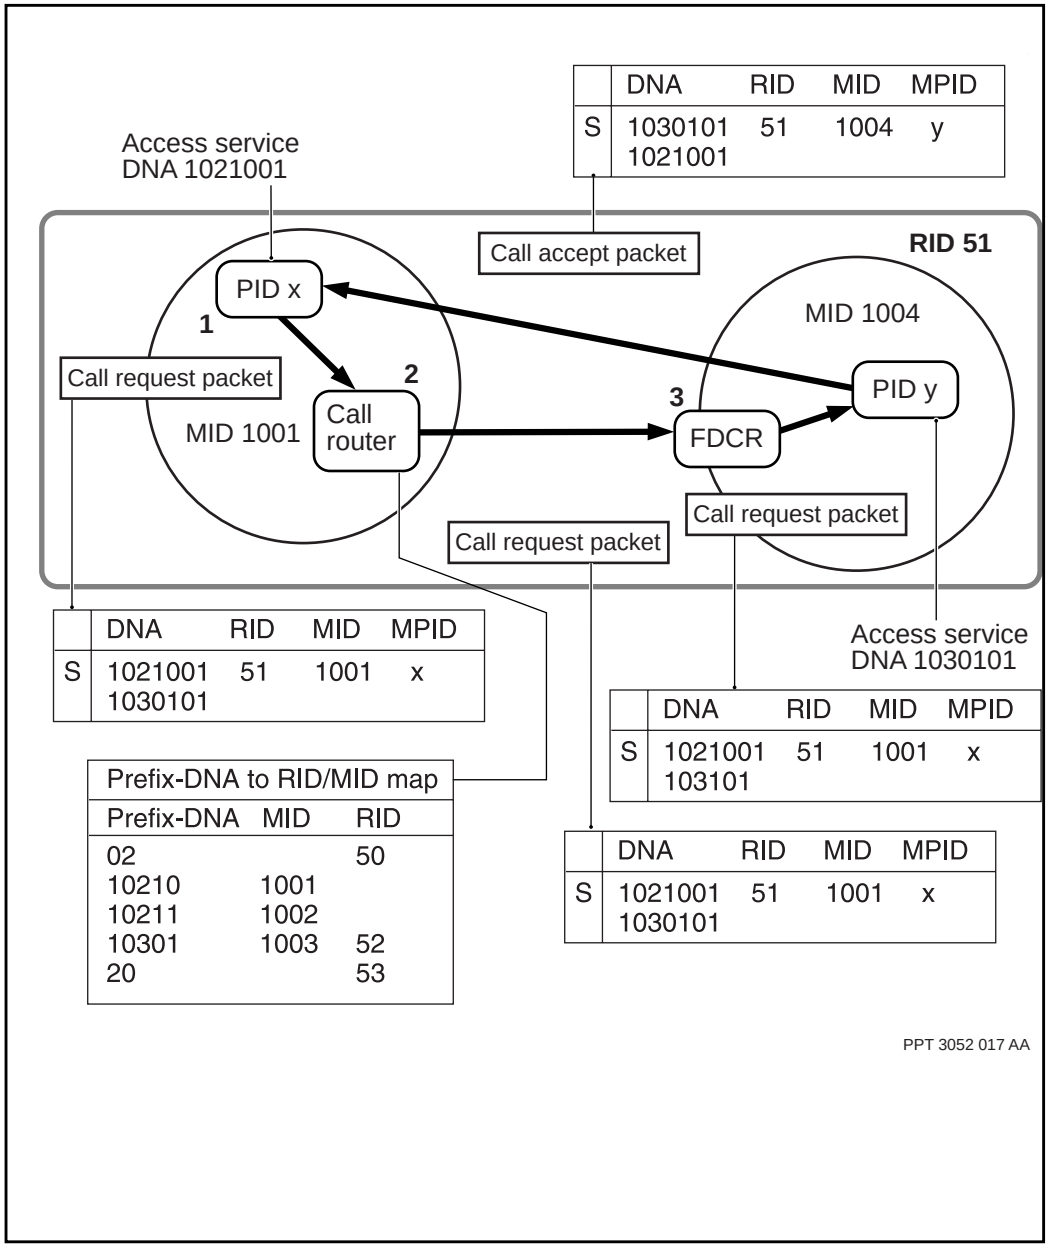

Call establishment outside the RID subnet

Figure 10, “Call routing outside the subnet,” (page 44) shows the process of establishing a call between multiple RID subnets. The figure shows the four phases in the process:

1 The VC process of the access service creates the call request packet. This phase of the process operates identically to the single-RID subnet process. The call request packet contains the DNA (2062201) of the access service, RID 0, and the LMID (8B40) of the local CR.

2 In this case, the prefix-DNA (20) maps to RID 52, so the call request packet is updated with the RID value and sent through DPRS packet forwarding to the CR (at LMID 8080) in the closest Passport node in the destination RID subnet.

3 The translation phase of the process is identical to the translation of the prefix-DNA in the case of the single RID subnet. The second CR maps the prefix-DNA (20622) to MID 2002, and the packet is sent to the FDCR at LMPID 1.

4 Determining the access service for multiple RID subnets is identical to determining the access service for a single RID subnet.

Figure 10
Call routing outside the subnet

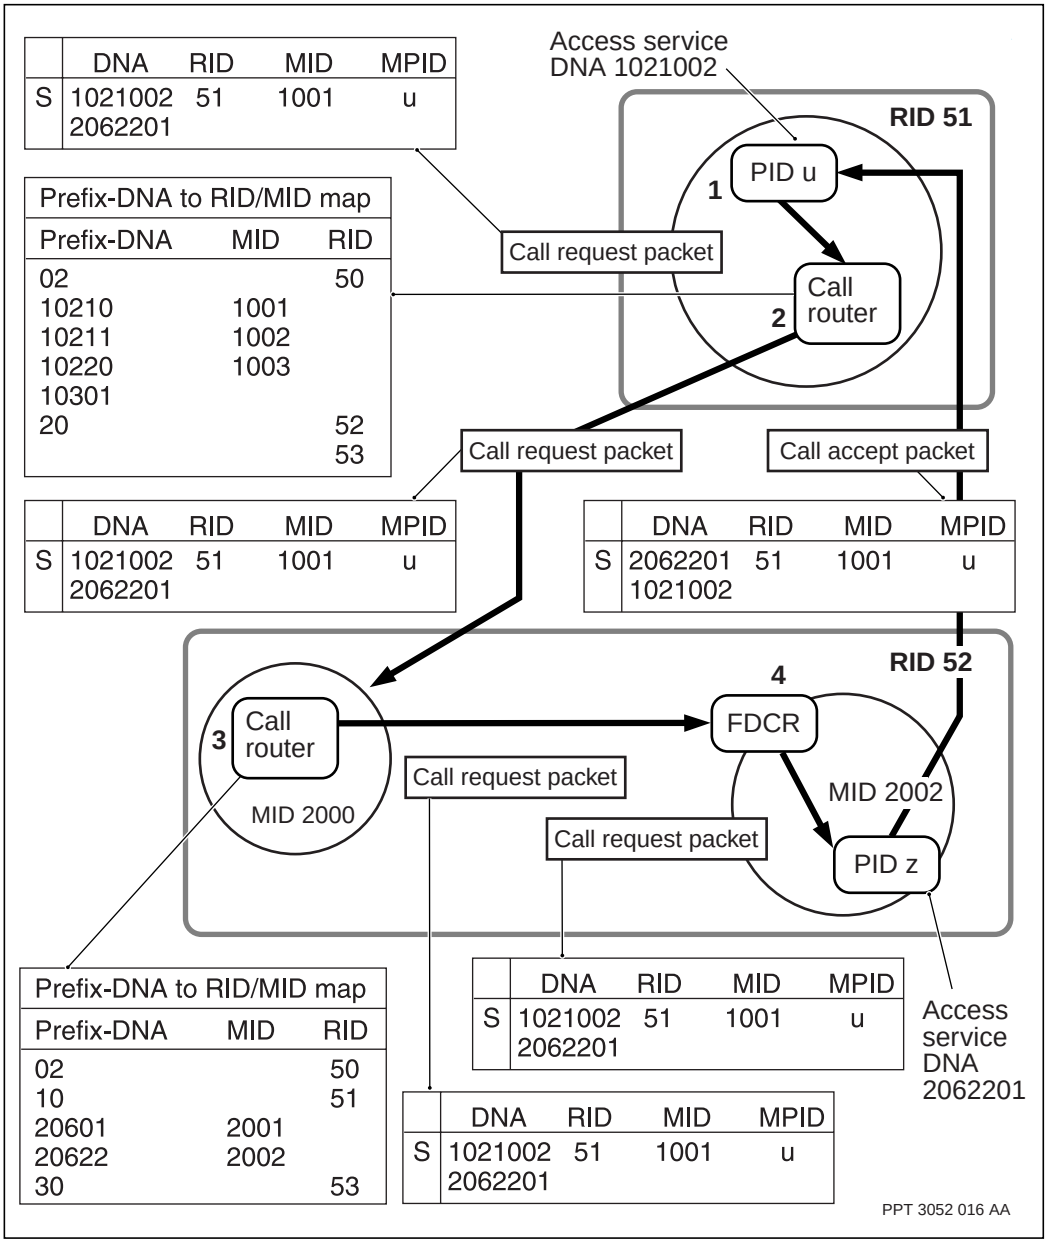

Packet forwarding

DPRS packet forwarding makes a decision on each node about where to send a packet to move it closer to its destination. DPRS packet forwarding forwards both data packets and call establishment packets. For more information on call establishment, see “Call establishment” (page 40).

Figure 11, “DPRS packet forwarding,” (page 46) shows the DPRS process of forwarding packets. DPRS checks the information in the packet header to determine which forwarding table to use. The forwarding table provides the link group identifier for indexing the link group tables. The link group tables list the links in the link group, and provide the following information for each link:

- congestion level
- destination FP
- destination queue identifier

DPRS uses the information provided by these attributes to select a link and send the packet to the queue for that link.

For more information on packet forwarding, see the following sections:

- “Packet header” (page 47)
- “Link group selection” (page 47)
- “Multimedia traffic forwarding” (page 48)

Figure 11
DPRS packet forwarding

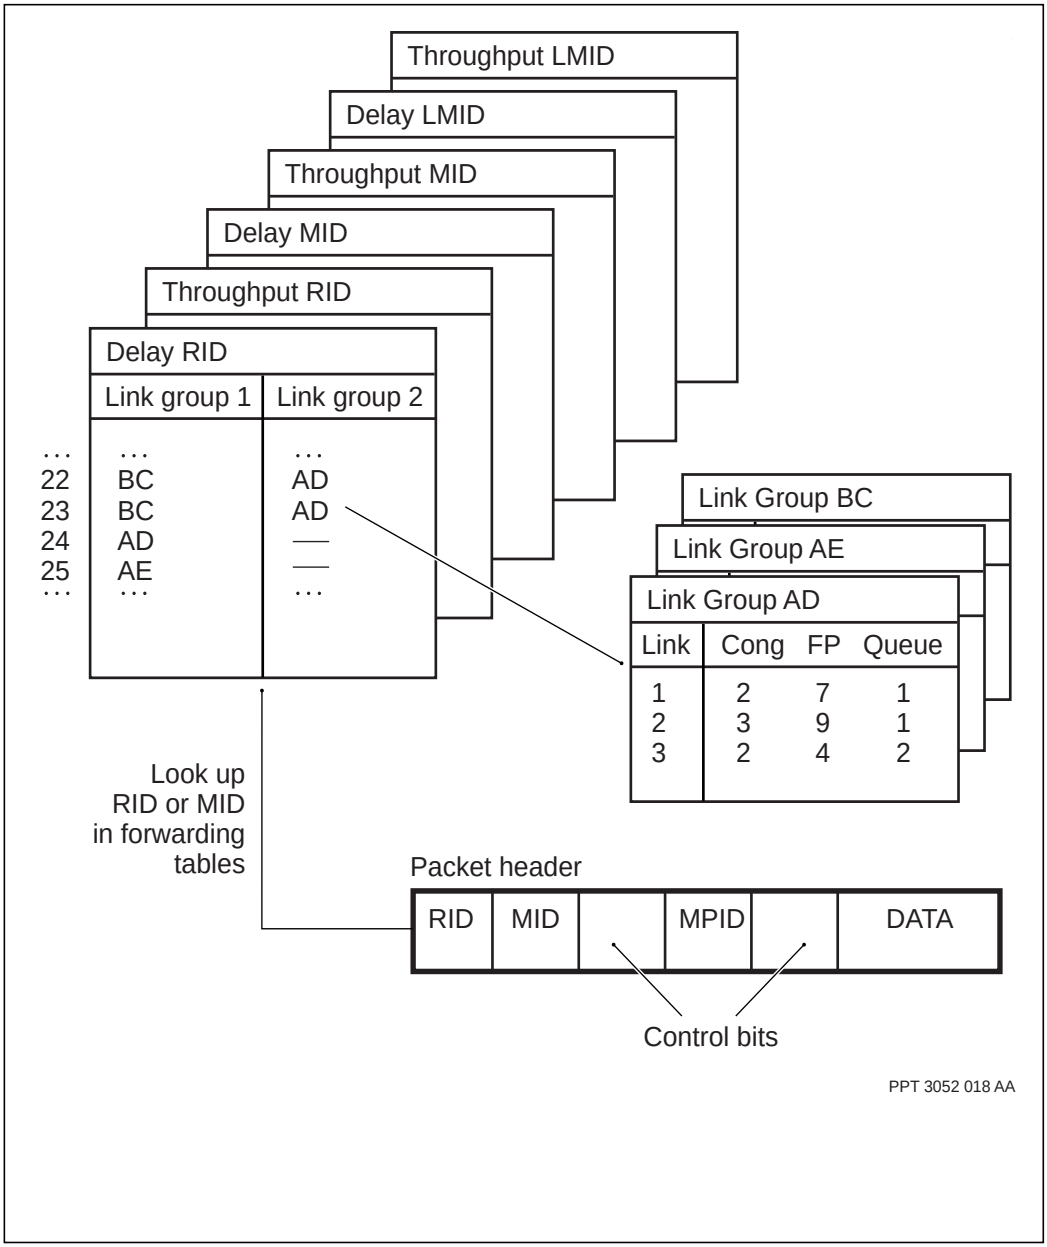

Packet header

DPRS checks the packet header to obtain the following information:

- the RID and MID addresses, which are used as an index into the appropriate forwarding table
- the RCOS indicator, which indicates the type of forwarding table to use, either delay or throughput

For comprehensive information on the packet header, see 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*.

Link group selection

DPRS selects the next-hop link group from the forwarding table using the load spreading traffic balancing algorithm. This path selection method is designed to randomize the choice of paths from the forwarding table, while ensuring that all packets belonging to the same traffic flow (or VC) are sent on the same path, and that the total number of traffic flows are statistically divided over the two paths.

Under congestion conditions, the loadspread algorithm allows high reliability traffic to overflow to an uncongested link in the alternative link group.

For more information on traffic balancing methods, see “Congestion under load spreading” (page 70) and “Congestion under load sharing” (page 71).

Link selection

After DPRS chooses the next-hop link group, it picks a link within the link group based on the traffic balancing algorithm in effect. The two provisionable algorithms are loadspread and loadshare (loadspread is the default).

The loadspread algorithm randomizes link choices, while ensuring that all packets belonging to the same VC are sent on the same link, as long as the link is not congested. In this way, the traffic flow remains ordered and experiences minimal delay variation. When the link becomes congested, the algorithm allows traffic to overflow onto other uncongested links in the group to avoid discarding packets.

The loadshare algorithm equalizes the use of links within a group. After a link group is chosen, the algorithm examines each link in the group, and sends the packet on the link with the lowest offered load value. The offered load value indicates how much traffic has been sent to the link, and is weighted by the inverse of the link speed.

For more information on traffic balancing methods, see “Congestion under load spreading” (page 70) and “Congestion under load sharing” (page 71).

Multimedia traffic forwarding

Multimedia traffic requires a different algorithm. This algorithm allows voice to be transported over frame relay with excellent transfer delay and delay variation characteristics. The algorithm only uses delay link groups, and does not allow excess traffic to overflow onto other links in the group under congestion conditions. The multimedia forwarding algorithm uses the same link group and link selection methods as load spreading (even if the forwarding algorithm is set to loadshare).

Routing in a network that contains Passport clusters

This section contains information on the following situations:

- “Routing from a Passport cluster to the backbone” (page 48)
- “Routing from the backbone to a Passport cluster” (page 50)
- “Routing from Passport cluster to Passport cluster” (page 51)

Routing from a Passport cluster to the backbone

RIDs, MIDs, and LMIDs of all nodes in the backbone and of other Passport clusters are not known to a Passport cluster. To access these nodes, a default routing protocol is implemented in a Passport cluster. It is in the form of forwarding up, that is, all packets destined to any unknown address outside of the cluster are forwarded to the backbone.

All information necessary to reach the backbone border nodes is advertised inside a cluster. The routing from one MID in a cluster to another MID in the backbone is called two-hierarchy level routing. Routing first occurs along the default path from the cluster source MID to the nearest backbone border node. Routing then occurs along the best path from the backbone border node to the destination MID. In the figure “Default routing from a Passport cluster to a

backbone" (page 49), path A-C-D is the default path from the source to the backbone. At node C, since the backbone portion of the path is unknown, the best link C-D is chosen. Once traffic is forwarded to the backbone border node, the optimal path is used from that point to the destination node.

Figure 12
Default routing from a Passport cluster to a backbone

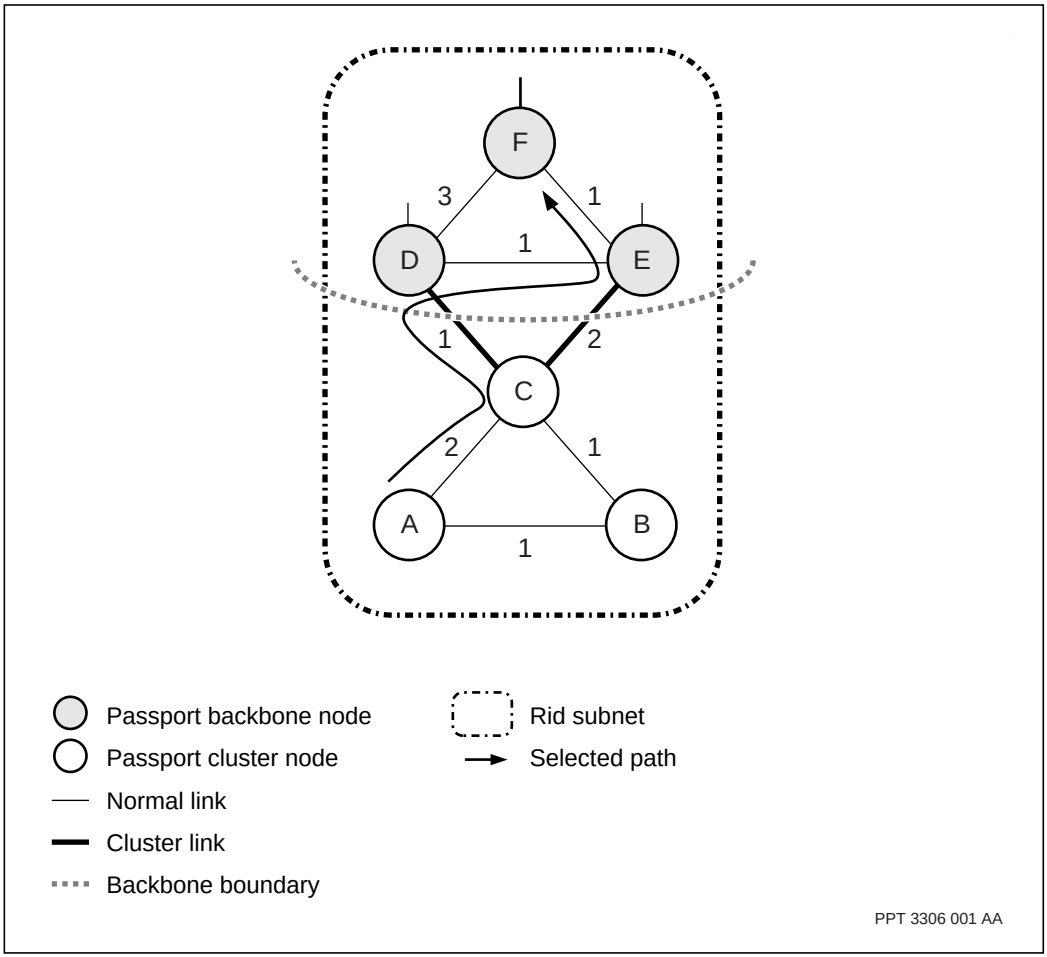

Routing from the backbone to a Passport cluster

From a backbone node, routing to a Passport cluster in the same RID subnet is the same as routing with other backbone nodes in the subnet. Reachable MIDs in clusters are learned from MID broadcasts propagated from the clusters and advertised to all backbone nodes in the subnet. With the reachability information, paths to reach each cluster MID are calculated independently by all subnet backbone nodes for each routing class of service (RCOS). Traffic destined to a MID inside a cluster can be optimally routed to the cluster. Traffic destined to a cluster in another RID subnet is sent through the existing RID routing protocol.

