

Multiservice Switch 7400/15000/20000

Layer 3 Traffic Management Fundamentals

NN10600-590

Nortel Networks Multiservice Switch 7400/15000/20000

Layer 3 Traffic Management Fundamentals

Publication: NN10600-590

Document status: Standard

Document version: 6.1S2

Document date: November 2004

Copyright © 2004 Nortel Networks.

All Rights Reserved.

Printed in Canada

NORTEL, NORTEL NETWORKS, the globemark design, the NORTEL NETWORKS corporate logo, DPN, and PASSPORT are trademarks of Nortel Networks.

Publication history

November 2004

6.1S2 Standard

General availability. Contains information on Nortel Networks Multiservice Switch 7400, 15000, and 20000 for the PCR6.1 release.

Contents

About this document	17
What's new in this document	17
<hr/>	
Chapter 1	
Layer 3 traffic management overview	19
IP Differentiated Services code point (DCSP)	20
MPLS Experimental Bits (EXP)	20
IP classification and marking	20
IP firewalls	20
IP policing	21
<hr/>	
Chapter 2	
Differentiated Services on a Layer 3 node	23
Per-hop-behaviors	23
Traffic class	23
Scheduling class	24
Drop precedence	24
Connection class	25
Per-hop-behavior categories	26
Default effort forwarding (DF) Per-hop-behavior	27
Assured forwarding (AF) Per-hop-behaviors	27
Class selector forwarding (CS) Per-hop-behaviors	28
Expedited forwarding (EF) Per-hop-behavior	28
<hr/>	
Chapter 3	
Differentiated Services in a Layer 3 network	29
Differentiated Services domains	29

- DifferentiatedServiceDomain default configurations 30
 - DifferentiatedServiceDomain/Multiservice (Dsd/ms) 32
 - DifferentiatedServiceDomain/classSelector (Dsd/cs) 33
 - DifferentiatedServiceDomain/assuredForwarding (Dsd/af) 33
 - DifferentiatedServiceDomain/packetVoice (Dsd/pv) 34
 - DifferentiatedServiceDomain/wirelessUmts (Dsd/umts) 34
 - DifferentiatedServiceDomain/custom (Dsd/cu) 35
- Interfaces between Differentiated Services domains 35
 - Domain core interface 37
 - Domain edge interface 38
 - Domain boundary interface 39

Chapter 4

Layer 3 traffic management in RFC 2547 networks 41

- RFC2547 PE node: MPLS tunnel entry data path 44
- RFC 2547 PE node: MPLS tunnel exit data path 50
- RFC2547 PE node: IP tandem data path on the Vrf 55
- RFC2547 PE node: IP tandem data path on the Rtr 61
- RFC2547 P node: MPLS tandem data path 64
- RFC 2547 P node: IP tandem data path 67

Chapter 5

Layer 3 traffic management in RFC 2764 networks 69

- DiffServ based layer 3 traffic management in RFC 2764 networks 71
 - RFC 2764 PE node: IP tunnel entry data path 71
 - RFC 2764 PE node: IP tunnel exit data path, DiffServ based 78
 - RFC 2764 PE node: IP tandem data path on the cVR, DiffServ based 84
 - RFC 2764 PE node: IP tandem data path on the VCG, DiffServ based 84
- CoS based layer 3 traffic management in RFC 2764 networks 85
 - RFC 2764 PE node: IP tunnel entry data path, CoS based 85
 - RFC 2764 PE node: IP tunnel exit data path, CoS based 90
 - RFC 2764 PE node: IP tandem data path on the cVR, CoS based 95
 - RFC 2764 PE node: IP tandem data path on the VCG, CoS based 95

DiffServ/Cos hybrid based Layer 3 traffic management in RFC 2764 networks	96
RFC 2764 PE node: IP tunnel entry data path, DiffServ/CoS hybrid based	96
RFC 2764 PE node: IP tunnel exit data path, DiffServ/CoS hybrid based	101
RFC 2764 PE node: IP tandem data path on the cVR, DiffServ/CoS hybrid based	106
RFC 2764 PE node: IP tandem data path on the VCG, DiffServ/CoS hybrid based	106

Chapter 6

Layer 3 traffic management in VIPR networks 107

DiffServ based layer 3 traffic management in VIPR networks	109
VIPR node: IP tandem data path, DiffServ based	109
CoS based layer 3 traffic management in VIPR networks	116
VIPR node: IP tandem data path, CoS based	116

Chapter 7

Layer 3 traffic management in native MPLS networks 123

Native MPLS PE node: IP tunnel entry data path	126
Native MPLS PE node: MPLS tunnel exit data path	130
Native MPLS PE node: IP tandem data path on the Rtr	133

Chapter 8

Layer 3 traffic management for local packets 137

DiffServ based Layer 3 traffic management for IP packets generated at a local source	137
DiffServ based Layer 3 traffic management for IP packets terminated at a local destination	138
Cos based Layer 3 traffic management for IP packets generated at a local source	138
Cos based Layer 3 traffic management for IP packets terminated at a local destination	139

Chapter 9

Layer 2 to Layer 3 interactions 141

- Layer 3 interactions with Layer 2 overview 141
 - Layer 3 interactions with ATM 141
 - Interaction of ATM end-to-end loopback and connection class 144
 - Layer 3 interactions with Ethernet 145
 - Layer 3 interactions with Frame Relay (FR) 147
 - Layer 3 interactions with Point-to-Point Protocol (PPP) 149
-

Chapter 10

IP classification and marking 151

- Ingress IP packet classification and marking 151
 - Differentiated Services code point (DSCP) classification 151
 - Multi-field (MF) classification 152
 - Link classification 153
 - Packet marking 153
 - Egress IP packet classification and marking 154
 - Differentiated Services code point (DSCP) classification 154
 - Packet marking 154
-

Chapter 11

IP firewalls 155

- Overview of IP flow filters 155
 - IP flow filter definition 157
-

Chapter 12

IP policing 159

- Overview of IP policing services 159
 - Policing services benefits 161
 - Policing process 161
 - Traffic stream selection 161
 - Traffic metering 162
 - Policing actions 162
 - Packet marking 162
 - Traffic meter and policing algorithm 162
 - Traffic meter and policing algorithm parameters 162
-

Leaky bucket mechanism	163
Egress policer drop precedence	165
Policing statistics	166
Policer types	166
Dual rate policer	167
Single CIR policer with all excess packets dropped	168
Single CIR policer with all excess forwarded	169
Single EIR rate policer	170
Color blind and color aware policer	170
Combined traffic class policer	171
EF minimum rate guaranteed policer	171
Policer as a DiffServ statistics collector	172

Chapter 13

IP Class of Service (CoS)

177

Overview of IP CoS	178
Packet classification at ingress	180
Layer 2 classification	181
DSCP-based classification	182
Flow-based classification	184
IP CoS policies	185
Packet treatment at egress	186
Packet marking	186
Class-based packet forwarding	187
Frame relay DTE class-based forwarding	189
CoS to QoS mapping over multiple DLCIs	189
CoS to QoS mapping over a single DLCI	190
IP-optimized DLCI class-based forwarding	192
ATM MPE class-based forwarding	193
Gigabit Ethernet class-based forwarding	194
Point-to-point protocol class-based forwarding	198
IP CoS over virtual media	198

Chapter 14

IP CoS to IP DiffServ migration

199

Overview of CoS to DiffServ migration	199
---------------------------------------	-----

- IP DiffServ advantages 200
- Determining migration requirements 201
 - Determining the carrier DiffServ domain 201
 - Determining the IP CoS to DiffServ PHB mapping 202
 - Determining the carrier ATM VCC migration 202
- Migration stages 205
 - Carrier network migration 205
 - Individual customer VPN migration 206
 - Single customer router (cVR) migration 208
- Comparison between IP CoS and IP DiffServ provisioned attributes 208

List of figures

- Figure 1 Example: Layer 3 service differentiation on AtmMpe using connection class 26
- Figure 2 Diffserv domain interface relationships 36
- Figure 3 DSCP marking through a core interface 37
- Figure 4 DSCP marking through a domain edge interface 38
- Figure 5 DSCP marking through a domain boundary interface 40
- Figure 6 RFC 2547 network 42
- Figure 7 RFC2547 PE node: MPLS tunnel entry data path 45
- Figure 8 RFC2547 PE node: MPLS tunnel exit data path 51
- Figure 9 RFC2547 PE node: IP tandem data path on the Vrf 56
- Figure 10 RFC2547 PE node: IP tandem data path on the Rtr 61
- Figure 11 RFC2547 P node: MPLS tandem data path 64
- Figure 12 RFC 2764 network 70
- Figure 13 RFC 2764 PE node: IP tunnel entry data path, DiffServ based 72
- Figure 14 RFC 2764 PE node: IP tunnel exit data path, DiffServ based 79
- Figure 15 RFC 2764 PE node: IP tunnel entry data path, CoS based 86
- Figure 16 RFC 2764 PE node: IP tunnel exit data path, CoS based 91
- Figure 17 RFC 2764 PE node: IP tunnel entry data path, DiffServ/ CoS hybrid based 97
- Figure 18 RFC 2764 PE node: IP tunnel exit data path, DiffServ/ CoS hybrid based 102
- Figure 19 VIPR network 108
- Figure 20 VIPR node: IP tandem data path, DiffSev based 110
- Figure 21 VIPR node: IP tandem data path, CoS based 117
- Figure 22 Native MPLS network 124
- Figure 23 Native MPLS PE node: IP tunnel entry data path 126
- Figure 24 Native MPLS PE node: MPLS tunnel exit data path 130
- Figure 25 Native MPLS PE node: IP tandem data path on the Rtr 133
- Figure 26 Filtering based on source address 156
- Figure 27 Filtering based on destination address 157
- Figure 28 Ingress/Egress Traffic Processing 160
- Figure 29 Leaky bucket mechanism 164
- Figure 30 Dual rate policer 167

Figure 31	Single CIR policer with all excess dropped	168
Figure 32	Single CIR policer with all excess forwarded	169
Figure 33	Single EIR rate	170
Figure 34	EF minimum rate guaranteed	172
Figure 35	IP CoS assignment on transmission path	181
Figure 36	Packet classification at ingress	183
Figure 37	CoS to QoS mapping on multiple DLCIs	190
Figure 38	Emission priority mapping on a single DLCI	192
Figure 39	CoS to QoS mapping on multiple VCCs	194
Figure 40	Ingress on gigabit Ethernet; egress on PQC-based ATM	196
Figure 41	Ingress on PQC-based ATM; egress on gigabit Ethernet	197

List of tables

Table 1	Default PHB configuration by domain type	31
Table 2	RFC2547 PE node: MPLS tunnel entry data path	47
Table 3	RFC2547 PE node: MPLS tunnel exit data path	53
Table 4	RFC2547 PE Node: IP tandem data path on the Vrf	58
Table 5	RFC2547 PE node: IP tandem data path on the Rtr	63
Table 6	RFC2547 P node: MPLS tandem data path	65
Table 7	RFC 2764 PE node: IP tunnel entry data path, DiffServ based	74
Table 8	RFC 2764 PE node: IP tunnel exit data path, DiffServ based	81
Table 9	RFC 2764 PE node: IP tunnel entry data path, CoS based	88
Table 10	RFC 2764 PE node: IP tunnel exit data path, CoS based	93
Table 11	RFC 2764 PE node: IP tunnel entry data path, DiffServ/CoS hybrid based	99
Table 12	RFC 2764 PE node: IP tunnel exit data path, DiffServ/CoS hybrid based	104
Table 13	VIPR node: IP tandem data path, DiffSev based	112
Table 14	VIPR node: IP tandem data path, CoS based	119
Table 15	Native MPLS PE node: IP tunnel entry data path	128
Table 16	Native MPLS PE node: MPLS tunnel exit data path	132
Table 17	Native MPLS PE node: IP tandem data path on the Rtr	135
Table 18	Interactions with ATM at ingress	142
Table 19	Interactions with ATM at egress	143
Table 20	Interactions with Ethernet at ingress	145
Table 21	Interactions with Ethernet at egress	146
Table 22	Interactions with Frame Relay (FR) at ingress	147
Table 23	Interactions with Frame Relay (FR) at egress	148
Table 24	Interactions with Point-to-Point Protocol (PPP) at ingress	149
Table 25	Interactions with Point-to-Point Protocol (PPP) at egress	150
Table 26	Policer parameters	173
Table 27	IP CoS support on access media	179
Table 28	EgressCosTreatment emissionPriority to scheduling class mapping	188
Table 29	Example set of carrier network PHBs	202