Note: Even though Passport cluster nodes are visible to backbone nodes, Passport cluster nodes are never selected as tandem nodes on the path from one backbone node to reach another backbone node. This is illustrated in Figure 13, “Optimal path to reach a backbone node,” (page 51). Since Passport cluster node E is isolated from the backbone’s topology and E does not include backbone nodes in its reachable node summaries, the path computation performed on backbone node A would not pick up the path that tandems through E as the best path to reach D.

Figure 13
Optimal path to reach a backbone node

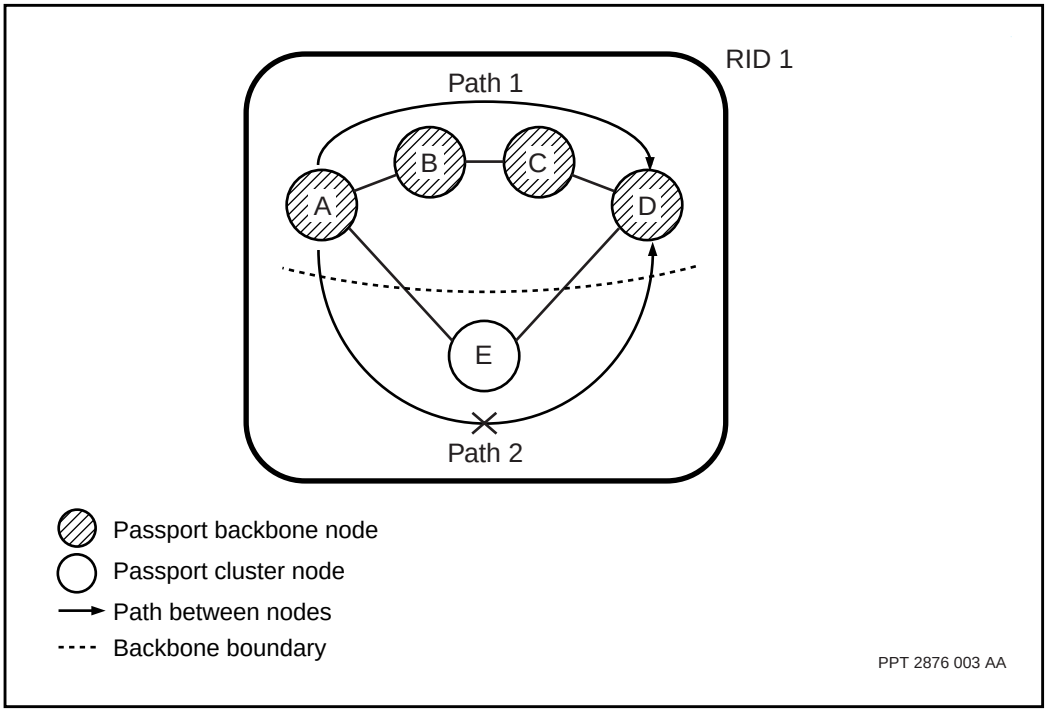

Routing from Passport cluster to Passport cluster

Routing between clusters involves two paths: from the cluster to the nearest backbone node via a default route, and from the backbone node to the destination cluster via the RID or MID routing protocol.

For adjacent clusters as shown in Figure 14, “Routing from one Passport cluster to another,” (page 52), two default routes are selected at node A, because the cluster link groups have equal metrics to the backbone. In this case, traffic is spread between two paths as A-C-D-E and A-D-E, although the total metric of path A-C-D-E is worse than that of A-D-E.

Figure 14
Routing from one Passport cluster to another

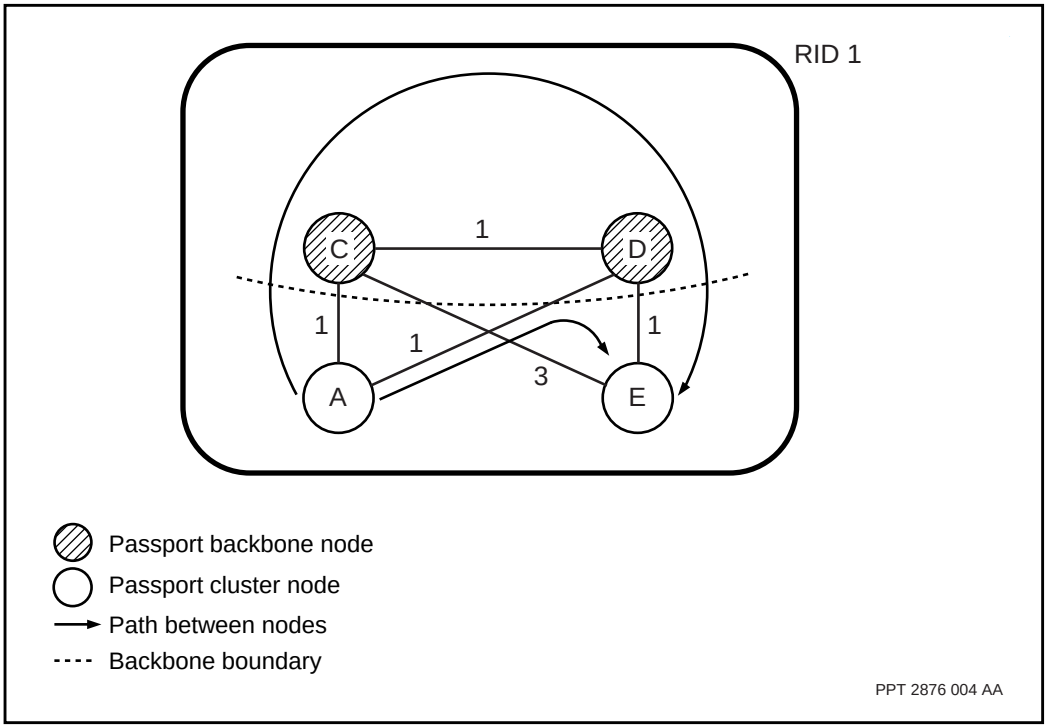

Chapter 3
Traffic management

For a description of the Dynamic Packet Routing System (DPRS) traffic management mechanisms, see the following sections:

- “Bandwidth management” (page 53)
- “RCOS routing” (page 59)
- “Variance” (page 61)
- “Congestion management” (page 70)
- “Tandem suppression” (page 71)
- “Delay and throughput cutoff” (page 74)

For further information on Passport traffic management mechanisms, see 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*.

Bandwidth management

DPRS uses the following bandwidth management mechanisms to ensure the efficient use of resources during packet routing:

- “Available bandwidth calculation” (page 54)
- “Preferred link selection” (page 55)
- “Packet forwarding algorithms” (page 55)
- “Dynamic trunk speed change” (page 58)

Available bandwidth calculation

In the process of creating the forwarding tables, the topology manager and DPRS must calculate the metrics of all the possible routes to each destination in the network. This metric calculation is based on the trunk information provided by the transport resource manager (TRM).

As well as obtaining link state information, TRM calculates the available bandwidth on each link. To do this, TRM must determine first whether a link is configured with reserved path-oriented routing system (PORS) bandwidth or a link override. An override is a provisioned trunk bandwidth or delay value that takes precedence over the measured bandwidth or delay.

If the trunk is provisioned with an override, TRM uses that value as the prevailing link bandwidth. (For more information about setting trunk overrides, see 241-5701-420 *Passport 7400, 15000, 20000 Trunking Guide*.)

If there is no override and there is reserved PORS bandwidth provisioned, the TRM applies the reserved PORS bandwidth to the link delay and throughput metrics as follows:

- For throughput metrics, TRM subtracts the provisioned percentage of bandwidth reserved for PORS from the overall trunk bandwidth. TRM then reports the remaining bandwidth to the topology manager. For example, if a 2 Mbyte link is 60% reserved for PORS, TRM calculates the DPRS bandwidth as 0.8 Mbyte.
- For delay metrics, TRM places a bias on a link with bandwidth reserved for PORS, to give a larger reported delay for connectionless routing. TRM does this by breaking the measured delay into two components:
  - — the emission delay, which is a function of the throughput of the link
  - — the propagation delay, which is the remaining delay after the emission delay is subtracted from the measured delay

To compute the delay for metric calculations, TRM divides the emission delay by the connectionless percentage of bandwidth. (The connectionless percentage of bandwidth is the bandwidth remaining after the PORS percentage is subtracted.) TRM then adds the propagation delay to the result, and this value determines the delay metric for the link.

Preferred link selection

After TRM calculates the available bandwidth of the links, it assembles the link information for all direct links between two Passport nodes to form a link group. In this process, TRM determines the preferred links for both throughput and delay routing for DPRS packet forwarding. Preferred links are those links in a link group which normally carry DPRS traffic for that RCOS.

When the default packet forwarding method, load spreading, is in effect, TRM selects the link with the largest bandwidth (in bits/second) in the link group as a throughput-preferred link. TRM also considers any link with at least half of the bandwidth of the highest throughput link as a throughput-preferred link.

If load sharing is in effect, TRM partitions the links in a link group into two sets: those with a delay greater than 150 milliseconds, and those with a delay less than or equal to 150 milliseconds. TRM chooses the set with the greatest aggregate bandwidth for throughput, and all the links in the set are marked as throughput-preferred. If both sets have the same aggregate bandwidth, TRM chooses the set with the lower delay.

TRM selects the link with the lowest delay in the link group as a delay-preferred link. It also considers any link with 1.5 times the delay of the lowest delay link (or less) to be a delay-preferred link. Any link with a delay of less than 20 milliseconds is also a delay-preferred link.

If a link is provisioned to carry 100% PORS traffic, TRM does not mark it as a throughput-preferred or delay-preferred link. If all links in the link group are 100% PORS, then TRM marks all of them as preferred.

Packet forwarding algorithms

Loadspread (the default algorithm), loadshare, and loadspreadfast are the three alternative algorithms you can choose for link selection during packet forwarding.

The loadspread algorithm is a method of distributing the traffic across the chosen link groups, and the preferred links in those groups. The loadspread algorithm keeps all the traffic for a particular VC in a link group, and on one link within the group.

The figure “Load spreading” (page 57) shows the load spreading mechanism. In the diagram, all packets for VCx sequentially follow each other across the same link in the link group. The same applies to the packets for VCw, VCy, and VCz. Packet forwarding uses bits of the destination VC address in the packet header to select the link for the packet. Since all packets for a VC are destined for the same address, all traffic for one VC stays on the same path through the network, and on the same link in each link group traversed.

The loadshare algorithm is a method of sharing traffic across the preferred links in a link group to make the entire link group’s bandwidth available to each VC. As shown in the figure “Load sharing” (page 58), traffic from a single VC is distributed across all links in a group.

Load sharing of VC traffic is sensitive to the capacity of each link in the link group. Links in the group are offered traffic based on their available capacity. Higher-bandwidth links receive a greater proportion of the traffic load than lower-bandwidth links.

In addition to load sharing traffic across all the preferred links in the group, DPRS continues to load spread traffic across up to two next-hop link groups to the destination. This practice implies that traffic from a particular VC always remains within the same link group, minimizing the disordering of packets.

The loadspreadFast algorithm is another method of distributing the traffic across the chosen link groups and the preferred links in those groups. When the loadspreadFast algorithm is used, 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.

For more detail on engineering DPRS packet forwarding, see “Guidelines for packet forwarding” (page 81).

Figure 15
Load spreading

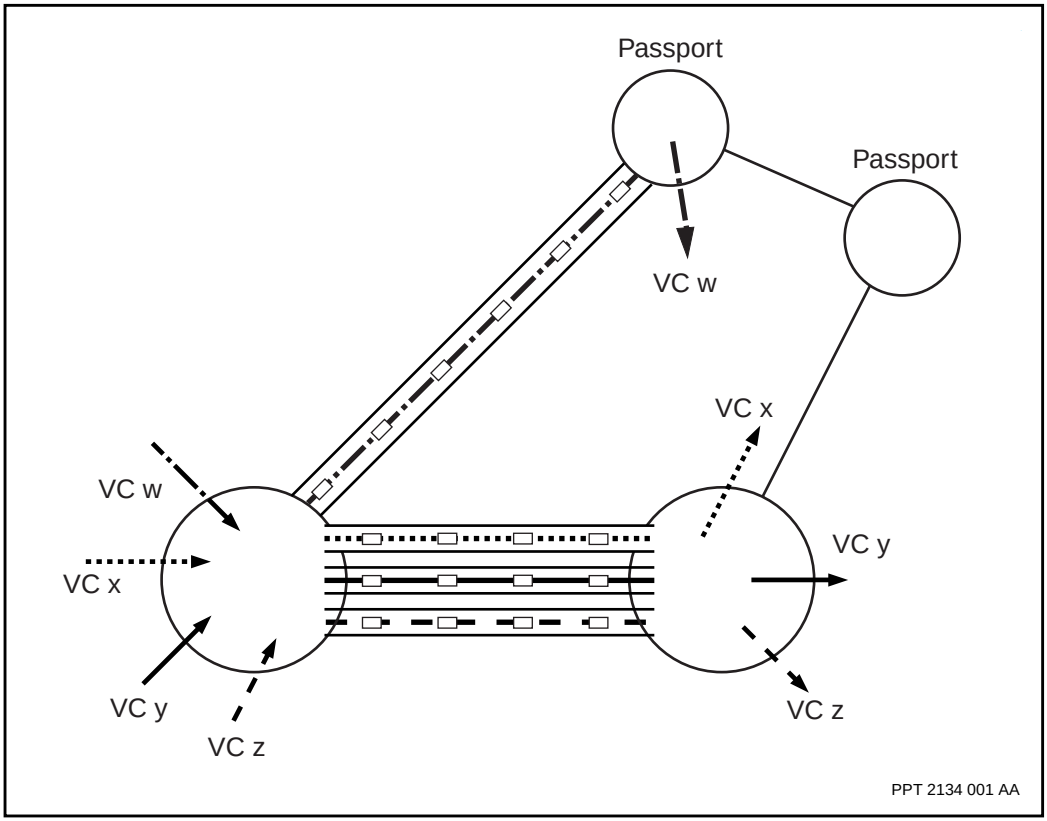

Figure 16
Load sharing

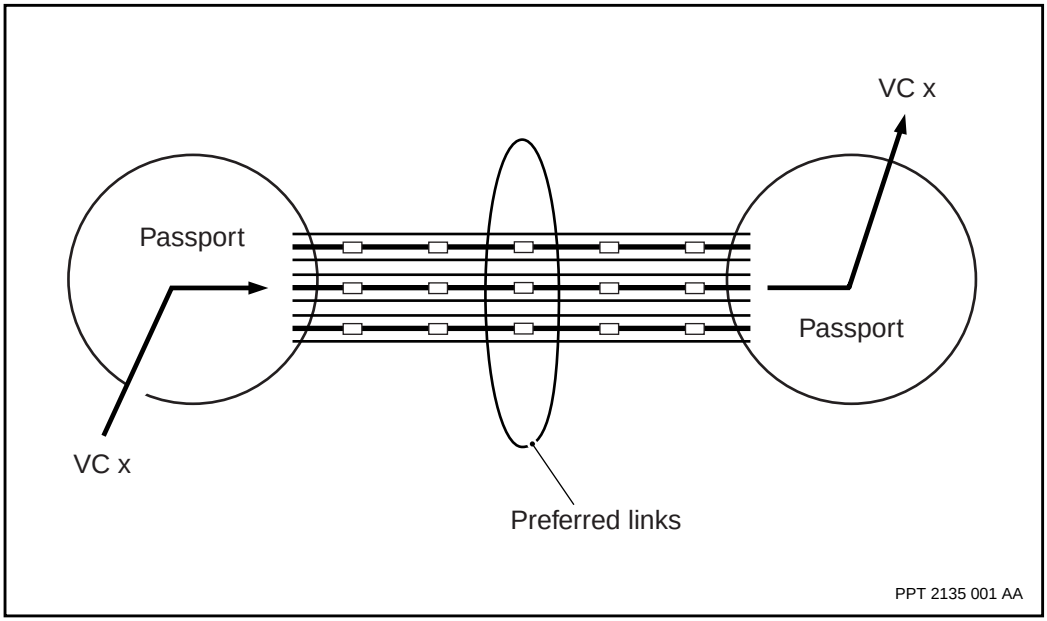

Dynamic trunk speed change

By default, any dynamic trunk speed change that occurs is not reported to DPRS. This situation means that the change in speed does not affect the calculation of metrics. However, you can enable and customize the speed change reporting mechanism so that speed changes can be included in the metric calculation process. This capability allows DPRS to adjust traffic flows in response to the changes. For more information on dynamic trunk speed change, see 241-5701-420 *Passport 7400, 15000, 20000 Trunking Guide*.

For more information on RCOS routing, see the following sections:

- “RCOS routing attributes” (page 59)
- “Multimedia RCOS routing” (page 60)

RCOS routing

DPRS supports routing class of service (RCOS) routing, which allows traffic to be routed according to the characteristics and needs of the individual services.

RCOS routing attributes

Routing class defines whether the packet requires high speed (minimum delay), high throughput (bandwidth), or high emission priority and minimum delay variance. Time-sensitive data is typically sent using the delay RCOS or the multimedia RCOS. Large quantities of data that are not so time-critical are commonly sent using the throughput RCOS. As frame relay multimedia data is sensitive to delay variance, it is sent on delay class paths with the highest emission priority service available on the path.

The DPRS system of RCOS classification includes the three attributes: delay, throughput, and multimedia. These attributes are indicated through a combination of the minimum delay (MD) bit in the DPRS packet header and the emission priority (EP) bit in the Passport common header. The table “RCOS attributes” (page 60) explains the attributes and shows the combinations of packet header bits that result in the RCOS values.

DPRS supports a type of routing known as default RCOS. This term means that if a delay path is not available for delay or multimedia traffic, that traffic is routed on a throughput path.

Table 4
RCOS attributes

RCOS attribute	Packet header bit	Characteristics
Throughput	MD=0 (throughput) EP=0	This attribute determines whether a packet takes the high throughput (greater bandwidth) route. Large quantities of delay-insensitive data are transmitted using the throughput RCOS. Throughput RCOS traffic is sent over the normal emission trunk queue.
Delay	MD=1 (delay) EP=0	This attribute determines whether a packet takes the least delay (high speed) route. Delay-sensitive data is transmitted using the delay RCOS. Delay RCOS traffic is sent over the high emission trunk queue.
Multimedia (For the Passport 7400 series switch only	MD=1 EP=1 (high)	This attribute indicates whether the packet uses the multimedia RCOS. The multimedia RCOS ensures that traffic (such as frame relay multimedia data) is guaranteed a high class of service as it is transported throughout the network. Multimedia RCOS is routed on the interrupting queue.

Multimedia RCOS routing

The multimedia RCOS ensures that delay-sensitive multimedia traffic is guaranteed a high quality of service in transport throughout the network. DPRS packet forwarding sends multimedia RCOS traffic through delay-preferred links. In hybrid networks, the multimedia class of service is maintained on interrupting trunks throughout its route.

If a packet contains multimedia data for a frame relay service, the EP is set in the Passport common header. When this bit is set, DPRS packet forwarding handles the packet as follows:

- At the DPRS ingress point, the traffic enters the network through a Passport access switch, and the multimedia traffic routing depends on the Passport trunk type:

  - — When the multimedia traffic frames are routed to a frame-cell interrupting trunk, the packets are sent to the interrupting link transmit queue.
  - — When the frames are routed to a frame-cell HDLC trunk, the packets are sent to the high priority link transmit queue.
  - — When the multimedia traffic is routed to a Passport trunk over ATM, the packets are sent to the ATM link transmission queue. There, the traffic is prioritized according to the service category of the VCC.
- At the DPRS tandem, or the egress point where the traffic leaves the network, DPRS receives the multimedia traffic from the high-priority link receive queue (for frame-cell interrupting trunks), the normal-priority link receive queue (for frame-cell HDLC trunks), or the ATM link receive queue. DPRS then sends the packet to the appropriate link transmit queue or the destination function processor (FP).

Variance

Variance 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.

For more information on variance, see the following sections:

- “Multipaths” (page 62)
- “Variance safety check” (page 62)
- “Traffic distribution” (page 64)
- “Network routing with variance” (page 64)
- “Variance with the delay RCOS” (page 65)
- “Variance with reserved PORS bandwidth” (page 66)
- “Multipaths with more hops” (page 68)
- “Variance values” (page 68)

Multipaths

In the default behavior of the Passport routing system, every path for routing 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.

Variance 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.

You can provision variance values for both the delay and throughput RCOS. The variance value for an RCOS is the maximum allowable difference in metric between the best path and a second path. The value is 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:

- has a metric within the provisioned range of the best path metric
- is loop-free (does 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.