Table 30	Example of CoS index mapping to a DiffServ PHB	202
Table 31	Traffic mapped to ATM SC, with a change to interactive data traffic preference	204
Table 32	IP CoS and IP DiffServ software activation attributes	209
Table 33	IP CoS and IP DiffServ provisioning components	209
Table 34	IP CoS and IP DiffServ layer 2 configuration attributes	209
Table 35	IP CoS and IP DiffServ layer 3 configuration attributes	210
Table 36	IP CoS and IP DiffServ Phb definition attributes	210
Table 37	IP CoS and IP DiffServ drop precedence selection attributes	210
Table 38	IP CoS and IP DiffServ local packet marking attributes	210
Table 39	IP CoS and IP DiffServ Dscp/Tos classification attributes	211
Table 40	IP CoS and IP DiffServ Dscp marking attributes	211
Table 41	IP CoS and IP DiffServ Flow/Mf classification attributes	211

About this document

This document provides fundamental conceptual information pertaining to the components that make up Layer 3 traffic management.

This guide is for anyone who performs the following tasks: planning, installing and provisioning, operating and maintaining IP and MPLS systems.

This guide assumes that you are familiar with the concepts of internetworking, particularly with the IP suite, MPLS suite, and their common uses.

What's new in this document

The following features were added to this document:

- “Differentiated services (DiffServ)” (page 18)
- “IP Diffserv support on 4-port gigabit Ethernet and 16-port OC-3/STM-1 POS and ATM FPs” (page 18)
- “IP policing” (page 18)

Other changes made to this document include the following:

- Updated section “ATM MPE class-based forwarding” (page 193) to explain an automated backup mechanism for an ATM MPE connection failure.
- Updated Table “IP CoS support on access media” (page 179) with accurate application-specific integrated circuit (ASIC) terminology.

Differentiated services (DiffServ)

This new Layer 3 Traffic Management Configuration Nortel Networks publication (NTP) includes DiffServ information from previous Nortel Networks Multiservice Switch 7400/15000/20000 releases and for PCR6.1. The following sections were added to this document:

- “Differentiated Services in a Layer 3 network” (page 29)
- “IP classification and marking” (page 151)

The following section was updated:

- “IP CoS to IP DiffServ migration” (page 199)

IP Diffserv support on 4-port gigabit Ethernet and 16-port OC-3/STM-1 POS and ATM FPs

The following sections were updated for this feature:

- “IP policing” (page 21)
- “Multi-field (MF) classification” (page 152)

IP policing

The following section was added to this document:

- “IP policing” (page 159)

Chapter 1

Layer 3 traffic management overview

Use this section to obtain an overview of the functional areas associated with Nortel Networks Multiservice Switch Layer 3 Traffic Management.

The Layer 3 traffic management architecture on the Nortel Networks Multiservice Switch allows the network service provider to configure IP and MPLS traffic management capabilities in an easy and effective way. Advanced layer 1 and layer 2 traffic management capabilities supported by the Multiservice Switch, such as MPS scheduling on the MSA and MSAS FPs, WFQ scheduling on the GE FPs, the ATM and Frame Relay bundles, are abstracted from the user to ensure the proper QoS treatment is applied to the traffic irrespective of the types of control and functional processors deployed.

Navigation

- “IP Differentiated Services code point (DCSP)” (page 20)
- “MPLS Experimental Bits (EXP)” (page 20)
- “IP classification and marking” (page 20)
- “IP firewalls” (page 20)
- “IP policing” (page 21)

For information on how to provision these components, refer to NN10600-591 *Nortel Networks Multiservice Switch 7400/15000/20000 Layer 3 Traffic Management Configuration*.

IP Differentiated Services code point (DSCP)

The six bit Differentiated Services code point (DSCP) field in the IP packet header is used to select a Per-hop behavior (PHB) for the packet when it is forwarded by an IP router (VirtualRouter, Router, VpnRouteForwarder). The assigned PHB determines how the packet is treated relative to others when it is forwarded by the node.

The possible DSCP values range from 0-63, but only twenty-one of these have standard definitions. The remaining values are for local or experimental use

For more information on the IP Differentiated Services feature, refer to “Differentiated Services in a Layer 3 network” (page 29) in this document.

MPLS Experimental Bits (EXP)

The three bit experimental (EXP) field in the MPLS packet header is used to select a Per-hop behavior (PHB) for the packet when it is forwarded by an MPLS label-switched router (LSR). The assigned PHB determines how the packet is treated relative to others when it is forwarded by the node.

The possible EXP values range form 0-7 and do not have standard definitions.

IP classification and marking

Ingress IP classification and marking allows the DSCP field of IP packets to be modified after it is received at an interface prior to PHB selection.

Egress IP classification and marking allows the DSCP field of IP packets to be modified after PHB selection before it is transmitted at an interface.

Ingress classification is based on some portion of the layer 2, layer 3, or layer 4 headers. Egress classification is based the DSCP field only.

IP firewalls

IP flow filters are critical components in the creation of a secure network environment. System network administrators can decide which IP packet flows will be permitted or denied entry to the network. Permit or deny actions are based on the most specific match to the IP packet flow’s source IP address, destination IP address, or both the source and destination IP addresses.

For more information on IP flow filters as applied to the Nortel Networks Multiservice Switch, refer to “IP firewalls” (page 155) in this document.

IP policing

In a shared network resources environment it is very important to provide network operators with the mechanism to limit rate of traffic that traverses the network. This service is provided by the Multiservice Switch IP Policing feature. Using this feature network provider can monitor traffic flow entering the interface and either discard the packet or increase drop precedence for packets violating service agreement.

For more information on the Policing feature, refer to “IP policing” (page 159) in this document.