Variance safety check

A safety check is applied to any candidate unequal path within the provisioned variance value to ensure that the path does not result in a loop. Non-minimum alternative paths can be guaranteed to be loop-free as long as the next hop in the path brings the traffic closer to its destination. The equation used to guarantee that the path is loop-free is as follows:

(metric of non-minimum path - metric of the first hop link group in that path) < metric of the best path

For example, as the figure “Multipath safety check” (page 63) shows, path A-B-D is an unsafe path, since traffic sent from node A to node B destined for node D is routed back to node A, causing a routing loop. (Node B perceives its best path to node D to be through node A.) The metrics from the figure “Multipath safety check” (page 63) in the safety-check equation are as follows:

Metric (B, D) = 7 >= Metric (A, D) = 5

Although the path A-B-D qualifies as a next-best metric path to the destination within the variance value, it is not selected since it does not pass the safety-check.

In contrast, the path A-C-D metrics from the figure “Multipath safety check” (page 63) result in the equation

Metric (C,D) = 4 < Metric (A,D) = 5

Since this path passes the safety check, it is selected as a safe alternative path.

Figure 17
Multipath safety check

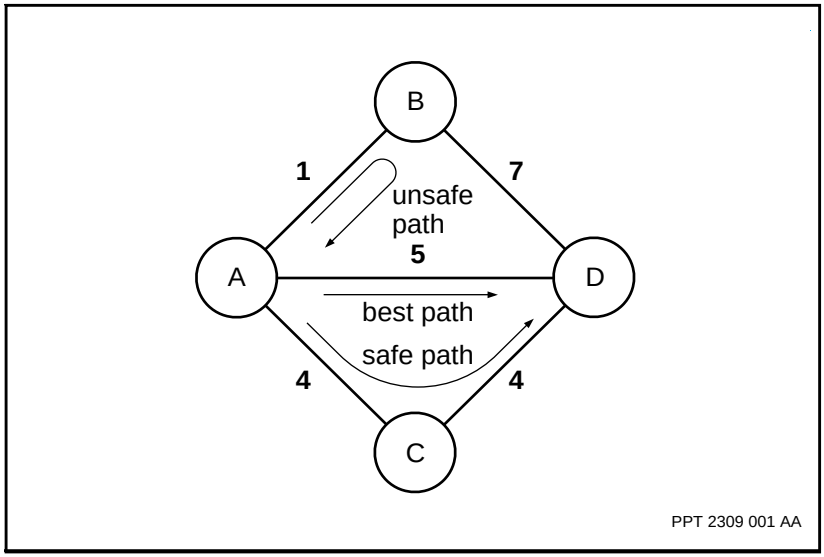

Traffic distribution

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

For example, if two paths have throughput metrics of 1 000 000 and 1 500 000, the first path receives 3/5 of the throughput traffic flows. The resolution of 6.25% results in 62.5% of the traffic flows using the better path. (This result does not always correspond to 62.5% of the actual total volume of traffic, since the traffic flows (VCs) can be of different sizes.)

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

Network routing with variance

When variance is deployed in a network, there is

- more spreading from edge nodes

  Edge nodes deployed around the periphery of the network backbone tend to be connected to the network with lower throughput links. This situation causes the first hop from the edge node to the network to be a more significant part of the overall metric of its paths. The application of the safety check at the edge nodes typically provides a safe alternative path for many network destinations. There is better spreading from the edge of the network, and better spreading of traffic across the backbone.

- minimal fine-tuning

  When the deployment of variance in a network still does not produce a significant number of alternative paths that meet the safety check criteria, minor adjustments in the form of overrides can be made to create more multipaths. In the same network without variance, a significant override adjustment is needed to create multipaths.

- multiple alternative paths

  In a small topology, the application of the safety check formula can result in the asymmetrical selection of multipaths. For example, the safe path from node A to node D in the figure “Asymmetrical multipaths”

(page 65) does not necessarily qualify as a safe path from node D to node A. In a small network topology, the safety check allows for multipaths only in one direction. In the overall network, there are more paths to choose from and it is more common to find a safe path in both directions.

Figure 18
Asymmetrical multipaths

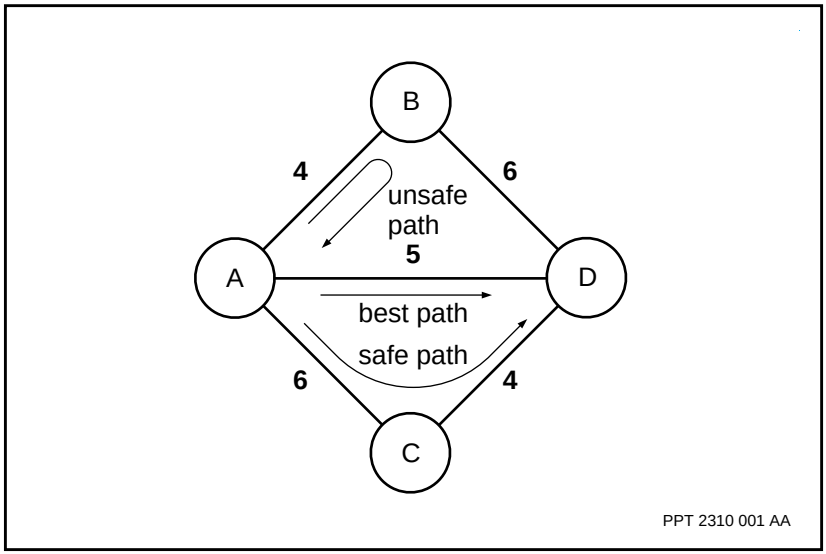

Variance with the delay RCOS

Delay metrics are more sensitive than throughput metrics to minor variation in delays between different link groups. This sensitivity means that equal-cost multipaths are rare for the delay RCOS in default routing behavior unless overrides are used.

When variance is in effect, two paths with minor differences in their delays can be used as equal-cost multipaths, allowing for more effective spreading of traffic. The unequal-cost alternative path must meet the safety check before variance can select both paths.

In the figure “Variance with delay RCOS (ms refers to milliseconds)” (page 66), all links have the same speed (1917000 bit/s), so throughput metrics are equal; however, one of the four links has a slightly larger delay value than the others. The path delay metrics on the two paths from node A to node D are 2000 and 2500. If node A is provisioned with a delay variance value of 25%, delay traffic can go over both paths. Since the throughput path metrics of the two paths are equal, throughput traffic flows are evenly spread across the two paths.With variance provisioned, delay traffic flows are also spread across the two paths, with slightly more traffic (9/16 of the traffic flows) on the better path.

Figure 19
Variance with delay RCOS (ms refers to milliseconds)

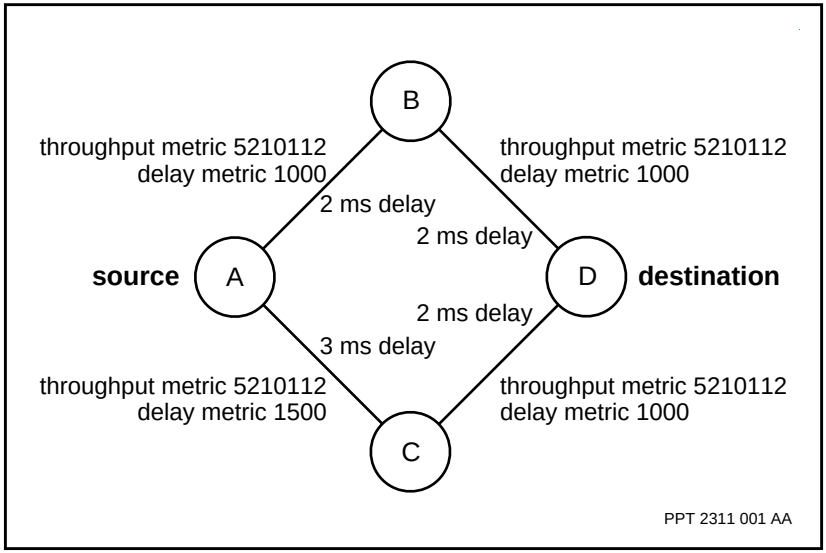

Variance with reserved PORS bandwidth

Even if two paths between two nodes have equal bandwidth, the amount of bandwidth reserved for PORS can differ on each path. Since the percentage of bandwidth allocated to PORS is subtracted from the overall trunk bandwidth to determine the bandwidth available for DPRS, this situation can result in the loss of multipaths between the two nodes. Variance can play a role in this situation to reestablish multipaths between the two nodes.

In the figure “Variance with reserved PORS bandwidth” (page 67), all links have a measured throughput of 1 917 000 bit/s and a measured delay of 2 milliseconds, but the calculated metrics of the two paths from node A to node D are different because of differing percentages of reserved PORS bandwidth. The best total path throughput metric is 20856832, and the next best path throughput metric is 23461888. With a throughput variance value of 13% provisioned on node A, both paths can be used for throughput traffic.

In the figure “Variance with reserved PORS bandwidth” (page 67), the best total path delay metric is 1070000 and the next best is 1337000. Provisioning a delay variance value of 25% allows both paths to be used for delay traffic. Since the throughput metrics of the two paths are close, the throughput traffic flows are evenly spread across the two paths. Delay traffic flows are also fairly evenly spread across the two paths, with slightly more traffic (9/16 of the traffic flows) on the better path.

Figure 20
Variance with reserved PORS bandwidth

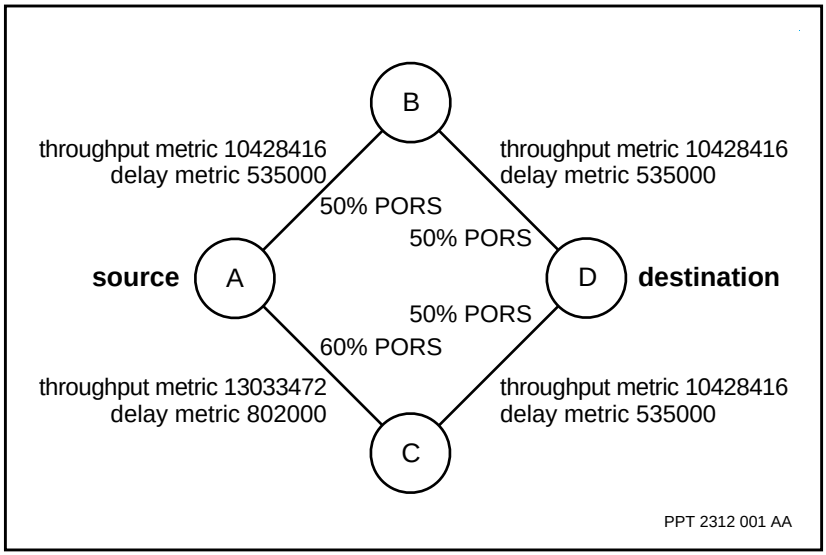

Multipaths with more hops

Even an alternative path with a greater number of hops can be used with variance to provide an additional path. For example, in "Multipaths with more hops" (page 68), the total path metrics from node A to node E are 10 and 15. A variance value of 50% allows both paths to be used, spreading the traffic flows so that 5/8 of the traffic takes the best path and 3/8 takes the next best path. The safety check verifies that the best metric from node C to node E is less than the best metric from node A to node E. At node C, the safety check for the alternative path back through node A fails, since the best metric from the next-hop node, node A, to node E is not less than the best metric from node C to node E. The safety check prevents a loop between node A and node C.

Figure 21
Multipaths with more hops

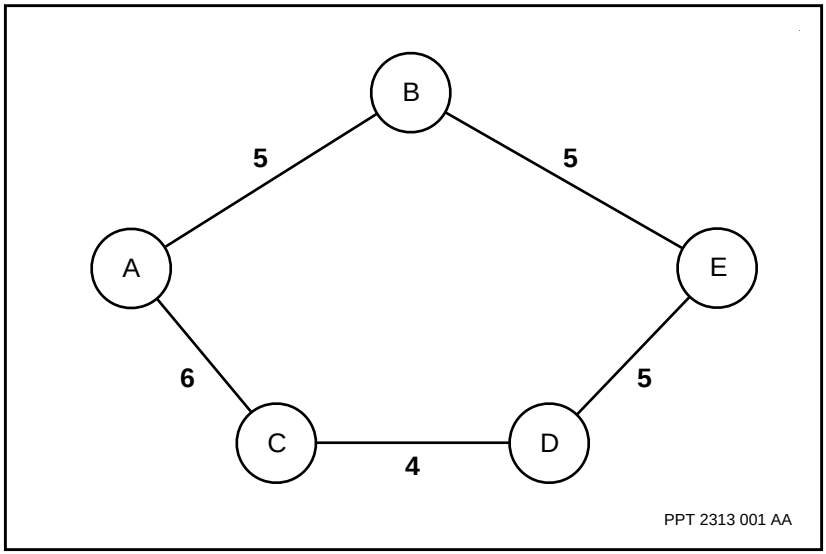

Variance values

To create situations in which an unequal-metric multipath can be used for traffic spreading, you can provision variance. The *variance* attribute value is the maximum difference in metric between the best path and the next-best path, defined as a percentage of the best-path metric.

You can provision two *variance* attribute values independently for each RCOS, to be applied to different ranges of delay and throughput metrics. Typically, a larger variance is more appropriate for small-delay paths than for large-delay paths. For example, with a 2 millisecond delay best path, a 4 millisecond delay can be acceptable for a second path. In this case, you provision a 100% variance value. However, for a 30 millisecond delay best path, a 100% variance allows a second path with a 60 millisecond delay, which can be unacceptable.

To implement this concept, you can provision a *variance* attribute value for high and low ranges of the delay metric. The *variance* attribute value for low-delay metric paths is used when the total best-delay path metric is less than the provisioned value for the *delayMetricRangeBoundary* attribute. The value for high-delay metric paths is used for all other delay paths. If no value is provisioned for low-delay paths (it is set to the default of 0), the provisioned value for high-delay paths is used for all paths. If the maximum value of 2147483647 is provisioned for this attribute, the *variance* attribute provisioned value for the delay RCOS and the low-delay value are used for all delay paths.

Similarly, it is possible to provision a larger variance for a high-throughput path (low-throughput metric) than for a low-throughput path (high throughput metric), so you can provision *variance* attribute values for high and low ranges of throughput metrics. The value for low-throughput metric paths is used when the total best-throughput path metric is less than the value for the *tputMetricRangeBoundary* attribute, and the value for high-throughput metric paths is used for all other throughput paths. If no value is provisioned for low-throughput paths (it is set to the default of 0), the provisioned value for high-throughput paths is used for all paths. If the maximum value of 2147483647 is provisioned for this attribute, the *variance* attribute provisioned value for the throughput RCOS and the low-throughput value are used for all throughput paths.

The *variance* attribute value can be defined as

- 0, the default, to specify that only equal-metric paths can be used, and no alternatives are permitted
- a percentage value to be applied as the percentage of the best-path metric

- 9999% to specify infinite variance, so that any safe path can be used as an alternative path.

Congestion management

When congestion occurs in the network, DPRS uses various control mechanisms. The methods of handling congestion differ, depending on whether load sharing or load spreading is in effect.

For more information on congestion management, see the following sections:

- "Congestion under load spreading" (page 70)
- "Congestion under load sharing" (page 71)
- "Congestion of multimedia traffic" (page 71)

Congestion under load spreading

When one link in the link group is congested, both normal and high-reliability traffic overflows onto other, uncongested links in the group. Traffic is designated as normal or high reliability through the normal reliability (NR) bit in the DPRS packet header.

If there are no uncongested links in the group, packet forwarding attempts to perform overflow routing of high-reliability traffic when an equal-cost metric path exists (or a second non-equal path, if variance is in effect). Normal-reliability traffic continues to travel on the initially-selected link. In this case, normal-reliability packets are either delivered or discarded, based on the packet's discard priority.

Originally, the access service derives the packet's discard priority from a combination of the discard priority (DP) bit in the packet's common header and the priority (Pri) bit in the DPRS header. If the priority bit is set, the service increases the original DP value from the common header by one level to make the packet less likely to be discarded. The new DP value is then used throughout the network. This discard status is compared to the link's level of congestion indicated in the link table to determine whether to discard the packet (if overflow routing is not possible).

At this stage of congestion, packet forwarding sets the forward congestion indication (FCI) bit in the Passport common header. This FCI bit indicates that congestion has occurred as the packet traversed the network. When the packet reaches its destination, the VC maps the FCI bit to a backward congestion indication (BCI) in the reverse traffic direction of the same connection.

Congestion under load sharing

When a link is congested, the loadsharing algorithm permits overflow onto another link within the group. If the traffic has high-reliability status, DPRS attempts to route the traffic onto an equal-cost alternative path.

Congestion under LoadSpreadFast

When the link selected is congested, a comparison is made between the discard priority of the frame and the discard level of the link. A frame with a priority higher than the discard level of the link has the FCI bit set and then is sent out of the link. Otherwise, the frame is discarded.

Congestion of multimedia traffic

DPRS uses special packet forwarding behavior for multimedia traffic to guarantee ordered packet delivery with minimal delay. DPRS uses load spreading for multimedia traffic to keep all traffic for a particular VC on the same link in a link group. Under congestion, multimedia traffic does not overflow to another link in the link group, or to an alternative link group, because the resulting packet disordering is unacceptable to the multimedia application.

For more information on Passport congestion management, see *241-5701-400 Passport 7400, 15000, 20000 Networking Overview*

Tandem suppression

Tandem suppression 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, intrasubnet (MID) tandem traffic is blocked. The only traffic routed to the node is traffic destined for the node itself, or any neighboring RIDs.

Tandem suppression is particularly useful for small access Passport nodes and CPE edge nodes. The feature can control traffic flows by

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

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

The tandem suppression option is only effective for routing within a RID subnet (when MID routing is operating). For example, in “Tandem suppression” (page 73), node C is assigned the tandem suppression option, so traffic is not routed to that node from within RID 10 unless it is destined for node C. When a node is provisioned for tandem suppression, any affected DPRS traffic flows are immediately rerouted.

If the link between node C and node F (shown as a dashed line in “Tandem suppression” (page 73)) exists, the tandem traffic destined for RID 20 (or from RID 20 to other nodes in RID 10) is not suppressed. Instead, the traffic flows through node C. Since it is unlikely for a small Passport node or CPE device to be connected to two subnets, this is not likely to be a problem.

When adding a new node at the periphery of the network, it is important to be aware of the tandem-suppression configuration of the connected nodes. If the new node connects to only tandem-suppressed nodes, it is effectively severed from the network. For example, in “Isolated node” (page 73), node C becomes isolated if nodes A and D are tandem suppressed.

> ***Note:*** It is important to set up a robust network that can remain operational even when links go down. For example, assume that only node D in “Isolated node” (page 73) is tandem suppressed. If the link between node A and node C goes down, node C is still effectively isolated. This example shows a non-robust topology.

Figure 22
Tandem suppression

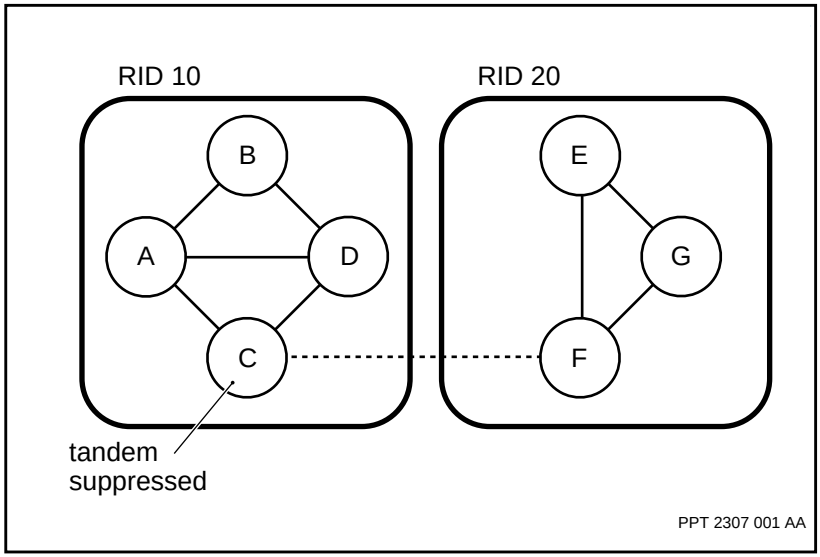

Figure 23
Isolated node

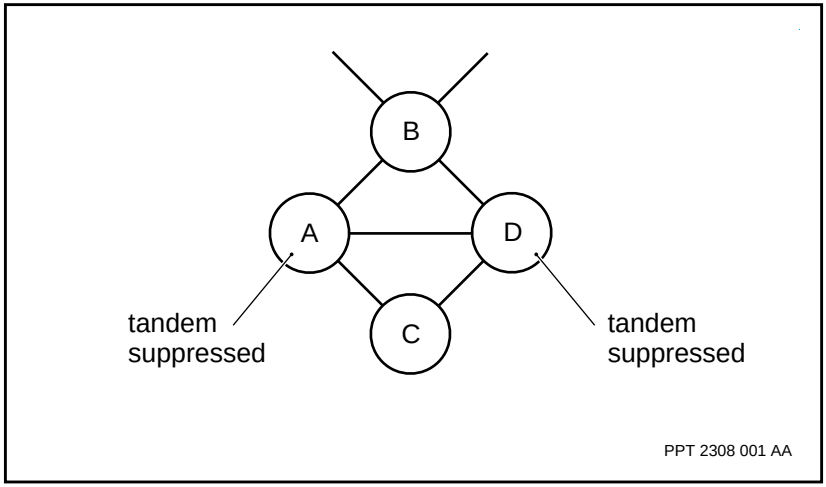

Delay and throughput cutoff

You can provision the delay and throughput cutoff values as a mechanism for fine tuning routing within a DPRS network. The RID routing system uses these values to determine at which point a switch is unreachable. If the metric value to reach the switch exceeds the cutoff for that RCOS, the switch is unreachable for that RCOS. (Even if a switch is unreachable on the delay RCOS because its delay metric exceeds the delay cutoff value, it can still be reachable on the throughput RCOS. The reverse can also apply.

The value of the delay and throughput metric cutoffs can affect the amount of routing traffic propagated throughout the network. The lower the cutoff value, the sooner a failed RID is unreachable, resulting in fewer routing table updates circulating around the network. This situation subsequently reduces the number of broadcasts that circulate around a RID subnet.

In a Passport-only network, when there are many nodes in all Passport subnets, it is unlikely for any RID in the network to disappear completely from the network. For this reason, changing the RID cutoff values is likely to have no effect on RID routing control traffic levels.

There are additional implications of a lower cutoff value. If a RID is unreachable in one RCOS, packet forwarding relies on default RCOS routing for routing traffic destined to the RID. (In default RCOS routing, DPRS sends delay traffic on a throughput path if there is no delay path available.) This situation does not always result in the best path being selected for the chosen RCOS.

Chapter 4
Network planning and engineering

For information on network planning for a Dynamic Packet Routing System (DPRS) network, see the following sections:

- “RID subnets” (page 75)
- “Guidelines for RID subnets” (page 78)
- “Guidelines for packet forwarding” (page 81)

For information on deployment strategies for a DPRS network, see the following sections:

- “RID filters” (page 84)
- “RID subnet splits” (page 92)

For information on route determination rules, see “Routing determinism” (page 97).

RID subnets

When you group Passport nodes into RID subnets it allows for very large network growth due to

- the efficient use of addressing space
- the advantages of hierarchical routing
- the efficient propagation and processing of network routing information

The network can support a maximum of 126 RID subnets. In a Passport-only network where Passport clusters are not deployed, each RID subnet can support an engineering maximum of 500 Passport nodes. If there are access modules (AM) in this network, each RID subnet can support an engineering maximum of 300 Passport nodes.

RID subnet design plays a key network scaling role in DPRS networks. Anticipation of significant network growth can help determine the RID subnet topology of the network. If you expect the number of Passport nodes in a RID subnet to increase significantly over time, you need to consider immediately partitioning the network into multiple RID subnets. This process can avoid the need for partitioning at a later date, affecting more services.

The decision to use one RID subnet or multiple RID subnets within a Passport network depends on many factors, and is often a network-specific decision. When you determine the breakdown of a network into RID subnets, RID subnet resiliency often plays a large part in the decision. Additionally, in large global networks, the geographical characteristics of the network can have an impact on the RID subnet partitioning.

You can partition a global network, such as the one shown in the figure "Global network with RID subnets" (page 77) into RID subnets. If you can divide the different geographical regions into separate RID subnets you will have a more resilient topology under link failure scenarios.

RID subnets provide an effective mechanism for regionalized networks to interconnect as seamlessly as possible, with minimum configuration and maximum redundancy. In addition, you can use RID filters to provide security at RID subnet borders, and control traffic flows through the network. (See "RID filters" (page 84).)

Figure 24
Global network with RID subnets

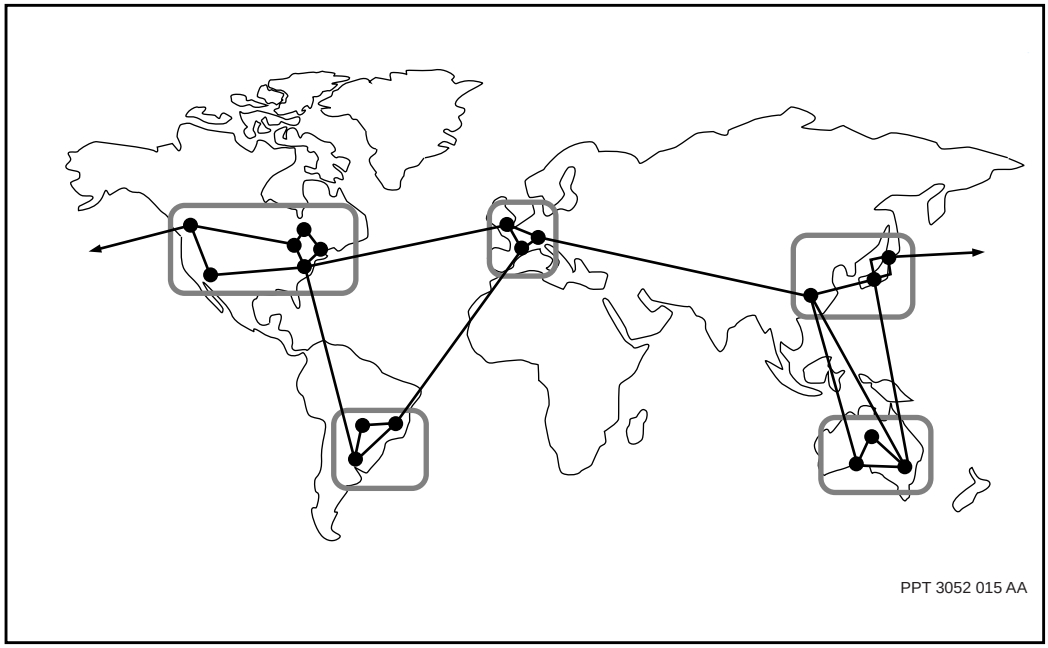

Passport clusters

Passport clusters are groups of Passport access switches connected to a RID subnet's backbone. Clusters do not exchange topology information with other clusters or the backbone, and only limited routing information is exchanged with the backbone.

A complete subnet, including subnet backbone nodes, Passport clusters and DPN-100 AM clusters, is limited by the number of MIDs in that subnet. Since Passport clusters are not included in the backbone's topology view, customers with networks that have large RID subnets can continue to grow their subnets beyond the engineering guidelines for existing subnets.

For information on configuring and deploying Passport clusters, refer to "Provisioning Passport clusters" (page 121).

Passport cluster paths

Under each of the RID, MID, and CS components, a new attribute *routeType* is introduced that shows these paths as learned paths or default paths. A learned path to a RID, MID, or CS is the best or next best path among all paths available to the destination found by DPRS. A default path is the best or next best path from a cluster node to the backbone. Since all the paths that reach a RID, MID, or CS are found independently by DPRS at a backbone node, the routeType value is always “learned” if these commands are issued at a backbone node. At a cluster node where the RID, MID, or CS is reachable within the local cluster, the paths to reach it are the best and next best paths within the local cluster, making the routeType value “learned”. All DPRS traffic to RIDs, MIDs, and CSs outside the cluster is forwarded to the backbone, making the paths to those destinations the default paths and the value of routeType “default”. Default route information is displayed if the value of routeType is “default”

Under each MID component, a new attribute *midType* is introduced. The value for this attribute can be “em”, “am”, “cluster”, and “unknown”, and is determined as follows:

- if the MID is a Passport backbone node, the *midType* is “em”
- if the MID is a Passport cluster node, the *midType* is “cluster”
- if the MID is a DPN-100 AM module, the *midType* is “am”
- if the MID is not known to the cluster (for example, the MID outside of the cluster), the *midType* is “unknown”

Guidelines for RID subnets

The following engineering guidelines apply to RID subnets:

- A Passport network supports 126 RIDs.
- A RID subnet supports an engineering maximum of up to 500 Passport nodes in a Passport-only network where Passport clusters are not deployed. In an interworked Passport and DPN-100 network, a RID subnet can support up to 300 backbone Passport nodes and 1600 DPN-100 access modules (AM).

- In an interworked Passport and DPN-100 network, a RID subnet with backbone and cluster Passport nodes supports an engineering maximum of up to 200 Passport nodes in the backbone, and 300 Passport nodes in each cluster.
- Each RID subnet can support a provisioned MID addressing space from 1 to 1909. The total number of MIDs in a RID subnet then, including all Passport backbone nodes, Passport cluster nodes, and DPN-100 AMs, cannot exceed 1909 MIDs. Each RID subnet is a separate routing area, allowing for MID re-use in different RID subnets.

To maintain the resilient connectivity between all nodes within a RID subnet, each Passport node must have at least two connections to two other Passport nodes within the same RID subnet. The only exception to this rule is singly connected CPE Passport nodes, which become isolated if a link failure occurs.

For example, in “Interconnecting RID subnets” (page 80), no single link failure in RID 20 or RID 30 can cause either subnet to be severed. (An exception is the link to the CPE node I, which can become isolated if the link between node D and node I fails.) The configurations in the figure show valid examples of interconnecting subnets:

- One or more Passport nodes in a RID subnet can connect to one or more Passport nodes in another RID subnet (for example, the connections between RIDs 20 and 30).
- A single Passport node in a RID subnet can connect to one or more Passport nodes in another RID subnet (for example, the connections between node A and nodes B and C). Each of the connections constitutes a separate link group to the neighbor RID subnet.
- A single Passport node in one RID subnet can connect through one or more connections to a single Passport node in another RID subnet (for example, the connections between nodes E and F). The links make up a single link group on each of the Passport nodes.

Figure 25
Interconnecting RID subnets

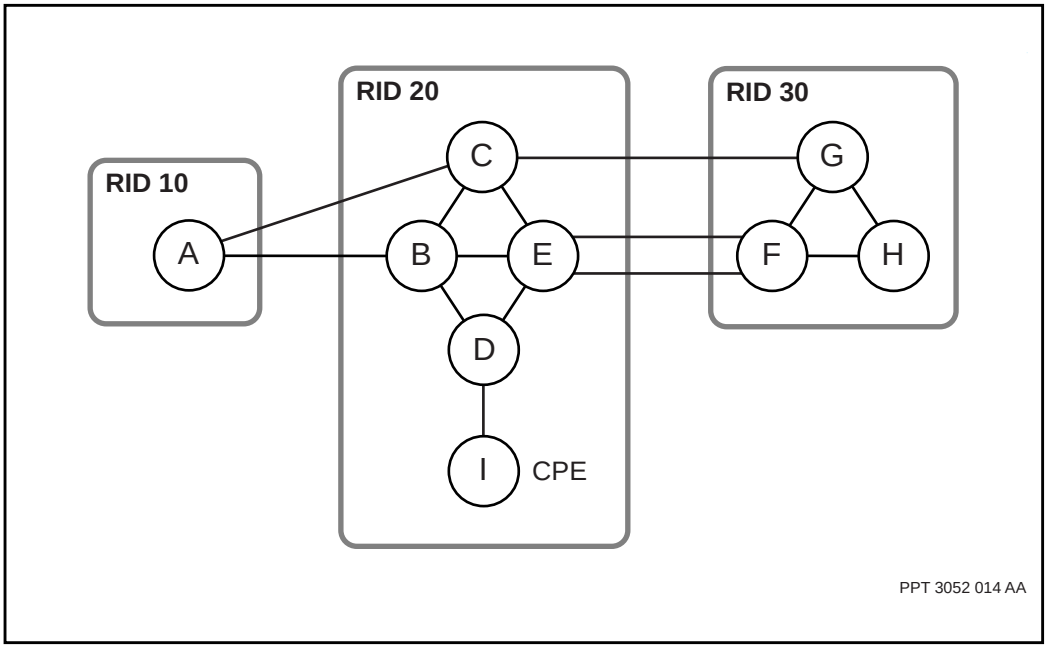

Guidelines for Passport clusters

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

- Passport trunks cannot be provisioned between clusters.
- Clusters should be configured around a topology region's or RID Subnet's backbone.
- RID subnets or topology regions should be interconnected via backbone nodes, not cluster nodes.
- Passport cluster nodes should be connected to at least two backbone nodes for redundancy purposes.
- DPN-100 CSRMs associated with a RID subnet are allowed to be connected to the backbone portion of a subnet only. In a mixed Passport and DPN-100 network, this means that all DPRS call request packets

generated by an access service in a cluster will be routed to a CSRM via the backbone nodes. No connections from a Passport cluster to a DPN-100 CSRM or RM are allowed. To prevent connections to DPN-100 RM or CSRMs from being accidentally established, links to these modules are automatically disabled on a cluster node.

- DPN-100 AM clusters are allowed to be subtended off of Passport clusters, and can also be connected to the backbone portion of a RID subnet. Routing behavior remains as before.
- All backbone nodes to be connected to a cluster node must first be upgraded to the same software level as the cluster node. If a backbone node has not been upgraded to the same software level as the cluster node to which it is attached, an alarm will be generated and trunks between the cluster node and backbone node will fail to stage.

Guidelines for packet forwarding

Each of the three packet forwarding algorithms, loadspread, loadspreadFast, and loadshare, is effective under different network configurations and circumstances. The packet forwarding algorithms are described under the following topics:

- "Selecting a packet forwarding method" (page 81)
- "Load spreading guidelines" (page 82)
- "Load sharing guidelines" (page 83)

Selecting a packet forwarding method

In load spreading, VC traffic flows are kept on the same links in a link group. The loadspread algorithm is optimal when many VCs are defined using access speeds that are less than the backbone trunk speeds. Load spreading is more effective when all links in a link group have approximately the same throughput and there are no large frame relay VCs. (A large frame relay VC is any VC whose CIR equals or exceeds 30 to 40% of the capacity of a single trunk anywhere along the VC path.)

An additional advantage of the loadspread algorithm is that there is no packet disordering, since all the traffic for a particular VC is routed along the same path through the network. This practice provides more efficient CPU consumption at the egress point.

The loadspreadFast algorithm can be selected to enable hardware forwarding and get improved performance on ATM IP and PQC12 based FPs.

> ***Note:*** For Passport 15000 and 20000, loadspreadFast is the only packet forwarding policy available.

In load sharing, all the links in the link group share the traffic flow from each VC. Use load sharing under the following conditions:

- different links in the group have different bandwidths
- access bandwidth of a VC exceeds the bandwidth of individual link capacity

Load sharing is a good solution when access port speeds equal the backbone speeds. Load sharing proportionally shares traffic across links in a link group based on link capacity. Since the packets from a VC are sent on different links across a link group, minor disordering can occur. However, the destination VC provides ordered delivery to the access service.

Load sharing is advantageous since it provides a more even distribution of frame relay traffic across all links in a link group. This balancing effectively provides an inverse multiplexing capability for each packet over the total capacity of the link group. This method enables a single VC to have a bandwidth greater than that of an individual link.

Load sharing is also more sensitive to PORS traffic. Reserved PORS bandwidth on a link affects the link's capacity to carry DPRS traffic. This situation can lead to different links in the group carrying different volumes of DPRS traffic.

Load spreading guidelines

Since the loadspread algorithm attempts to evenly spread all VC traffic across all links within the link group, it is important to ensure that the links have relatively equivalent characteristics. For the throughput RCOS it is important to ensure that the links have relatively equivalent bandwidth. For the delay RCOS, it is important to ensure that the links have relatively equivalent round trip delay values. In addition, since load sharing ideas take the reserved PORS

bandwidth into consideration, except in the selection of its preferred links, the same amount of reserved PORS bandwidth must be configured on each link. See “Preferred link selection” (page 55).

When loadspread is the selected packet forwarding algorithm, consider the following guidelines:

- Use equal-speed links in link groups whenever possible.
- If PORS traffic will use 80% or more of trunk bandwidth, provision the trunk as PORS only. This recommendation is particularly relevant when the PORS traffic comes from a single PORS PLC.
- In a link group carrying PORS traffic, configure all trunks with the same amount of reserved PORS bandwidth (except for any link configured as PORS only).
- If you are using an ATM IP FP, use the loadspreadFast algorithm to enable hardware forwarding.

Load sharing guidelines

In a network with mixed PORS and DPRS traffic, load sharing can balance the overall load on the links in a link group. The effectiveness of this method depends on the proportion of PORS and DPRS traffic on each functional processor (FP).

Incoming PORS traffic flows (VCs) on the ingress trunk FP are statistically spread across the outgoing link groups and links. If the number of PORS paths is small, the traffic is not always shared evenly across the outgoing links in a link group. DPRS load sharing responds to any imbalance in PORS traffic spreading by directing DPRS traffic to under utilized links in the link group.

When running the DPRS packet forwarding algorithm set to loadshare, the speed ratio of the links in the link group must, for high speed trunks, be set according to the following guidelines.

- Load sharing is not supported on OC12 trunks unless maxSubnetSize <= 512, and all trunks in the trunk group have the same bandwidth.
- Load sharing is supported on OC3 trunks at maxSubnetSize <= 4K if all the trunks in the trunk group have the same bandwidth. If maxSubnetSize <= 512, then a mix of 0C3 down to DS3/E3 is supported.

- Load sharing is supported on DS3 / E3 trunks at maxSubnetSize <= 4K if all the trunks in the trunk group have the same bandwidth. If maxSubnetSize <= 512, then a mix of DS3 / E3 down to DS1 / E1 is supported.

RID filters

Normally, all connected neighbor nodes that are in different RID subnets exchange RID routing information, in the form of routing table updates (RTUs). RID filters allow you to restrict the propagation of RID routing information between RID subnets. This restriction provides network security, as well as additional network efficiency in RID subnets. For a method of further enhancing network security with topology regions, see 241-5701-400 *Passport 7400, 15000, 20000 Networking Overview*.

RID import lists

RID filters ensure routing protection at designated boundaries within the network. In a network alliance situation, provisioning an import filter on the border nodes of two interconnected networks ensures that only the anticipated routing information is processed and distributed in the rest of the network. RID filters ensure that a RID subnet receives only the routing updates from a RID subnet that affect RIDs in the filter import list. This practice ensures more efficient use of bandwidth in the RID subnet, and reduces the amount of CPU consumed to process routing updates.

By default, new RID filters contain all RIDs, so they do not change the default behavior of the routing system. If a RID is removed from a neighbor node's filter import list, the neighbor node ignores RID routing information from that RID. For example, in the figure "RID import filters" (page 86), a RID import filter is provisioned on node B in RID 20 for its neighbor node in RID 30. Originally, the list included RIDs 10, 50, 60 and 70. When RID 60 is removed from the node B import list, node B ignores RID routing information about RID 60. In this situation no traffic destined for RID 60 uses the link between node A and node B. As long as RID 60 information does not arrive from any other source, node B does not put RID 60 in its forwarding table or send RID 60 traffic to node C in RID 30.

The figure "RID import filters used in merged networks" (page 87) illustrates individual networks connected through a carrier network. The carrier network has import filters provisioned on its border nodes, so it accepts only routing

information from the nodes that are in the RID import lists. Other networks provision import filters on their nodes in order to control the RID routing information that is routed from the carrier network to their network.

The carrier network includes RID 10 and 20 in its import list for network A and excludes RIDs 12 and 13. If the RIDs that belong to network A's lab network are propagated to the carrier network, they are not propagated within or beyond it since they are not in the carrier network's RID import filter list. This practice ensures that Network A cannot cause routing problems for the other networks.

Networks A and C can route calls to and from each other through the carrier network because network A includes network C's RIDs in its RID import list. Network C includes network A's RIDs in its RID import filter.

Network B is able to make calls only to the carrier network and is unable to connect to networks A and C because it includes only the RIDs from the carrier network in its RID import filter.