Chapter 2

Differentiated Services on a Layer 3 node

Use this section to learn about Differentiated Services Per-hop-behaviors (PHBs).

Navigation

- “Per-hop-behaviors” (page 23)
- “Per-hop-behavior categories” (page 26)

Per-hop-behaviors

Per-hop-behaviors (PHBs) define the treatment of a packet. The attributes of each PHB determine the importance of the packets relative to each other. The attributes are:

- “Traffic class” (page 23)
- “Scheduling class” (page 24)
- “Drop precedence” (page 24)
- “Connection class” (page 25)

Traffic class

The *trafficClass* attribute of a per-hop-behavior lets you aggregate per-hop-behavior component instances into groups that specify the same scheduling attribute values. The *trafficClass* attribute must be equal to an existing instance of component *Dsd TrafficClass*.

For more information, see the description of attribute *Dsd Phb trafficClass* in NN10600-060 *Nortel Networks Multiservice Switch 7400/15000/20000 Component Reference*.

Scheduling class

The Scheduling class of a per-hop-behavior determines how a packet is scheduled at the egress interface relative to other packets. *Scheduling class* is provisioned as a number from 0 to 7 depending on the number of queues supported at an interface.

In general, a higher value means that a packet is less likely to be delayed. A lower value means a packet is more likely to be delayed. The actual scheduling of packets depends on the scheduling mechanism at the egress interface.

For more information, see the description of attributes *Dsd Phb schedulingClass8Queues* and *schedulingClass4Queues* in NN10600-060 *Nortel Networks Multiservice Switch 7400/15000/20000 Component Reference* and the section on traffic management on the 4-port gigabit Ethernet FP in NN10600-551 *Nortel Networks Multiservice Switch 7400/15000/20000 FP Configuration Reference*.

Drop precedence

The *dropPrecedence* attribute of a per-hop-behavior controls the loss sensitivity of packets relative to other packets in a queue. The drop precedence can be provisioned to *low*, *medium*, or *high*.

In general, packets with a drop precedence of *low* are less likely to be discarded when the queue is congested. Packets with a drop precedence of *high* have a lower probability of being forwarded and are most likely to be discarded. The actual loss sensitivity of packets depends on the discarding mechanism of the queue.

For more information, see the description of attribute *Dsd Phb dropPrecedence* in NN10600-060 *Nortel Networks Multiservice Switch 7400/15000/20000 Component Reference*.

Connection class

The connection class assigned to the packet is used to select one of up to four link layer connections at the interface where the packet is transmitted. Per-packet selection of egress connections is supported only if these conditions are all true:

- the egress media is ATM or Frame Relay
- more than one link layer connection is configured to the next hop
- each connection is configured with a different connection class value

Normally, the packet is transmitted on the connection whose connection class is equal to the connection class assigned to the packet. If that connection class does not exist, the connection having the next lower connection class is used. If that connection does not exist either, the connection class having the next higher connection class is used.

For *AtmMpe*, the connection class is specified by the *ipCos* attribute of the *AtmConnection* subcomponent. For *IpoDlci*, the connection class is specified by the *ipCos* attribute of the *FrConnection* subcomponent.

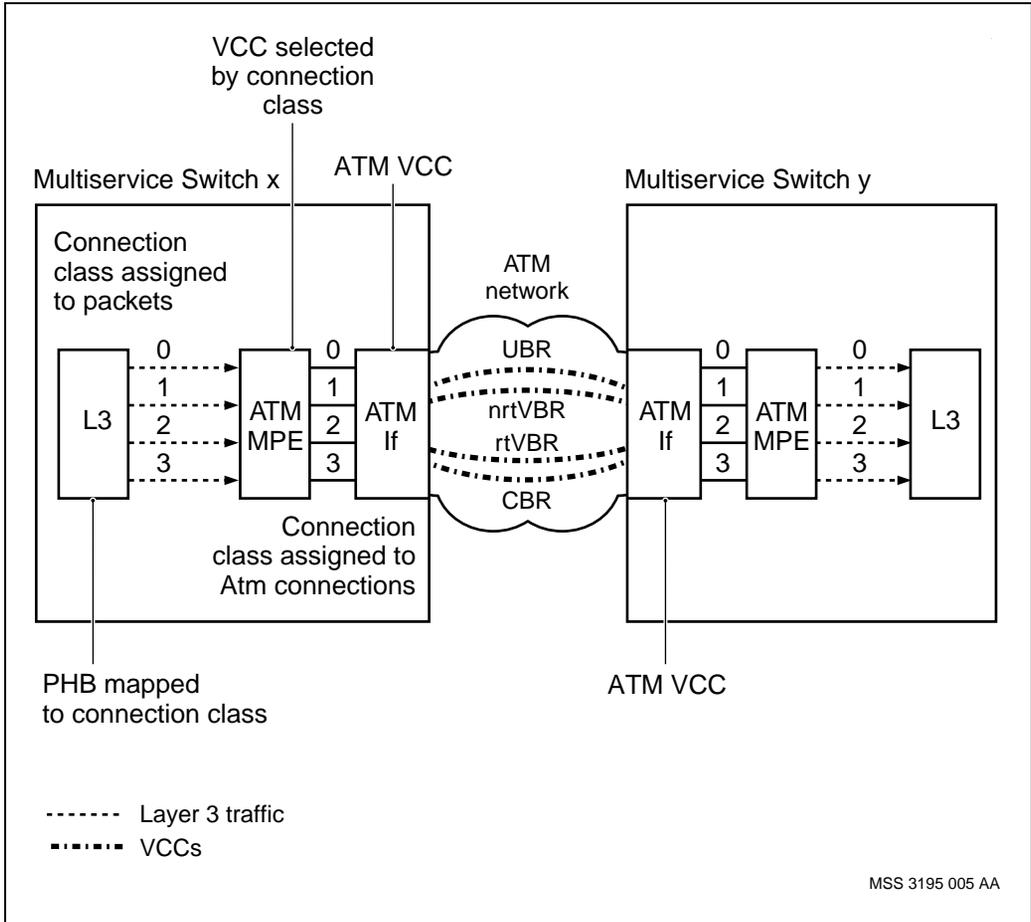
For *FrDte*, the connection class is specified by the *ipCos* attribute of the *StaticDlci* and *DynamicDlciDef* subcomponent. For *IpDlciGroup*, the connection class is specified by the *ipCos* attribute of the *FrConnection* subcomponent.