Figure 26
RID import filters

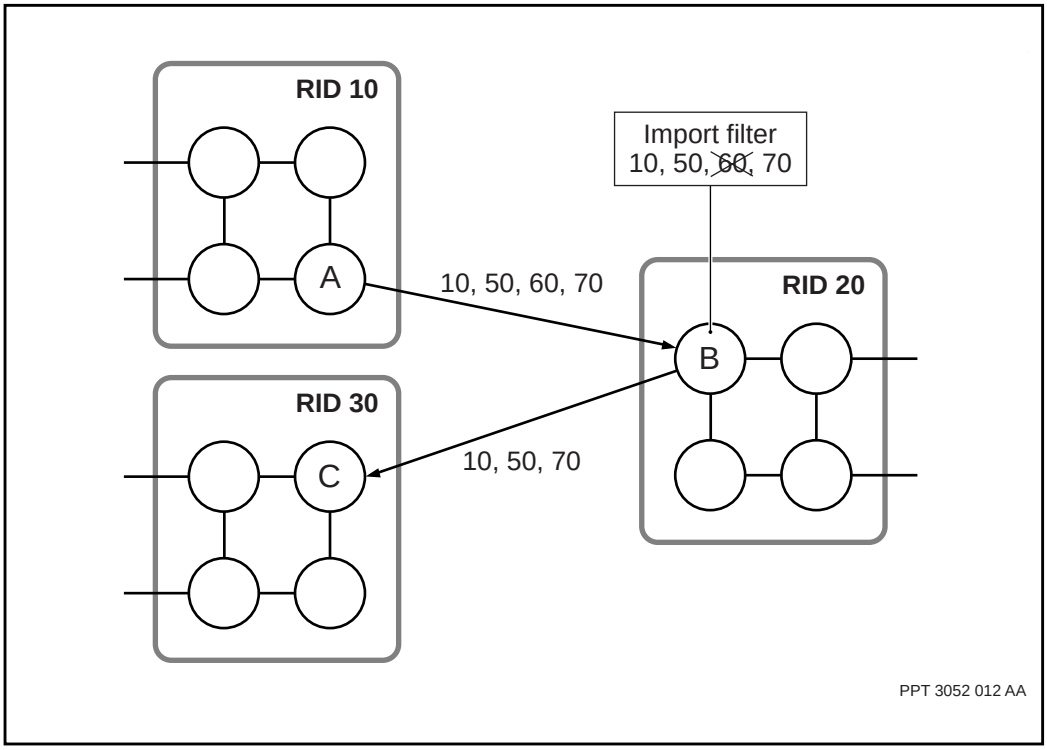

Figure 27
RID import filters used in merged networks

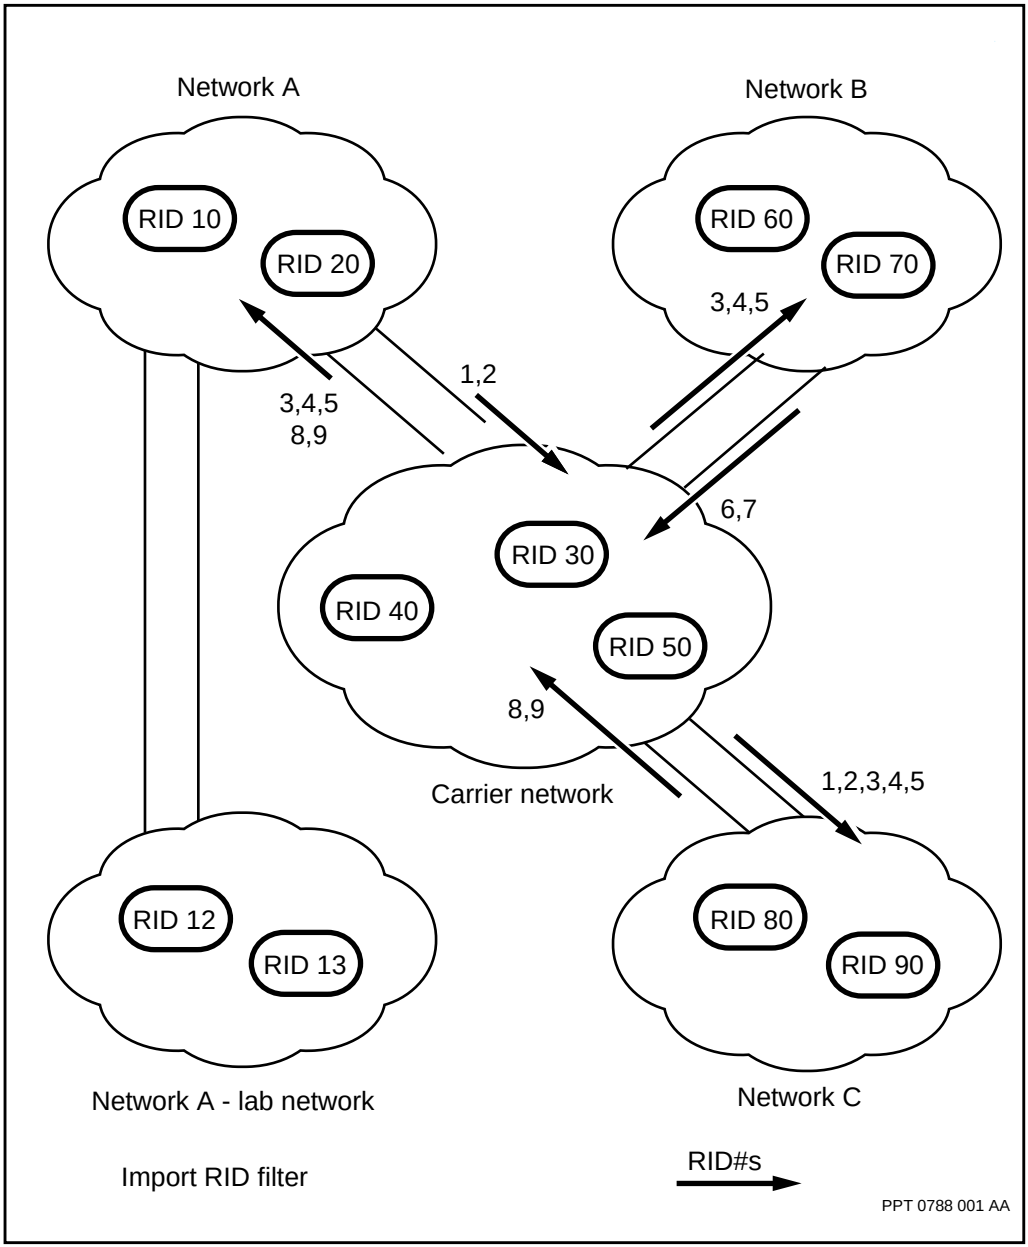

RID export filters

Export filters provide an additional level of efficiency over the routing of control traffic in the network. The use of export filters shifts the control of routing information to the originating gateway node.

A RID export filter controls the traffic that crosses an inter network trunk. If a RID is removed from the export list for a neighbor node, information about that RID does not go to the neighbor. In the example in the figure "RID export filters" (page 88), an export filter list is provisioned on node A in RID 10. When RID 60 is removed from the export list, no information about RID 60 goes from RID 10 to RID 20.

Figure 28
RID export filters

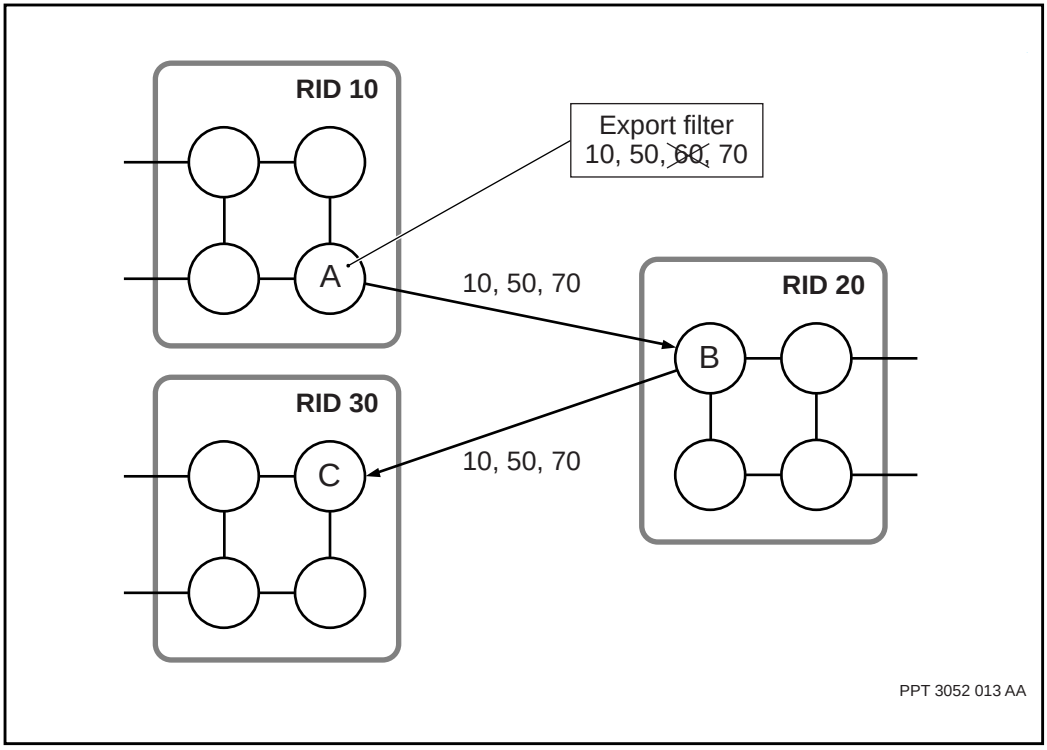

The figure “RID import and export filters in merged networks” (page 90) represents one network that spans two geographically separate locations. The carrier network carries information between the two locations but no traffic terminates on the carrier network.

In the illustration, networks A1 and A2 only use the carrier network to tandem traffic back and forth. The RIDs of the carrier network do not need to be included in its RID export filters because the carrier network provides a backbone network for networks A1 and A2 and its own routing system can remain transparent to the other networks.

To prevent routing updates for the lab network RIDs 12 and 13 propagating to the carrier network, network A1 provisions an export filter that excludes RIDs 12 and 13.

Figure 29
RID import and export filters in merged networks

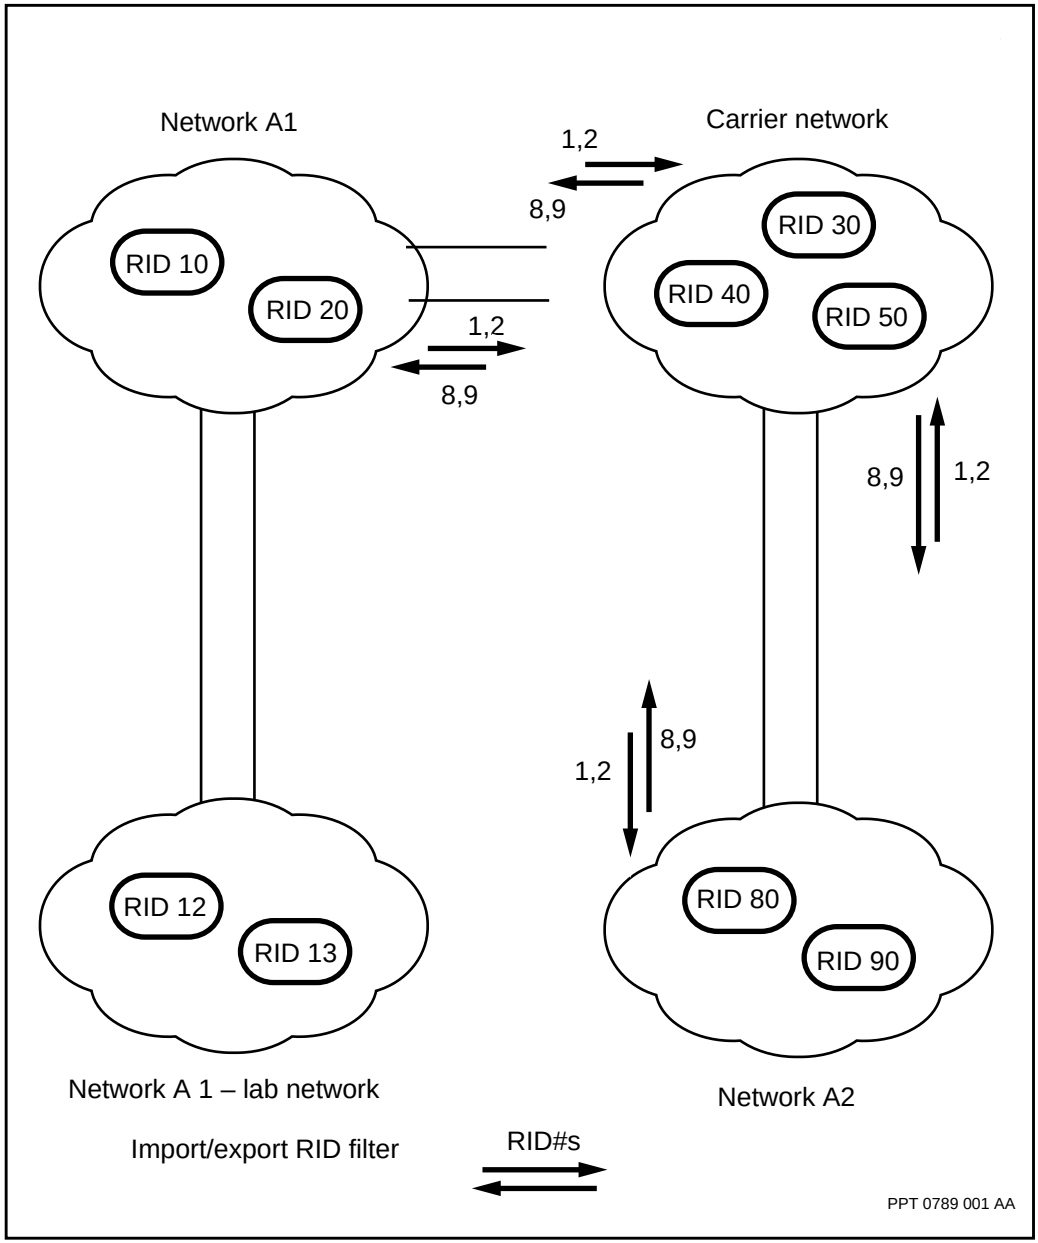

RID filters creation

By default, RTUs are imported from and exported to all neighbors unless they are provisioned with import or export RID filters.

If RID filters are incorrectly provisioned, routing paths can be lost. Issuing a ping command from the border nodes of the RID subnet to all the RIDs that the nodes should be able to reach will ensure RID filtering is correctly provisioned.

RID subnet splits

RID splitting is the process of redeploying some nodes in a large RID subnet to a separate RID subnet. the figure “Splitting a RID subnet” (page 93) shows a simple example of RID splitting. In the figure, stage 1 shows the original RID subnet. In stage 2, one node has been split off to form a separate RID subnet.

The RID redirection capability of the call redirection server (CRS) provides an efficient method of handling the splitting process with minimal operation impact. During the RID-splitting process, the CRS temporarily redirects calls destined for nodes in the original destination RID to the new RID. After the relocation process is complete, the call router databases in both RIDs are updated, and the CRS is no longer needed. For a detailed explanation of RID redirection, see 241-5701-410 *Passport 7400, 15000, 20000 Call Redirection Server Guide*.

> ***Note:*** You can merge two RID subnets using a procedure similar to the procedure for RID splitting. The main difference between RID splitting and RID merging is that you update the call router database to reflect the new RID configuration before migrating the last node of the merged RID. The last node to migrate from the merged RID is the one on which the CRS is defined. For details on merging RID subnets that contain clusters, see “Merging two RID subnets that contain Passport clusters” (page 95)

For more information on RID splitting, see the following sections:

- “Preparing to split a subnet” (page 93)
- “Splitting a RID subnet” (page 122)
- “Splitting a RID subnet that contains Passport clusters” (page 94)

Figure 30
Splitting a RID subnet

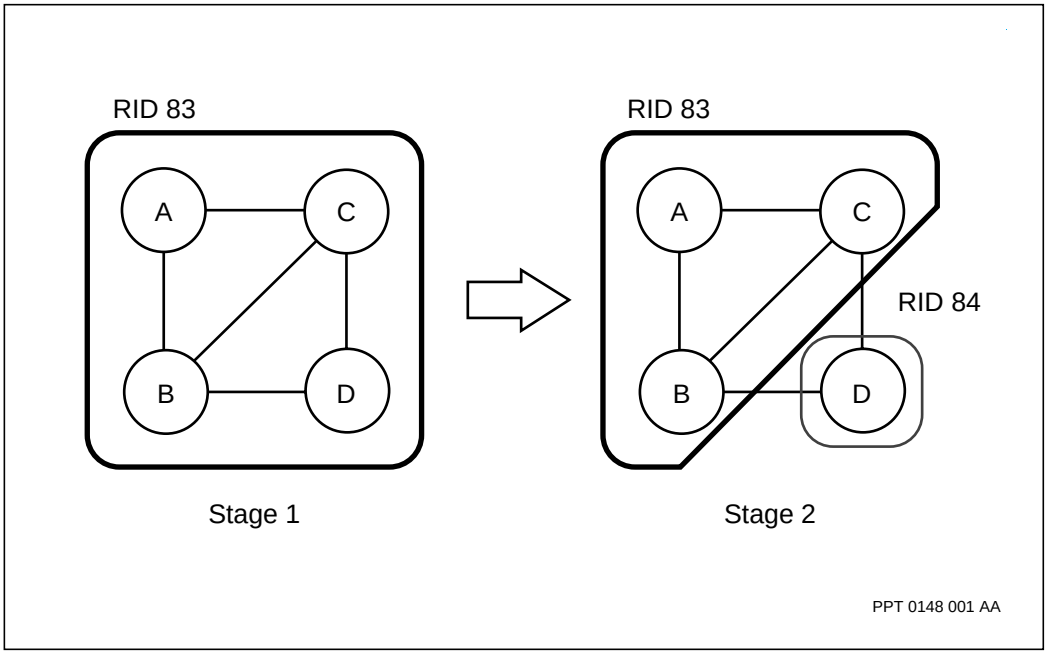

Preparing to split a subnet

Before you begin the RID-splitting process, complete the following tasks:

- Make sure you do a traffic analysis to understand how traffic will flow when the existing RID subnet becomes two RID subnets. Network analysis and planning is needed before beginning the RID-splitting process to avoid problems in intersubnet and intrasubnet traffic flow.
- Make sure that the call router database is consistently and accurately provisioned across the original RID subnet.
- Plan the RID split with the objective that a configuration of multiple islands of nodes belonging to the same RID is not allowed. All the nodes that belong to a RID must directly connect to a node of the same RID at all times (except when only the first node has been split to the new RID subnet).

- Before you move a node, make sure that the RID/MID combination for the node to be moved is not represented in the *secondaryRidMid* component of the CRS database. If so, you must update the CRS database to reflect the backup RID (across both the original and new RID subnets) as the node is moved.
- If the MID value of the node to be moved needs to be changed, do it either before or after the RID splitting operation, not during the process.
- Make sure that the MIDs of the nodes to be moved do not already exist in the new RID subnet.

Splitting a RID subnet that contains Passport clusters

The operating and provisioning steps for merging or splitting a RID subnet that contains clusters remains unchanged from the steps for RID subnets that do not contain clusters. However, the order in which the Passport nodes have their RID value changed is important. During the migration, each cluster node must maintain at least one path to reach a backbone node within its own RID subnet.

When splitting a RID subnet that contains clusters, start by changing the RID of the backbone nodes. Otherwise, some cluster nodes will be isolated from the rest of the network. After all of the backbone nodes for the new RID subnet have been changed, start changing the RID of the cluster nodes, ensuring that the remaining cluster nodes maintain a path to the RID subnet in the backbone. Disconnect clusters with different RIDs once all RID changes are completed.

Figure 31
Example of splitting a RID subnet that contains clusters

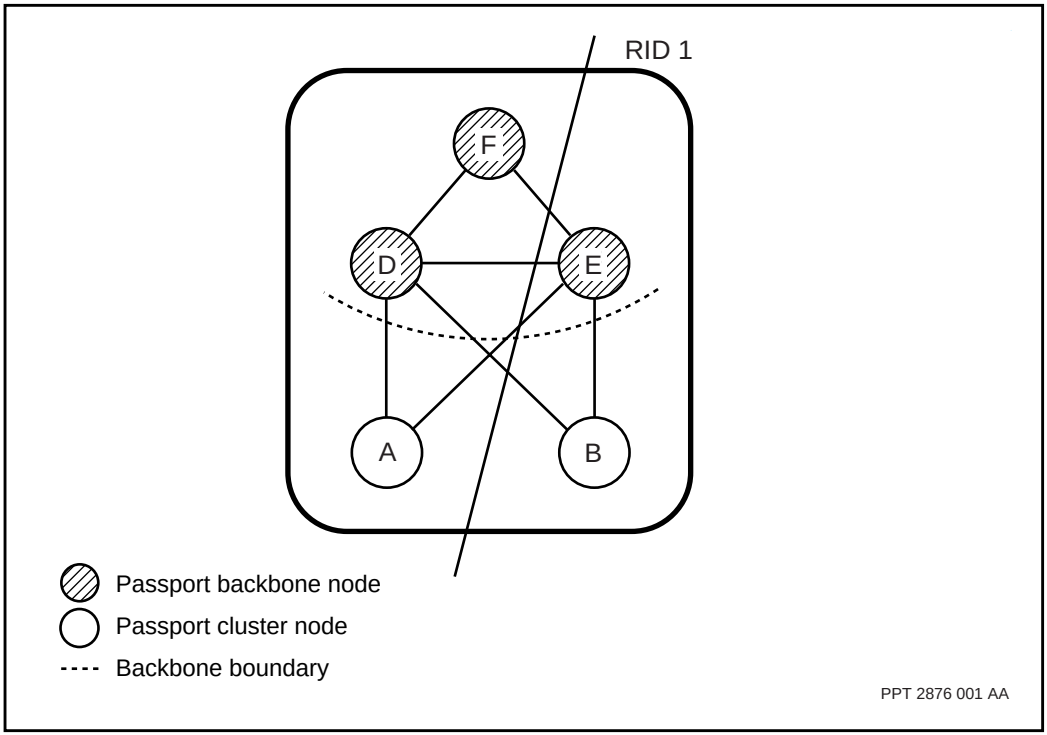

To split the RID subnet containing Passport clusters shown in “Example of splitting a RID subnet that contains clusters” (page 95), one would take the following steps:

1 Change the RID of node E

2 Change the RID of node B

3 Disconnect link A-E and link B-D

Merging two RID subnets that contain Passport clusters

When merging two RID subnets that contain clusters, start by changing the RIDs of nodes on the subnet border, with the RIDs of backbone nodes being changed first, followed by that of cluster nodes. Follow this by changing the RIDs of backbone nodes and cluster nodes which lie on the newly created subnet border.