IP multicast and IP broadcast packets are not legible for multi link connection control. Each IP packet of this kind is copied and transmitted individually on each of the available layer 2 connection to the next IP hop.

If each link layer connection is configured to provide different policing and scheduling characteristics, the forwarding behaviors of IP packets are differentiated as they transit through the link layer network. In general, the scheduling characteristics of the connections should be configured to be better for numerically greater connection class values. For example, if four *AtmConnection* components are associated with *AtmIf* components configured for ATM service classes UBR, nrtVBR, rtVBR, and CBR, the *ipCos* attributes of the *AtmConnection* components should be either 0, 1, 2,

and 3, respectively, or 1, 1, 3, and 2, respectively. Refer to Figure 1, “Example: Layer 3 service differentiation on AtmMpe using connection class,” (page 26)

Figure 1
Example: Layer 3 service differentiation on AtmMpe using connection class



Per-hop-behavior categories

There are four general categories of per-hop-behaviors.

- “Default effort forwarding (DF) Per-hop-behavior” (page 27)

- “Assured forwarding (AF) Per-hop-behaviors” (page 27)
- “Class selector forwarding (CS) Per-hop-behaviors” (page 28)
- “Expedited forwarding (EF) Per-hop-behavior” (page 28)

Default effort forwarding (DF) Per-hop-behavior

Packets classified with default effort forwarding (DF) are given “best effort” quality-of-service treatment. Packets with this PHB are considered the least urgent and are the most likely packets to be discarded in a DiffServ enabled network.

Assured forwarding (AF) Per-hop-behaviors

RFC 2597 (“Assured Forwarding PHB Group”) describes this series of Per-hop-behaviors. Assured forwarding (AF) provides four bandwidth classes and three discard priorities. There are 12 assured forwarding PHBs that are grouped into four classes. Each class has three levels of loss sensitivity.

- af43
- af42
- af41
- af33
- af32
- af31
- af23
- af22
- af21
- af13
- af12
- af11

Class selector forwarding (CS) Per-hop-behaviors

RFC 2474 (“Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers”) describes this series of Per-hop-behaviors. There are 8 class selector (CS) PHBs. Each value is associated with a different level of urgency. The default forwarding DF PHB and the CS0 PHB are considered equivalent.

- cs7
- cs6
- cs5
- cs4
- cs3
- cs2
- cs1
- cs0

Expedited forwarding (EF) Per-hop-behavior

RFC 3246 (“An Expedited Forwarding PHB”) describes this Per-hop-behavior. Packets classified with expedited forwarding (EF) are considered delay sensitive and the least likely to be discarded. This PHB provides a very high quality of service treatment with special considerations for delay intolerance.

Chapter 3

Differentiated Services in a Layer 3 network

Use this section to learn about conceptual information pertaining to Differentiated Services in a Layer 3 networks.

Navigation

- “Differentiated Services domains” (page 29)
- “DifferentiatedServiceDomain default configurations” (page 30)
- “Interfaces between Differentiated Services domains” (page 35)

Differentiated Services domains

IP networks can be organized into DiffServ regions and DiffServ domains. A DiffServ domain is a group of routers (VirtualRouter, Router, or VpnRouteForwarder) that share the same per-hop-behavior definitions. A DiffServ region is a network of interconnected DiffServ domains.

Nortel Networks Multiservice Switch nodes provide six different pre-configured Differentiated Services domains. Each router (VirtualRouter, Router, or VpnRouteForwarder) is configured with one of the following domains:

- Multi Service domain (ms)
- ClassSelector domain (cs)
- AssuredForwarding domain (af)
- PacketVoice domain (pv)
- WirelessUmts domain (umts)

- Custom domain (cu)

Each differentiated service domain offers a subset of the PHBs available to the IP and MPLS packets being transported within the domain. Any domain may be customized by adding or deleting PHBs.

Your Nortel Networks technical support representative can help you determine the DiffServ domain that is the most suitable for your network.

DifferentiatedServiceDomain default configurations

A set of *PerHopBehavior* (*Phb*) and *TrafficClass* (*Tc*) subcomponent instances is automatically added when the *DifferentiatedServicesDomain* (*Dsd*) component is created. The selection of *Phb* and *Tc* subcomponent instances depends on the *Dsd* component instance.

The *Phb* subcomponent instances provide a user configured mapping from *Phb* to *Tc* and drop precedence. The *Tc* subcomponent instances provide a user configured mapping from *Tc* to scheduling classes for eight and four queues.

The *Phb/cs6* component instance is included in each default domain configuration because packets for network routing control are usually tagged for CS6.

The *Phb/df* or *Phb/cs0* component instance is included in each default domain configuration also. This component specifies the “best effort” or “default forwarding” treatment that is given to packets that are tagged DF or CS0, and to packets tagged with a Per-hop-behavior that is not supported in the domain.

In Nortel Networks Multiservice Switch systems, The EF and CS7 through CS1 *Phb* component instances are all configured to provide “low” drop precedence by default. The CS0 and DF *Phb* component instances are configured to provide “high” drop precedence by default. The AF1*, AF2*, AF3*, and AF4* series of *Phb* component instances are each configured to provide three levels of drop precedence by default.

Regardless which default domain configuration is used, the *Phb* and *Tc* subcomponents can be added, removed, and modified. Custom *Phb* components with instance values in the range 0 - 63 can be created also.

If an *Rtr* does not have a *Dsd* subcomponent, by default the router operates as if the *Dsd*/custom domain had been added. If a *Vrf* does not have a *Dsd* subcomponent, by default the *Vrf* operates using the *Dsd* configuration inherited from the router. If a *Vrf* and its parent router do not have *Dsd* subcomponents, by default the *Vrf* also operates as if the *Dsd*/custom domain had been added.

Table 1, “Default PHB configuration by domain type,” (page 31) lists the per-hop-behaviors supported for each DiffServ domain.

Each PHB in an IP DiffServ domain is given a default traffic class, drop precedence, and connection class value when it is created.

Any default PHB can be removed from a domain except for DSCP=0. The scheduling class, drop precedence, and connection class can also be modified for each PHB.

Table 1
Default PHB configuration by domain type

DSCP		Multiservice (ms)	Class Selector (cs)	Assured Forwarding (af)	Packet voice (pv)	WirelessU mts (umts)	Custom (cu)
decimal format	RFC format						
56	CS7	yes	yes		yes	yes	
48	CS6	yes	yes	yes	yes	yes	yes
46	EF	yes			yes	yes	
40	CS5	yes	yes		yes		
38	AF43	yes		yes			
36	AF42	yes		yes			
34	AF41	yes		yes			
32	CS4	yes	yes				
30	AF33	yes		yes		yes	
28	AF32	yes		yes		yes	
26	AF31	yes		yes		yes	
24	CS3	yes	yes				

(Sheet 1 of 2)

Table 1 (continued)
Default PHB configuration by domain type

DSCP		Multiservice (ms)	Class Selector (cs)	Assured Forwarding (af)	Packet voice (pv)	WirelessU mts (umts)	Custom (cu)
decimal format	RFC format						
22	AF23	yes		yes		yes	
20	AF22	yes		yes		yes	
18	AF21	yes		yes		yes	
16	CS2	yes	yes				
14	AF13	yes		yes	yes	yes	
12	AF12	yes		yes	yes	yes	
10	AF11	yes		yes	yes	yes	
8	CS1	yes	yes		yes	yes	
0	DF	yes	yes	yes	yes	yes	yes
0	CS0	yes	yes	yes	yes	yes	yes

(Sheet 2 of 2)

A brief summary of the default configuration for each domain type is presented below.

DifferentiatedServiceDomain/Multiservice (Dsd/ms)

- critical CS7
- network CS6
- premium CS5, EF
- platinum CS4, AF41, AF42, AF43
- gold CS3, AF31, AF32, AF33
- silver CS2, AF21, AF22, AF23
- bronze CS1, AF11, AF12, AF13
- standard CS0, DF

The default domain configuration is based on the Nortel Networks System Requirements Documents (SRDs) for end-to-end Quality of Service (QoS) across various Nortel Networks data networking platforms. When packets are transmitted at an interface that supports eight queues, a separate queue is used to process each one of the traffic classes. When packets are transmitted at an interface that supports four queues, the critical and network traffic classes are aggregated onto one queue; likewise, the platinum, gold, silver, and bronze traffic classes are aggregated onto one queue. The aggregation of AF1*, AF2*, AF3*, and AF4* onto one queue is not compliant with RFC 2597.

DifferentiatedServiceDomain/classSelector (Dsd/cs)

- critical CS7
- network CS6
- premium CS5
- platinum CS4
- gold CS3
- silver CS23
- bronze CS1
- standard CS0

This default domain configuration is the subset of the Dsd/pp configuration that provides just the eight Class Selectors (CS) *Phb* component instances. See “DifferentiatedServiceDomain/Multiservice (Dsd/ms)” (page 32) for more information and limitations.

DifferentiatedServiceDomain/assuredForwarding (Dsd/af)

- network CS6
- platinum AF41, AF42, AF43
- gold AF31, AF32, AF33
- silver AF21, AF22, AF23
- bronze AF11, AF12, AF13
- standard DF

The default domain configuration is optimized to provide all twelve AF Phbs at interfaces with eight or four queues. This is the only default domain configuration that is not a subset of the Dsd/pp default domain configuration. When packets are transmitted at an interface that supports eight queues, two of the queues are unused. When packets are transmitted at an interface that supports four queues, the network and platinum traffic classes are aggregated onto one queue; likewise, the bronze and standard traffic classes are aggregated onto one queue. Packets tagged with CS6 and AF31 are given equal treatment at an interface with four queues, and packets tagged with AF13 and DF are given equal treatment also.

DifferentiatedServiceDomain/packetVoice (Dsd/pv)

- critical CS7
- network CS6
- premium CS5, EF
- bronze CS1, AF11, AF12, AF13
- standard DF

This default domain configuration is the subset of the Dsd/pp configuration that provides just the nine *Phb* component instances that are used by Media Gateway applications on Nortel Networks Multiservice Switch systems. The aggregation of AF PHBs onto one queue at interfaces that support four queues is avoided by providing just the AF1* Phb subcomponent instances. See “DifferentiatedServiceDomain/Multiservice (Dsd/ms)” (page 32) for more information and limitations.

DifferentiatedServiceDomain/wirelessUmts (Dsd/umts)

- critical CS7
- network CS6
- premium EF
- gold AF31, AF32, AF33
- silver AF21, AF22, AF23
- bronze CS1, AF11, AF12, AF13
- standard DF

The default domain configuration is the subset of the Dsd/pp configuration that provides just the fourteen *Phb* component instances that are used by Universal Mobile Telecommunications System (UMTS) in Nortel Networks Multiservice Switch systems. See “DifferentiatedServiceDomain/Multiservice (Dsd/ms)” (page 32) for more information and limitations.

DifferentiatedServiceDomain/custom (Dsd/cu)

- network CS6
- standard DF

This default domain configuration is a subset of the Dsd/pp configuration that provides minimum Differentiated Services processing. Like any of the other default domain configurations, it can be customized by adding, removing, and modifying *Phb* and *Tc* subcomponents. If the Phb/CS6 subcomponent instance is deleted from the Dsd/cu configuration, all packets are given default forwarding treatment on one queue. See “DifferentiatedServiceDomain/Multiservice (Dsd/ms)” (page 32) for more information and limitations.