Figure 32
Example of merging two RID subnets that contain Passport clusters

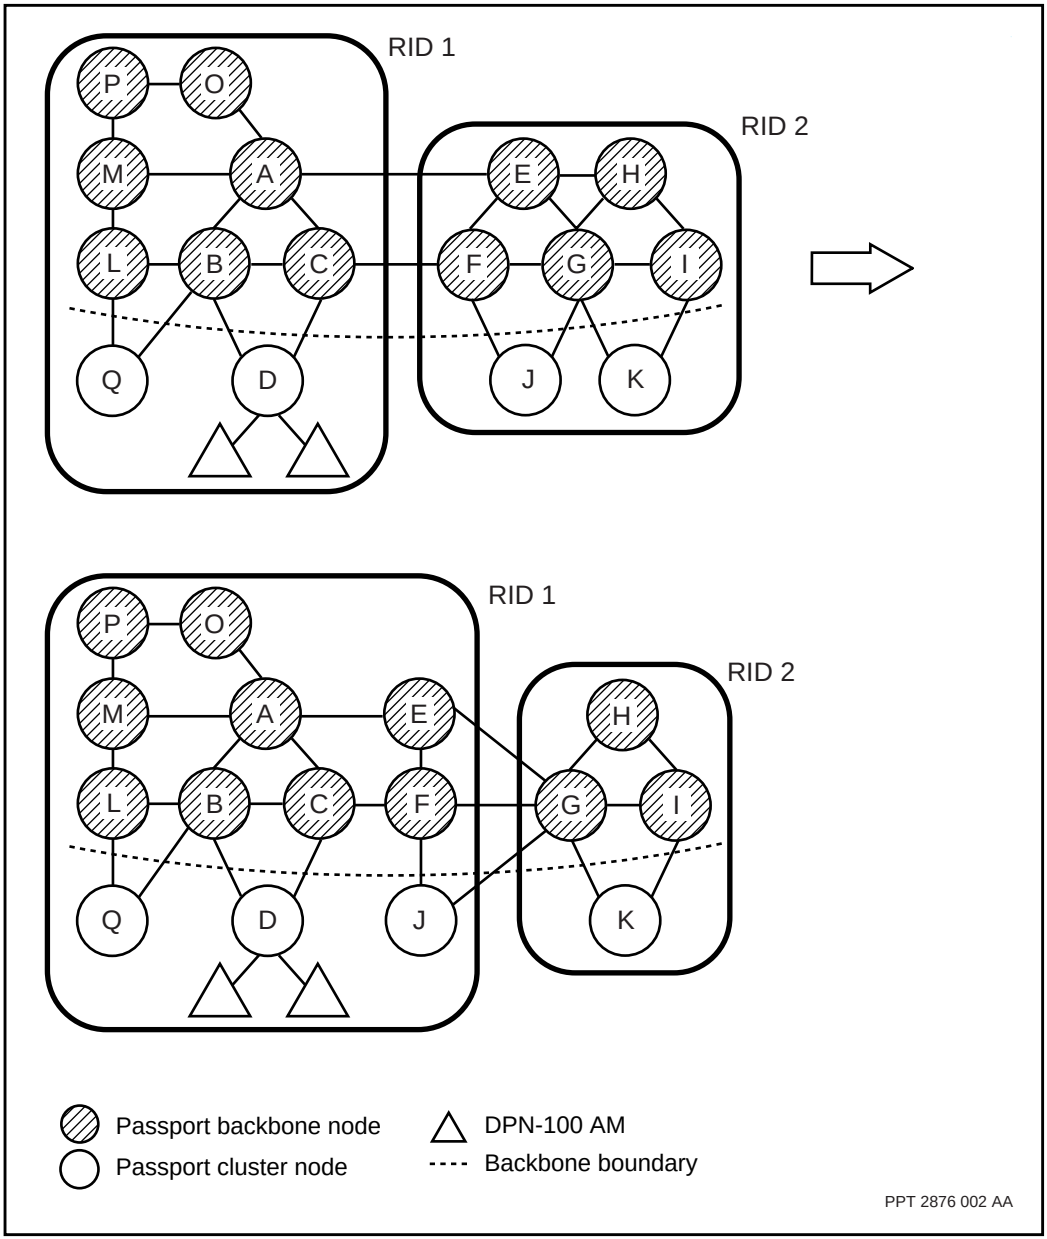

In “Example of merging two RID subnets that contain Passport clusters” (page 96), RID subnet 2 is to be merged into RID subnet 1, such that RIDS of all nodes in RID 2 are changed to RID 1. To merge the two RID subnets, the following steps would be taken:

1 Change the RIDs of backbone nodes E and F, as they lie on the subnet border.

2 Change the RID of cluster node J, which still lies on the subnet border.

3 A new subnet border is created as a result of steps 1 and 2.

4 Change the RID of backbone nodes H and G, as they lie on the new subnet border.

5 Change the RID of cluster node K.

6 Change the RID of the backbone node I.

Routing determinism

If there are more than two equal minimum-metric paths for a route, DPRS uses routing determinism algorithms, or rules, to choose the paths for the forwarding tables. This practice means that traffic patterns can be predictable. For example, the traffic routing will be the same after a link failure and recovery as it was before.

When the variance option is on, DPRS can select two or more next-best paths to a destination. In that case, the determinism rules choose between multiple, equal next-best paths.

These determinism rules are based on the location in the network, the destination, and the RCOS. The figure “Determinism rules” (page 98) shows the four possible routing situations and their corresponding determinism rules:

- Node to node. See “Node to node (rule 1)” (page 99).
- Node to gateway node (a node with a link to another RID subnet is a gateway node). See“Node to gateway node (rule 2)” (page 100).
- Gateway node to RID. See “Gateway node to RID (rule 3)” (page 101).

- Gateway node to gateway node. See “Gateway node to gateway node (rule 4)” (page 102).

Routing within a RID subnet depends on whether the destination is within the subnet or outside it. If the destination is inside the subnet, DPRS uses rule 1 to decide on the destination node. If the destination is outside the subnet, DPRS applies the rules to pick gateway nodes (rule 2) first, and then the rules to route between the nodes in a subnet (rule 1).

When DPRS is routing packets from gateway nodes, it uses rules based on the destination RID and next-hop RIDs (rule 3), and then next-hop MIDs (rule 4).

Figure 33
Determinism rules

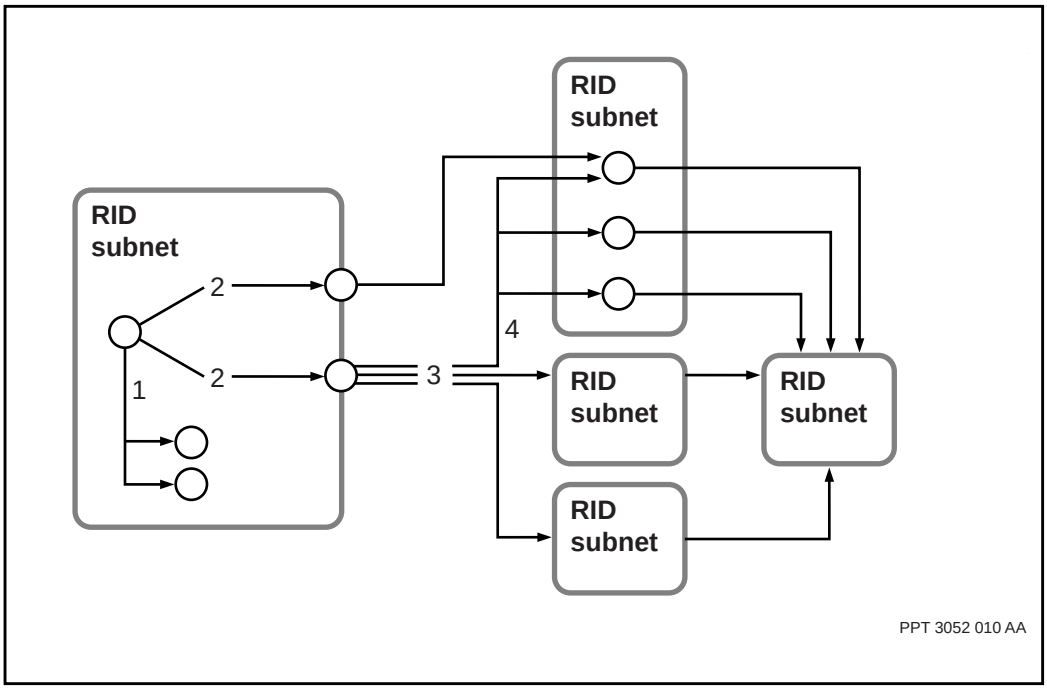

Node to node (rule 1)

When more than two equal paths exist to a destination within the subnet, the topology manager selects paths based on the following rules:

- Choose the two neighbor nodes with the minimum next-hop delay metrics.
- If there is a tie (three or more paths with equal minimum metrics), select the two with minimum next-hop throughput metrics.
- If there is still a tie, select the two paths with the highest Passport nodeIds (highest nodeId first, second-highest nodeId second).

In the topology in the figure "Node to node" (page 100), the topology manager selects multipath neighbors on node A to destination node G as follows:

- From node A, the set of possible multipath neighbors to node G is {B, C, D, E, F}.
- After the first rule of selecting minimum next-hop delays, the set is {C, D, E, F}.
- After the second rule of selecting minimum next-hop throughputs, the set is {D, E, F}.
- After the third rule of selecting the two highest neighbor nodeIds, the selected multipaths are {F, E}, in that order.

If the tie is for best path, these two paths are used, and appear in the table in this order. If the tie is for next-best path, the first path selected is used as the next-best path.

Figure 34
Node to node

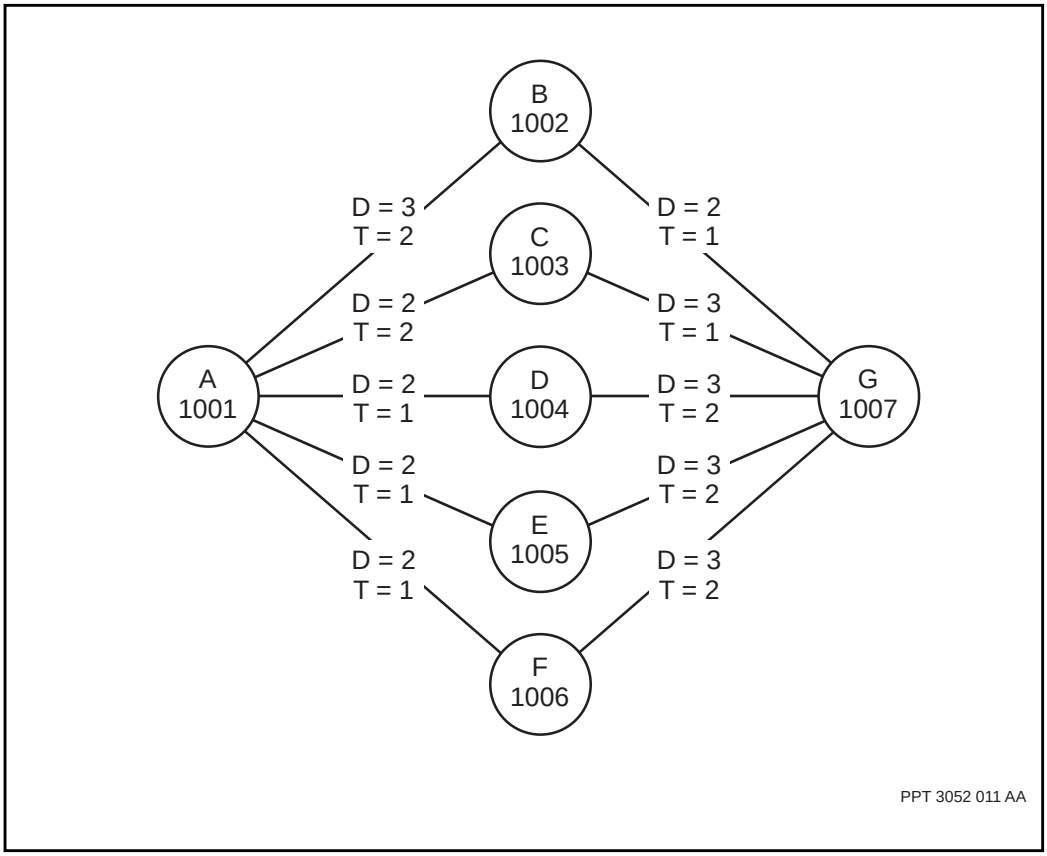

Node to gateway node (rule 2)

When there are two or more equal metric paths through different gateway nodes, DPRS selects

- first, the path through the gateway node with the lowest nodeId
- second, the path through the gateway node with the highest nodeId

If the tie is for best path, these two paths are used, and appear in the table in this order. If the tie is for next-best path, the first path selected (the one with the lowest gateway nodeId) is used as the next-best path.

In the figure “Node to gateway node” (page 101), DPRS selects paths through nodes B and D (in that order) for routing traffic from node A to RID 90, since nodeId 1002 is the lowest and nodeId 1004 is the highest.

When there are two different gateway nodes on equal paths to a destination, it is also possible that there are multiple paths to each gateway node or multiple gateway paths from one or more of the gateway nodes. When there are multiple paths to each gateway node, DPRS uses the first path selected by rule 1. When there are multiple gateway paths, DPRS uses the first path selected by either rule 3 or rule 4.

Figure 35
Node to gateway node

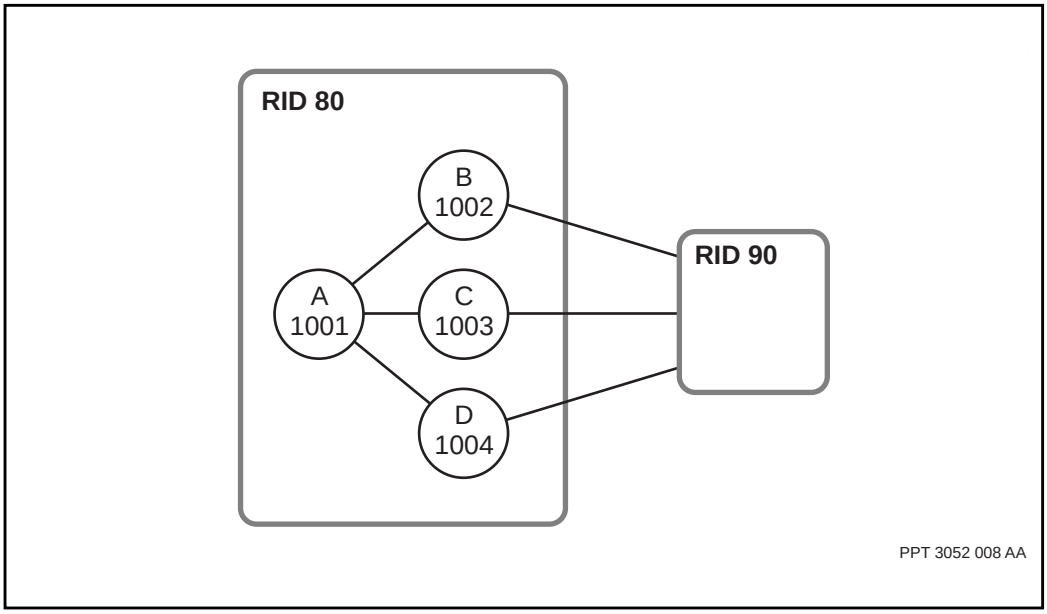

Gateway node to RID (rule 3)

When there are more than two equal cost paths through more than two gateways to other RIDs, DPRS selects

- for delay RCOS, the two paths through the neighbors whose RIDs are lowest if the destination RID is even and highest if the destination RID is odd

- for throughput RCOS, the two paths through the neighbors whose RIDs are highest if the destination RID is even and lowest if the destination RID is odd

In the figure “Gateway node to RID” (page 102), DPRS at node A picks the paths through RIDs 10 and 20 to reach RID 40, or the paths through RIDs 30 and 20 to reach RID 41 for the delay RCOS. For the throughput RCOS, the reverse is true.

Figure 36
Gateway node to RID

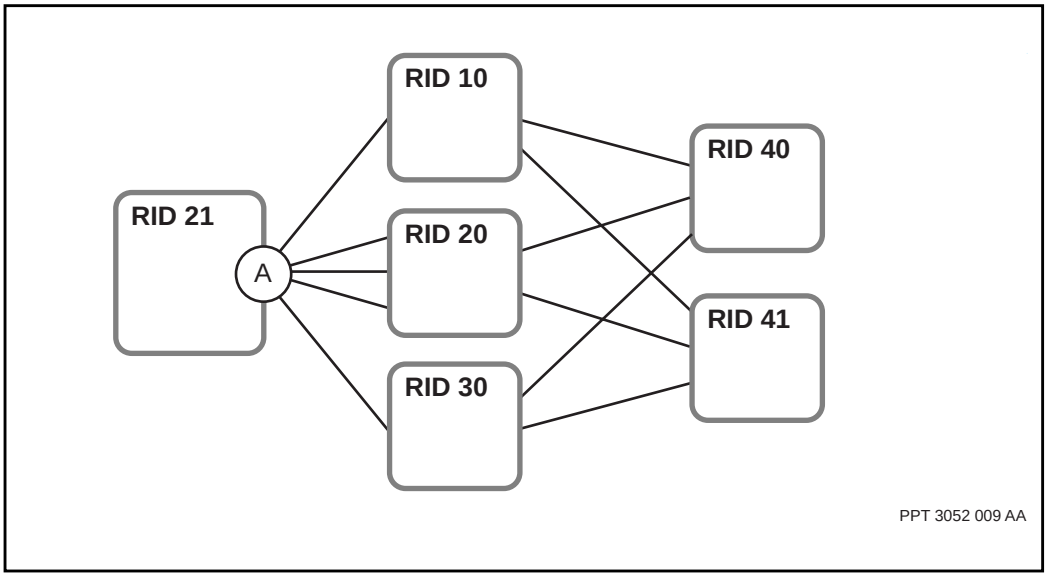

Gateway node to gateway node (rule 4)

In the figure “Gateway node to gateway node” (page 103), DPRS must select one of three equal paths from a gateway node through a neighboring RID to reach the destination RID.

When the best two paths are both through the same neighbor RID subnet then the path selection is based on neighbor MIDs. In this case, DPRS selects

- for delay RCOS, the two paths through the neighbors whose MIDs are lowest if the destination RID is even and highest if the destination RID is odd
- for throughput RCOS, the two paths through the neighbors whose MIDs are highest if the destination RID is even and lowest if the destination RID is odd

In the topology shown in the figure "Gateway node to gateway node" (page 103), all the links have equal metrics, so DPRS at node A picks nodes B and D, in that order, to reach RID 22 and 50 for the throughput RCOS. To reach RID 51, DPRS picks nodes C and D.

Figure 37
Gateway node to gateway node

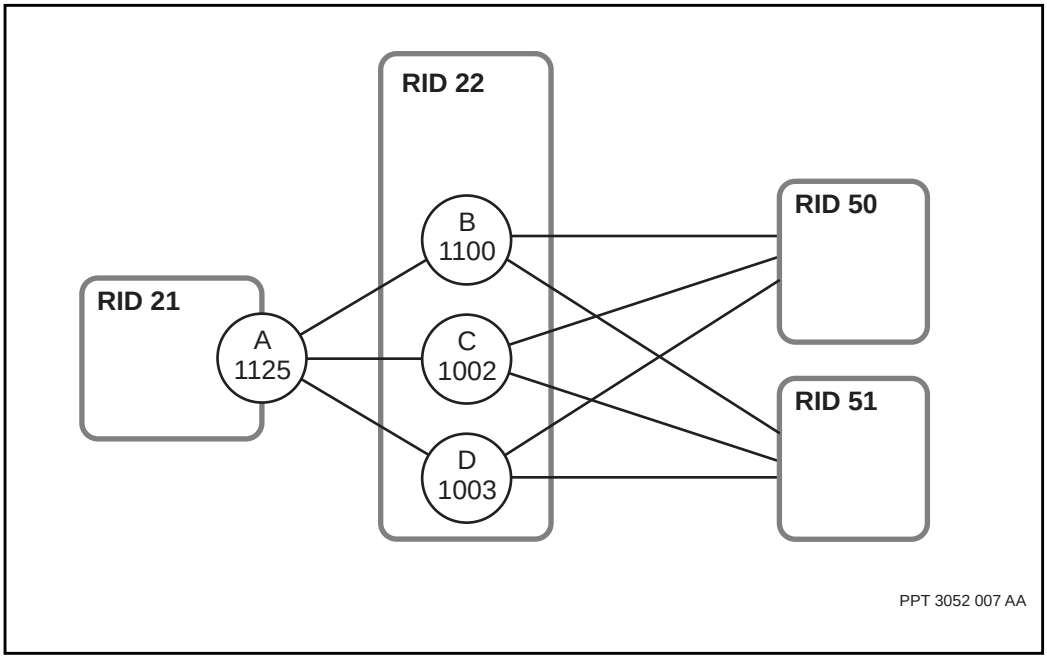

DPRS automatic route tester

The DPRS automatic route tester provides a built-in error detection capability that can constantly test and verify the Passport DPRS routing system. The route tester can be activated for one or both of delay and throughput COS (Class Of Service) on a per node basis to reduce the manual DPRS route verification workload of network operators and discover potential, although rare, DPRS routing system errors in a live network.