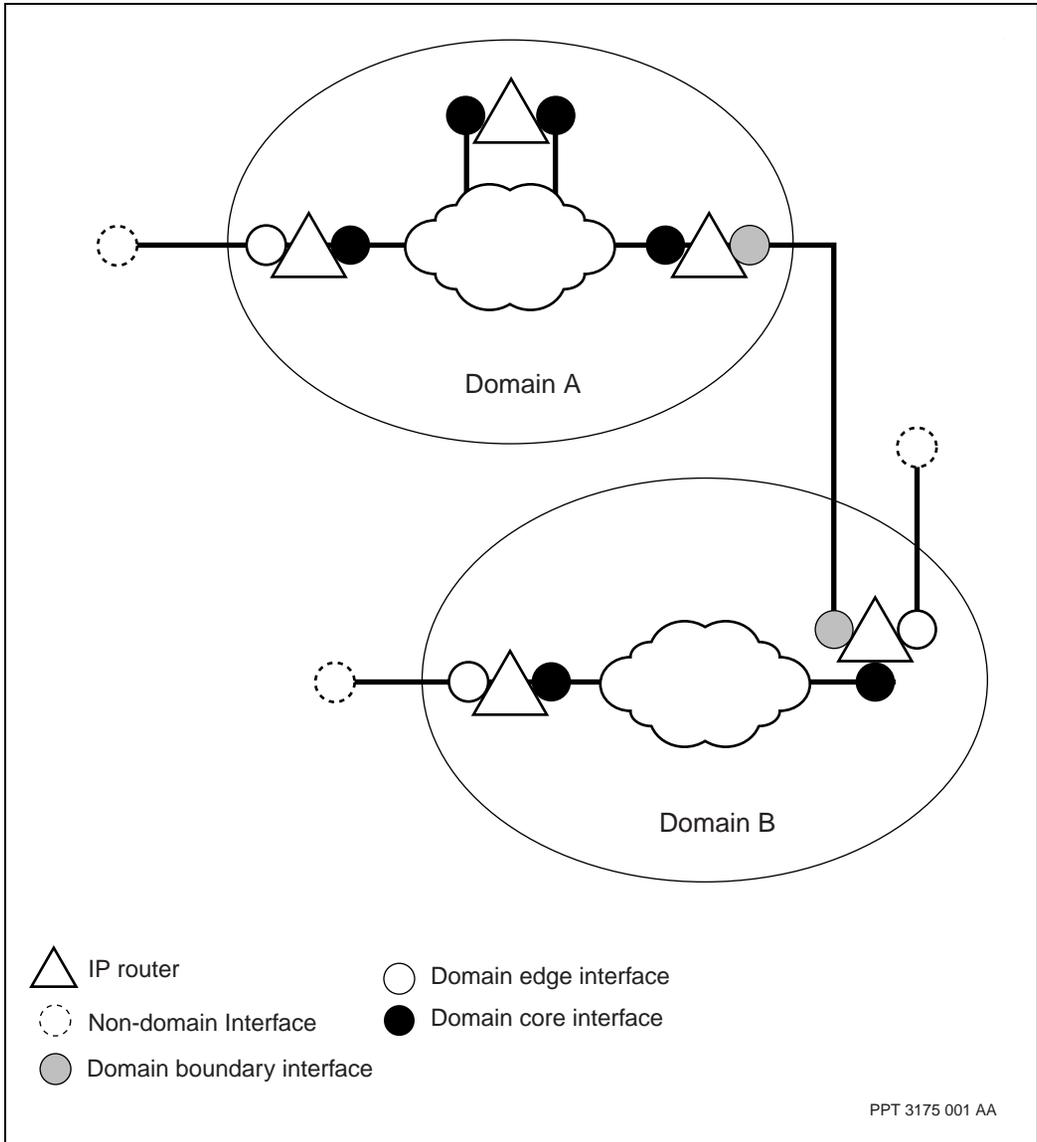
Interfaces between Differentiated Services domains

A large network can consist of many different interconnected Differentiated Services domains, and equipment within a DiffServ domain can be connected to equipment that does not support Differentiated Services (not in any DiffServ domain). This implies three kinds of DiffServ domain interface types.

- “Domain core interface” (page 37)
- “Domain edge interface” (page 38)
- “Domain boundary interface” (page 39)

Each interface has a different method of classifying and marking the DSCP field of the IP packets that are received and transmitted by the interface. Figure 2, “Diffserv domain interface relationships,” (page 36) shows where each interface type appears in a network.

Figure 2
Diffserv domain interface relationships

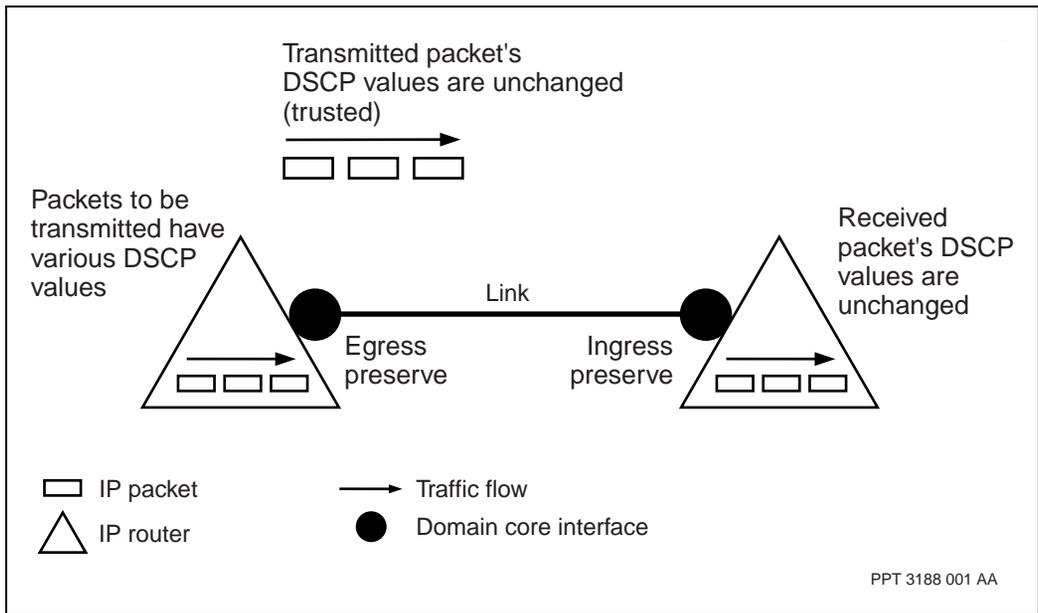


Domain core interface

A domain core interface transmits and receives packets from an IP router that is also in the same differentiated services domain. The DSCP value of a packet is preserved as it is received or transmitted through a core interface.