The automatic route tester operates by periodically sending DPRS ping packets to all RID and MID destinations as recorded in the DPN routing table. These packets are monitored to determine whether they have returned within 15 seconds, and analyzed for DPRS route information.

Any DPRS route is tested three consecutive times before being identified as having a problem. The first test uses RTD (round-trip-delay) Ping packets while subsequent tests use Path Ping packets. If no problems are found in the first or second test, any subsequent tests are skipped.

An alarm is raised by a particular node when the first DPRS routing problem is detected on that node. No subsequent alarms will be raised by that node even if additional problems are detected. A DPRS routing problem is considered to no longer exist only when a re-test indicates the problem is gone or when the RID/MID destination identified by this problem becomes unreachable. The alarm stays outstanding until all DPRS routing problems on the node no longer exist, in which case the alarm is cleared by the node that originally raised it. The same alarm will be raised if a future problem is detected.

The route testing interval sets the frequency (in seconds) at which DPRS route testing will be conducted for successive RID/MID destinations, and at which the re-testing is carried out for all already-detected DPRS routing problems The default value of 60 seconds is recommended.

The DPRS automatic route tester is disabled by default. To enable the route tester, see “Enabling the DPRS automatic route tester” (page 144).

Alarms

DPRS routing alarms help in resolving provisioning errors and indicate routing problems. The alarms include

- transmission error
- shared memory shortage
- temporary looping
- duplicate node IDs or MIDs in the RID subnet
- different RID values in the RID subnet
- incorrect metric cutoff provisioning
- missing updates from remote node
- software error in routing table maintenance

You can detect other provisioning errors by viewing the component and attribute values as they were provisioned.

For detailed information on alarms, see 241-5701-500 *Passport 6400, 7400, 15000, 20000 Alarms*.

Chapter 5
DPRS configuration

Configure DPRS to provide an efficient, connectionless routing system for delay-sensitive and high-throughbput traffic.

Prerequisites

- appropriate software release installed according to the procedures in the 241-5701-270 *Passport 7400, 15000, 20000 Software Installation Guide*.
- For general provisioning information and basic provisioning procedures, see 241-5701-600 *Passport 7400, 15000, 20000 Configuration Guide*.
- For a detailed description of the operator commands used in this chapter, see 241-5701-050 *Passport 7400, 15000, 20000 Commands*.
- For detailed information on components and attributes, see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

DPRS configuration task flow

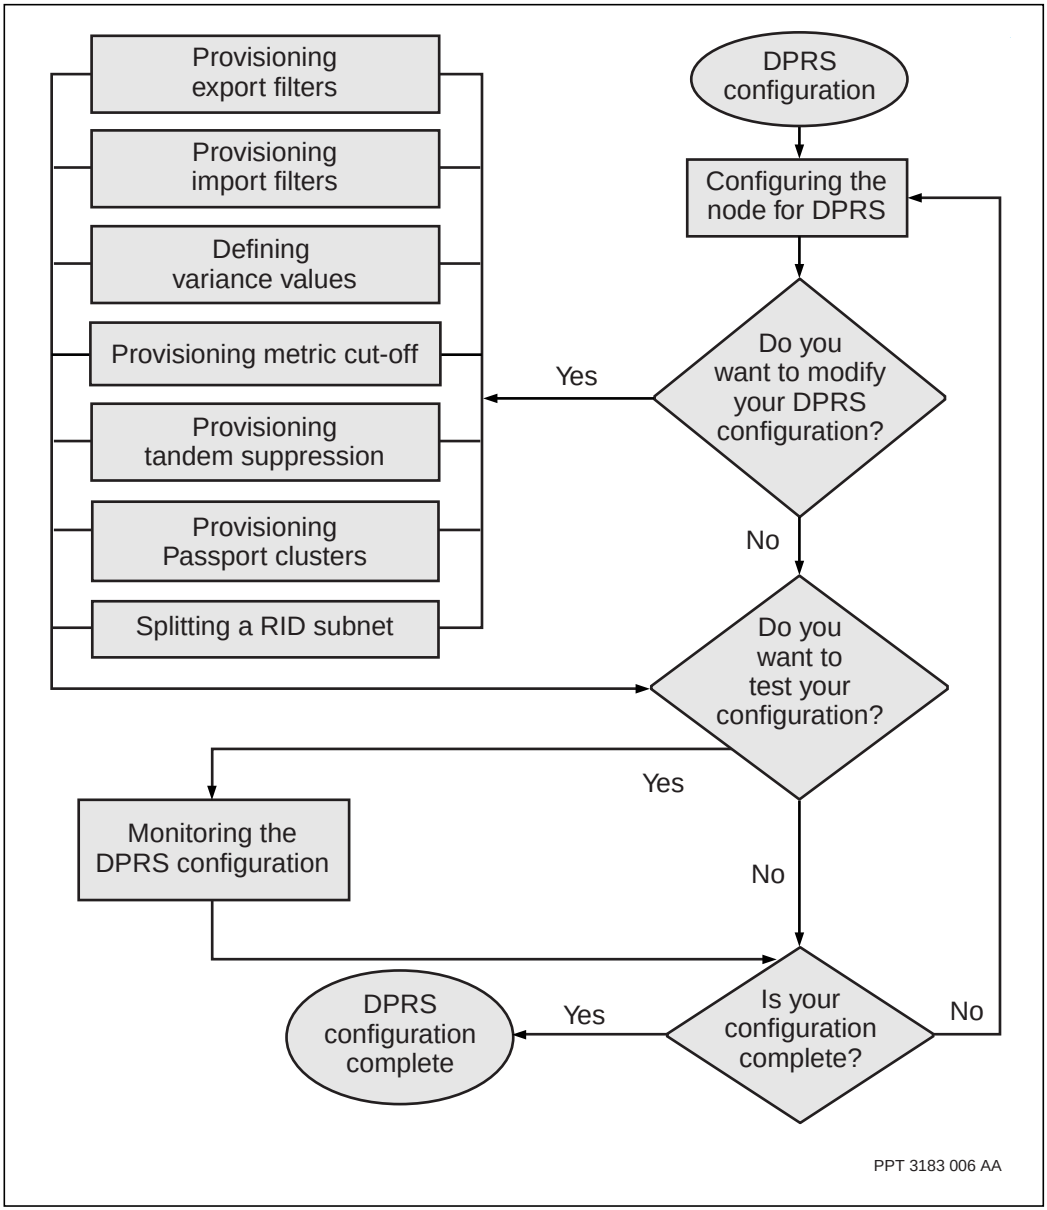

Configuring the node for DPRS

Provision DPRS on a Passport node to

- define the routing identifier (RID) for the node
- define the module identifier (MID) for the node
- define the traffic balancing algorithm

Prerequisites

- appropriate software release installed according to the procedures in the 241-5701-270 *Passport 7400, 15000, 20000 Software Installation Guide*.
- See “DPRS routing” (page 27) for more information on DPRS.
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1 Set the RID value for the node.

```
set rtg dpn routingid <RID_subnet>
```

2 Set the MID value for the node.

```
set rtg dpn moduleid <module_id>
```

3 Define the traffic balancing method for the node.

```
set rtg dpn forwardingPolicy <type>
```

Variable definitions

Variable	Value
<RID_subnet>	1 to 126. The RID value identifies the subnet containing this node and must be the same for all nodes that form one RID subnet, and unique amongst all RID subnets in the network.
<module_id>	1 to 1909. The MID value identifies the node and must be unique in the RID subnet.
<type>	loadshare, loadspread, or loadspreadfast (loadspread is the default)

Procedure job aid

Figure 38
Configuring the node for DPRS component hierarchy

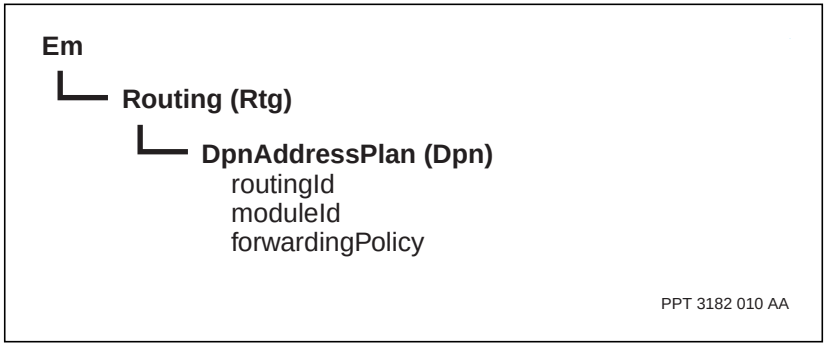

Provisioning import filters

Provision import filters to improve network security and efficiency by limiting the sharing of RID routing information between subnets.

Prerequisites

- You must complete the procedure "Configuring the node for DPRS" (page 110).
- See "RID filters" (page 84) for more information on import filters.
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1. Add a *ridFilter* component for the neighbor RID you want to filter.

   ```
   add rtg dpn ridFilter/<n>
   ```

   Note: All RIDs appear in the import list for the new component.

2. Specify the RIDs for which you want to receive RTUs.

   ```
   set rtg dpn ridFilter/<n> importRidList ! <rid_list>
   ```

Variable definitions

Variables	Values
<n>	neighboring RID from which to filter, in the range 1 to 126
<rid_list>	the list of RIDs that can be imported from this neighboring RID

Procedure job aid

Figure 39
Provisioning import filters component hierarchy

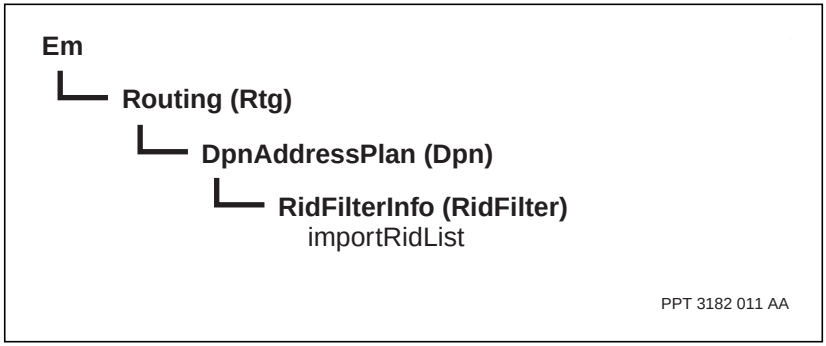

Defining variance values

Define variance values to improve traffic spreading across the Passport backbone.

Prerequisites

- You must complete the procedure "Configuring the node for DPRS" (page 110).
- See "Variance" (page 61) for more information about variance.
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1 Set the *delayMetricRangeBoundary* attribute.

```
set rtg dpn delayR <delay_value>
```

2 Set the *tputMetricRangeBoundary* attribute.

```
set rtg dpn tputR <tput_value>
```

3 Set the low *variance* attributes.

```
set rtg dpn variance co <cos> metricRange low
<lowrange_value>
```

4 Set the high *variance* attributes.

```
set rtg dpn variance co <cos> metricRange high
<highrange_value>
```

Variable definitions

Variable	Value
<delay_value>	a decimal number between 0 (the default) and 214783647, inclusive
<tput_value>	a decimal number between 0 (the default) and 214783647, inclusive
<cos>	delay or throughput

(Sheet 1 of 2)

Variable	Value
<lowrange_value>	a decimal number between 0 (the default) and 500, inclusive, or 9999
<highrange_value>	a decimal number between 0 (the default) and 500, inclusive, or 9999

(Sheet 2 of 2)

Procedure job aid

Figure 40
Defining variance values component hierarchy

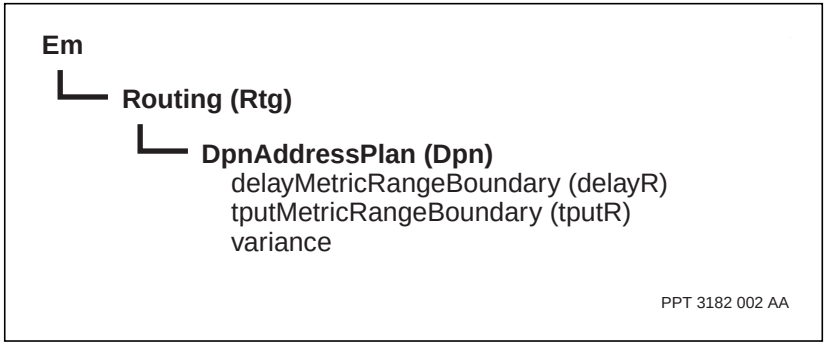

Provisioning metric cut-off

Provision metric cutoff to specify the metric value above which a RID is unreachable.

Prerequisites

- You must complete the procedure “Configuring the node for DPRS” (page 110).
- See“Topology metrics” (page 32) for more information on metrics.
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1 Set the value of the *delayMetricCutOff* attribute.

```
set rtg dpn delayMetricCutOff <delay_value>
```

2 Set the value of *throughputMetricCutOff* attribute.

```
set rtg dpn throughputMetricCutOff <thput_value>
```

Variable definitions

Variables	Values
<delay_value>	Decimal number between 64 and 250.
<thput_value>	Decimal number between 128 and 245.

Procedure job aid

Figure 41
Provisioning metric cut-off component hierarchy

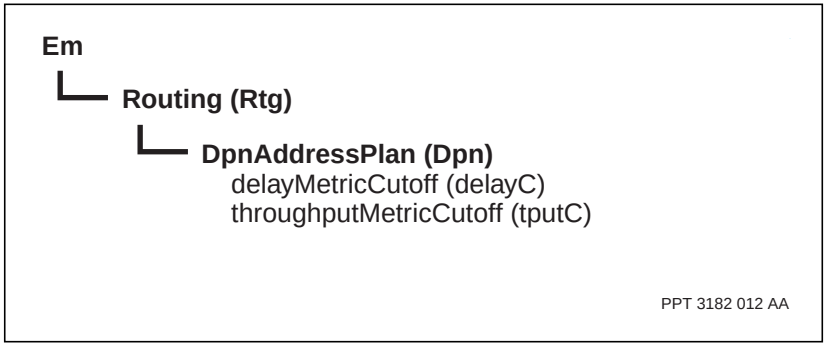

Provisioning tandem suppression

Provision tandem suppression to prevent tandem traffic from traversing the node and affecting DPRS and PORS traffic.

Prerequisites

- You must complete the procedure "Configuring the node for DPRS" (page 110).
- See "Tandem suppression" (page 71) for more information on tandem suppression.
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1 Set the *tandemTraffic* attribute to prevent tandem traffic or to allow tandem traffic.

```
set rtg tandemTraffic <permission>
```

Variable definitions

Variable	Value
<permission>	allowed (default) or denied

Procedure job aid

Figure 42
Provisioning tandem suppression component hierarchy

Provisioning export filters

Provision export filters to improve network security and efficiency by limiting the sharing of RID routing information between subnets.

Prerequisites

- You must complete the procedure "Configuring the node for DPRS" (page 110).
- See the section "RID filters" (page 84) for more information on export filters.
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1 Add a *ridFilter* component for the neighbor RID you want to filter.

```
add rtg dpn ridFilter/<n>
```

Note: All RIDs appear in the export list for the new component.

2 Specify the RIDs for which you want to send RTUs.

```
set rtg dpn ridFilter/<n> exportRidList <rid_list>
```

Variable definitions

Variables	Values
<n>	neighboring RID from which to filter, in the range 1 to 126
<rid_list>	the list of RIDs to export to the neighboring RID

Procedure job aid

Figure 43
Provisioning export filters component hierarchy

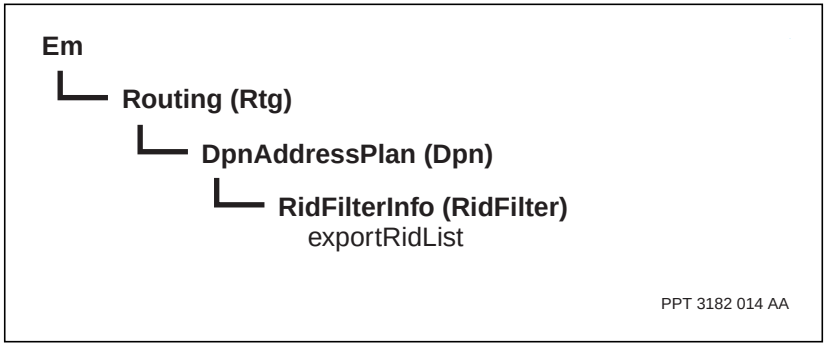

Provisioning Passport clusters

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

Prerequisites

- You must complete the procedure "Configuring the node for DPRS" (page 110).
- See the section "Passport clusters" (page 77) for more information on Passport clusters.
- Ensure both the planned cluster nodes and the planned backbone border nodes are running software that supports cluster functionality.
- For more information on DPRS components, see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

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

```
set rtg clusterNode yes
```

Note: The provisioning change is non-critical and takes effect immediately upon activation.

Procedure job aid

Figure 44
Provisioning Passport clusters component hierarchy

Splitting a RID subnet

Split a RID subnet to break a large subnet into smaller more manageable subnets.

Prerequisites

- You must complete the procedure “Configuring the node for DPRS” (page 110),
- See “RID subnet splits” (page 92) for more information on splitting a RID subnet.
- See 241-5701-410 *Passport 7400, 15000, 20000 Call Redirection Server Guide* for information about provisioning the CRS.
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1. Start with one Passport node at the edge of the original RID subnet.
2. Create a Crs component for the first node to be moved.

   `add Crs/<crs_id>`
3. Link the *Crs* component to the first node to be moved to the new RID.

   `add Crs/<crs_id> PAddr/<p_address>`
4. Create a Crs component for a node in the original RID.

   `add Crs/<crs_id>`
5. Link the Crs component to the node that is part of the original RID.

   `add Crs/<crs_id> PAddr/<p_address>`
6. Add an alternateRid subcomponent for the node in the original RID.

   `add Crs/<crs_id> AltRid`
7. Map the *AlternateRid* subcomponent of the original RID to the new RID.

   `set Crs/<crs_id> AltRid rid <subnet_id>`
8. Add an alternateRid subcomponent for the node in the new RID.

   `add Crs/<crs_id> AltRid`
9. Map the *AlternateRid* subcomponent of the new RID to the original RID.

```
set Crs/<crs_id> AltRid rid <subnet_id>
```

10 Provision the designated node to be in the new RID, and activate the node.

11 Gradually move more nodes from the original RID into the new RID. You can move more than one node at a time.

12 After all the nodes have been moved, update the call router databases in the two RID subnets to reflect the RID locations.

13 Verify that the call router databases are accurate.

14 Remove the *alternateRid* component from the CRS database.

```
del Crs/<crs_id> AltRid
```

15 Repeat these steps to create additional RID subnets.

Variable definitions

Variable	Values
<crs_id>	The identifier assigned for the Crs component.
<p_address>	The primary address of the node that the Crs is being used for.
<subnet_id>	The RID subnet number.

Procedure job aid

Figure 45
Splitting a RID subnet component hierarchy

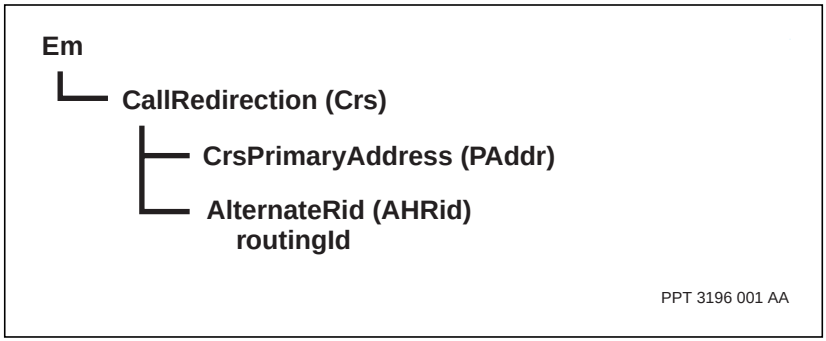

Chapter 6
DPRS monitoring

Monitor DPRS to view your network settings and identify problems with configuration.

Prerequisites

- “DPRS configuration” (page 107)
- For a detailed description of the operator commands used in this chapter, see 241-5701-050 *Passport 7400, 15000, 20000 Commands*.
- For detailed information on components and attributes, see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

DPRS troubleshooting task flow

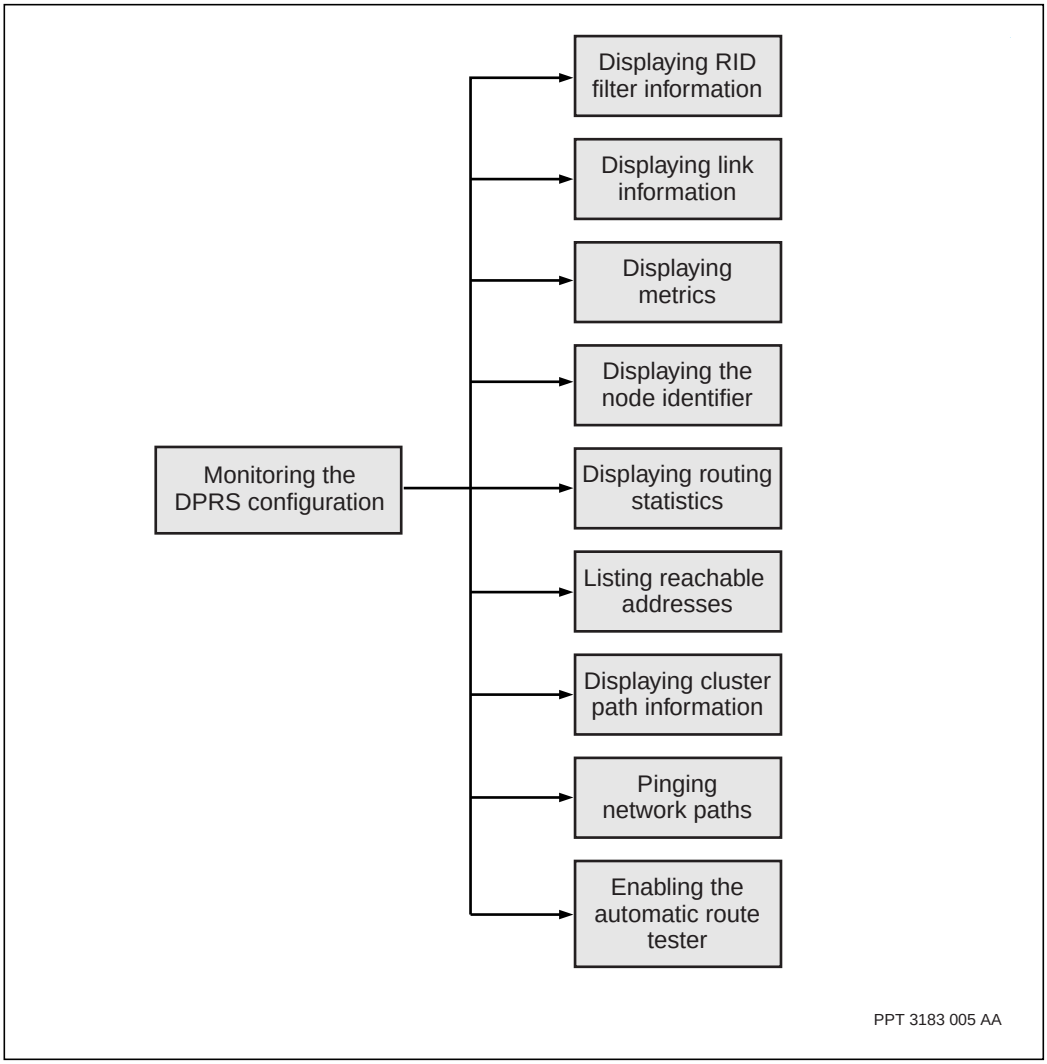

- “Displaying RID filter information” (page 138)
- “Displaying link information” (page 129)
- “Displaying metrics” (page 134)

Displaying the node identifier

Display the node identifier to check the node and region identifiers that base routing uses on this node.

Prerequisites

- You must have completed the procedure “DPRS configuration” (page 107)
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1 Display the node identifier and the region identifier for the current Passport node.

```
d -p mod
```

Procedure job aid

Figure 46
Displaying the node identifier component hierarchy

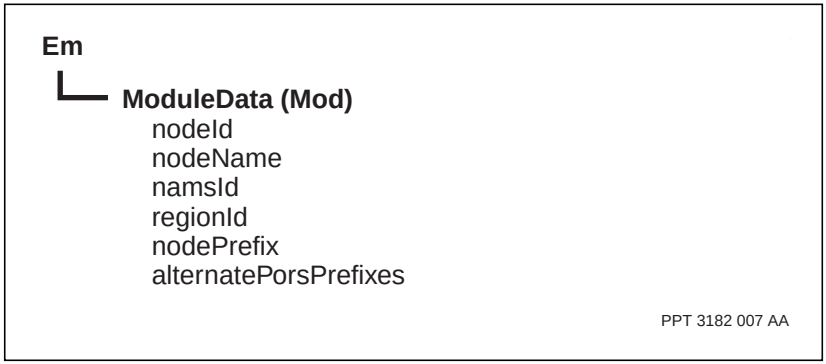

Displaying link information

Display link information to verify and troubleshoot your DPRS link configurations.

Prerequisites

- You must have completed the procedure "DPRS configuration" (page 107)
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Note: TRM link statistics are not updated for LoadSpreadFast.

Procedure steps

1. List all the links to neighboring nodes.

   ```
   l trm *
   ```

2. Display the links in a link group.

   ```
   l trm lg/<lg_name> lk/*
   ```

3. Display information about a link in a link group.

   ```
   d trm lg/<lg_name> lk/<link_number>
   ```

4. Display information about a logical network number. (The LNN is always 1 for DPRS.)

   ```
   d trm lg/<lg_name> LNN/1
   ```

Variable values

Variable	Value
<lg_name>	The name assigned to the link group.
<link_number>	The assigned link number.

Procedure job aid

Figure 47
Displaying link information component hierarchy

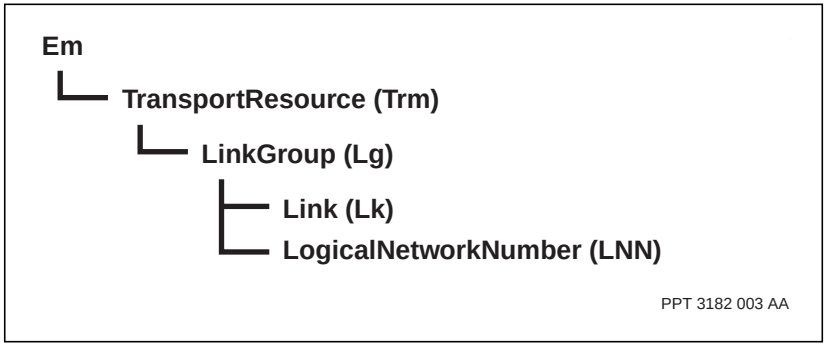

Listing reachable nodes

List reachable nodes to verify and troubleshoot the topology information of the subnet.

Prerequisites

- You must have completed the procedure “DPRS configuration” (page 107)
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1. Obtain a view of the nodes you can access (in all RID subnets in the topology region).

   ```
   l rtg top node/*
   ```

2. List the various link groups on a Passport node that connect to Passport node neighbors.

   ```
   l rtg top node/* lg/*
   ```

3. Display the metrics for a link group.

   ```
   d rtg top node/<node_name> lg/<lg_name>
   ```

Variable definitions

Variable	Value
<lg_name>	The name assigned to the link group.
<node_name>	The name of a Passport node.

Procedure job aid

Figure 48
Listing reachable nodes component hierarchy

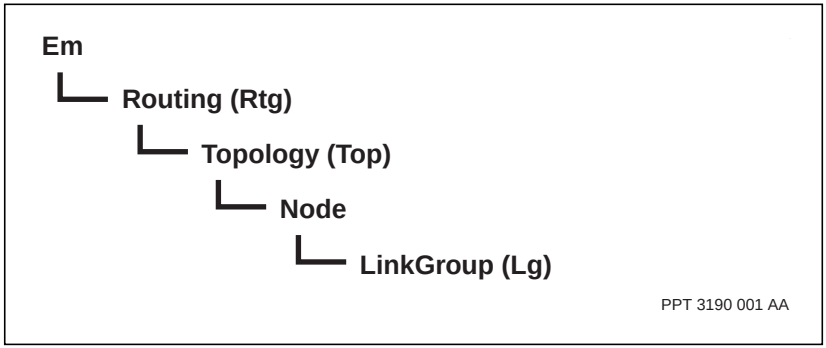

Listing reachable addresses

List reachable addresses to display the Passport destination call router (dcr), source call router (scr), and call redirection server (crd).

Prerequisites

- You must have completed the procedure "DPRS configuration" (page 107)
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1 List all reachable nodes in the subnet, as identified by their MIDs.

```
l rtg dpn mid/*
```

2 List all reachable Passport subnets, as identified by their RIDs.

```
l rtg dpn rid/*
```

3 List all reachable call servers.

```
l rtg dpn cs/*
```

Procedure job aid

Figure 49
Listing reachable addresses component hierarchy

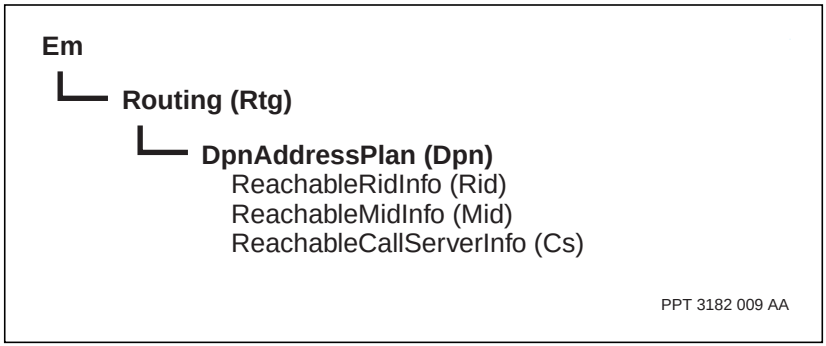

Displaying metrics

Display metrics to:

- view the DPN metrics that DPRS advertises to other RID subnets.
- view the Passport metrics DPRS uses in calculating best paths.

Prerequisites

- You must have completed the procedure “DPRS configuration” (page 107)
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1. Display the delay and throughput metrics and the next-hop link groups to reach a RID.

   ```
   d rtg dpn rid/*
   ```

2. Display delay and throughput metrics for a MID.

   ```
   d rtg dpn Mid/<mid_number>
   ```

Variable definitions

Variable	Value
<mid_number>	The identifier of a specific Mid.

Procedure job aid

Figure 50
Displaying metrics component hierarchy

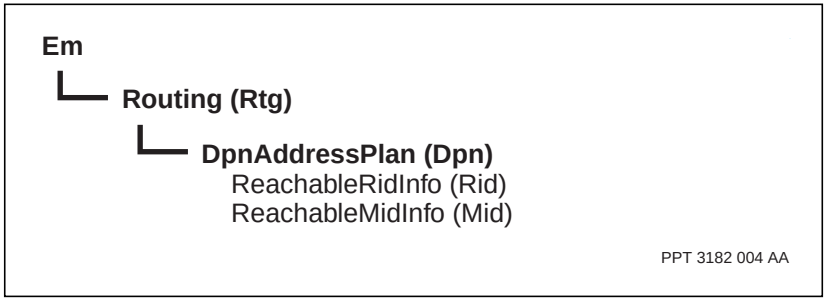

Displaying routing control statistics

Display the routing control statistics to

- list all the Passport MIDs in the subnet.
- display counts of the DPRS routing control traffic.
- display DPRS forwarding statistics that include the count of out-of-sequence packets, a total packet count, and the total divided into RCOS type.
- display a count of packets discarded because DPRS packet forwarding could not find a route for them.

Prerequisites

- You must have completed the procedure “DPRS configuration” (page 107)
- The forwardingPolicy must be set to loadSpread or loadShare for LP statistics to be displayed.
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1. Display control statistics on the routing of DPRS traffic on this node to all Passport nodes in the RID subnetwork.

   ```
   d rtg dpn
   ```

2. Display control statistics for a logical processor.

   ```
   d rtg dpn lpstats/<lp_number>
   ```

Variable definitions

Variable	Value
<lp_number>	The identifier of a specific logical processor.

Procedure job aid

Figure 51
Displaying routing control statistics component hierarchy

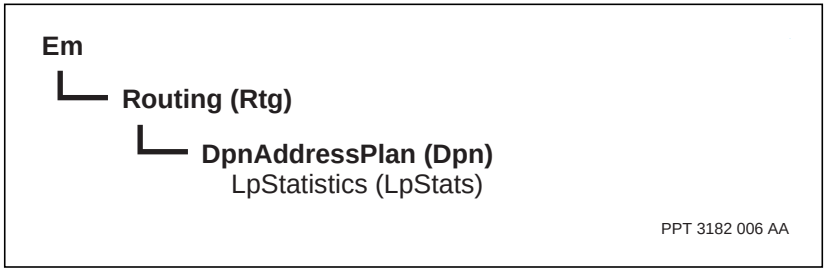

Displaying RID filter information

Display RID filter information to show the filter lists associated with a RID subnet.

Prerequisites

- You must have completed the procedure “DPRS configuration” (page 107)
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.
- You must have completed the procedure “Provisioning import filters” (page 112)
- You must have completed the procedure “Provisioning export filters” (page 119)

Procedure steps

1 Display the import and export lists for the RID filter associated with a RID subnet.

```
d -p rtg dpn ridfilter/<filter_num>
```

Variable definitions

Variable	Value
<filter_num>	The identifier of a specific RID filter.

Procedure job aid

Figure 52
Displaying RID filter information component hierarchy

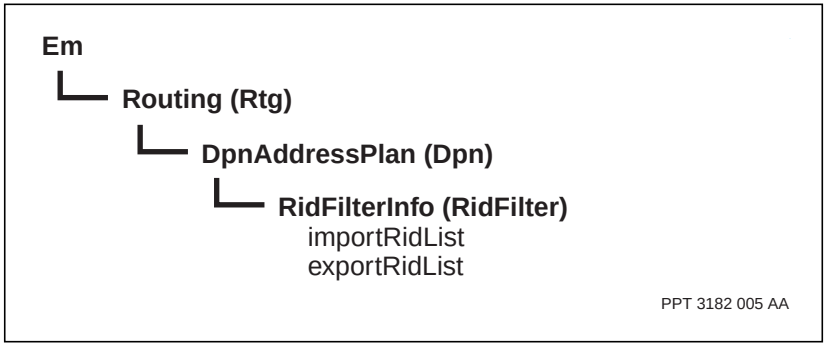

Displaying cluster path information

Display cluster path information to view path routing information for RIDs, MIDs, and call servers.

Prerequisites

- You must have completed the procedure "DPRS configuration" (page 107)
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure Steps

1 List all reachable MIDs in the entire RID subnet, including all Passport backbone and cluster MIDs, and AM clusters MIDs in the subnet.

```
list Rtg Dpn MID/*
```

2 Display information for the path being used to reach a RID.

```
d Rtg Dpn RID/<rid_number>
```

3 Display information for the path being used to reach a MID.

```
d Rtg Dpn MID/<mid_number>
```

4 Display information for the path being used to reach a call server (CS).

```
d Rtg Dpn CS/<callserver_num>
```

5 Display default route information including delay and throughput metrics, next hop link groups, and traffic proportions.

```
display Rtg Dpn defaultRoute
```

Variable definitions

Variable	Value
<rid_number	The identifier of a specific RID.
<mid_number>	The identifier of a specific MID.
<callserver_num>	The identifier of a specific call server.

Procedure job aid

Figure 53
Displaying cluster path information component hierarchy

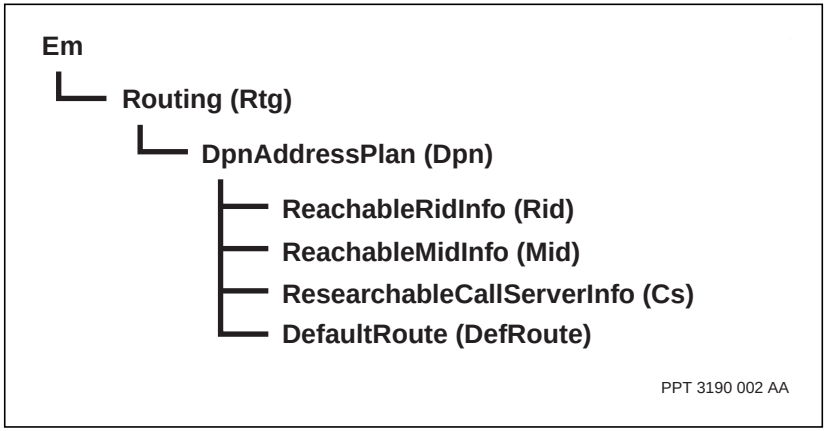

Pinging network paths

Ping network paths to

- determine the paths to the destination RID, MID, or both, based on the packet's RCOS (delay, throughput, or multimedia), reliability, or priority value
- measure the round trip delay in the network to any RID, MID, or both

Prerequisites

- You must have completed the procedure "DPRS configuration" (page 107)
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.
- For more information on the ping command, see 241-5701-050 *Passport 7400, 15000, 20000 Commands*

Procedure steps

1 Ping a MID to find the path with the highest throughput.

```
ping rtg dpn mid/<mid_id>
```

2 Determines how long it takes a packet to travel to a MID and back again (the round trip delay):

```
ping -rtd rtg dpn mid/<mid_id>
```

Variable definitions

Variable	Value
<mid_id>	The identifier of a specific MID.

Procedure job aid

Figure 54
Pinging network paths component hierarchy

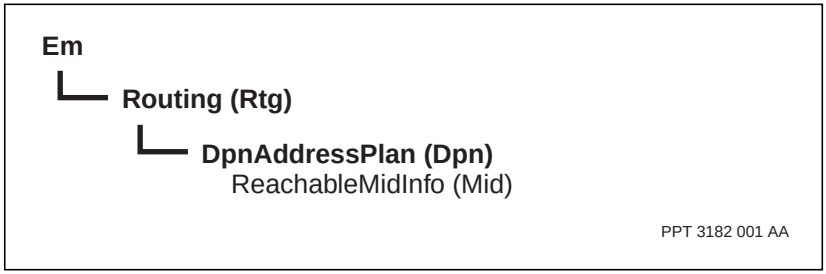

Enabling the DPRS automatic route tester

Enable the automatic route tester to constantly test and verify the Passport DPRS routing system.

Prerequisites

- You must have completed the procedure "DPRS configuration" (page 107)
- For more information on DPRS components see 241-5701-060 *Passport 7400, 15000, 20000 Components*.

Procedure steps

1. Enable the automatic route tester for DPRS delay-sensitive COS (Class of Service), throughput-optimized COS, or both.

   ```
   set rtg dpn art cosUnderTest <type_of_cos>
   ```

2. Adjust the route testing interval.

   ```
   set rtg dpn art testInterval <value>
   ```

   Note: Changing the default value of 60 seconds can significantly effect network performance. You should consult with a Nortel network engineer before altering any of the default values.

3. Turn the automatic route tester off.

   ```
   set rtg dpn art cosUnderTest none
   ```

Variable definitions

Variable	Values
<type_of_cos>	*delayCosOnly* for delay-sensitive COS, *tputCosOnly* for throughput-optimized COS, or *delayAndTputCos* for both
<value>	an interval value between 5 and 900 indicating the number seconds between successive tests.

Procedure job aid

Figure 55
Enabling the DPRS automatic route tester component hierarchy

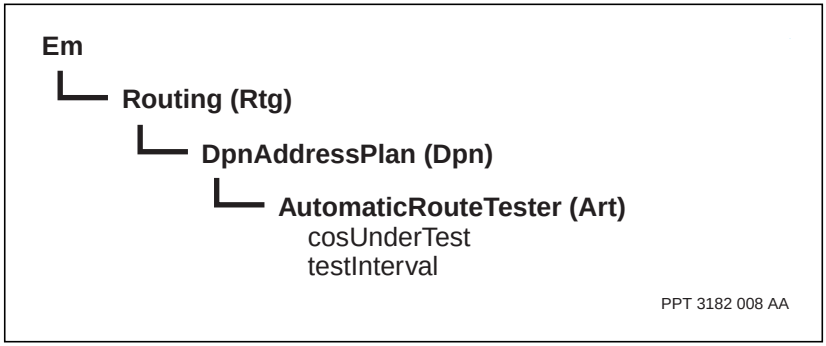

Passport 7400, 15000, 20000

Dynamic Packet Routing System Guide

Release 5.2

Publication: 241-5701-425
Document status: Standard
Document version: 5.2S1
Document date: November 2003
Printed in Canada

NORTEL NETWORKS