Figure 3, “DSCP marking through a core interface,” (page 37) shows how packets are treated at a core interface.

Figure 3
DSCP marking through a core interface

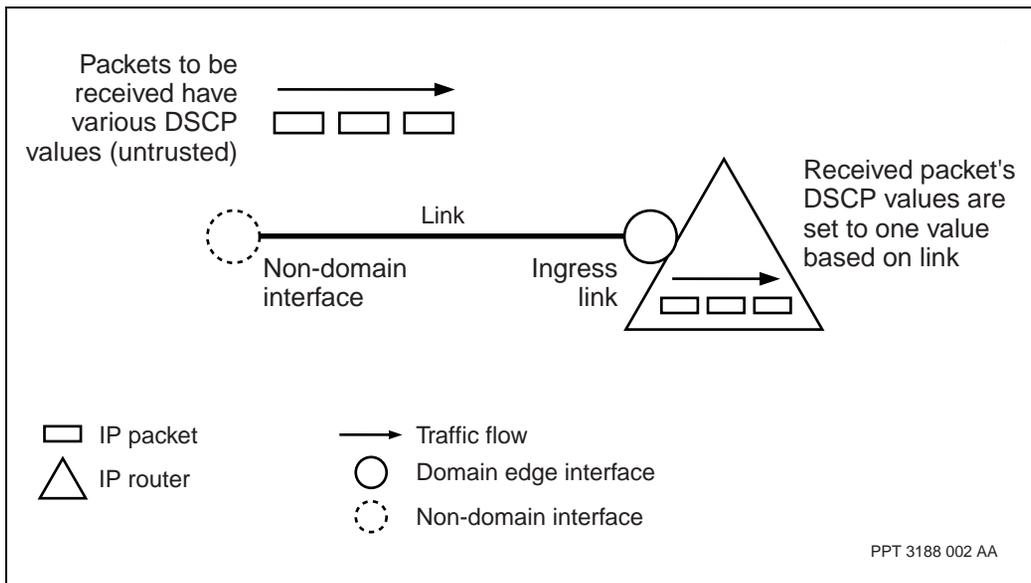


Domain edge interface

A domain edge interface transmits or receives packets from a device outside of a DiffServ domain. The DSCP value of a packet is changed as it is received by a domain edge interface according to the configuration of the link.

Figure 4, “DSCP marking through a domain edge interface,” (page 38) shows how the DSCP of the packet is changed.

Figure 4
DSCP marking through a domain edge interface



Domain boundary interface

A domain boundary interface transmits or receives packets from an IP router that is not in the same DiffServ domain but is in another DiffServ domain.

There are two options for handling the DSCP at a boundary interface.

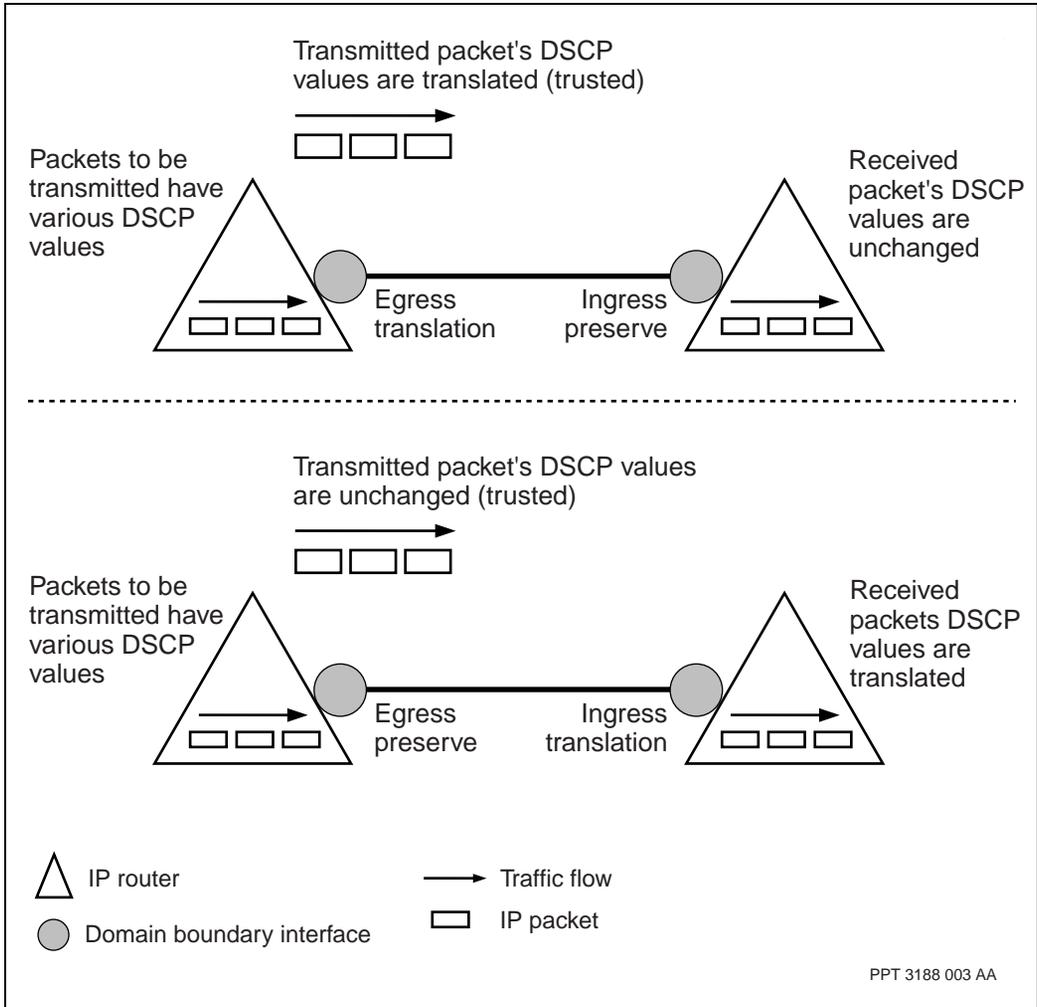
- 1 The DSCP of the packet may be translated as the packet is transmitted by the boundary interface.

This process is known as re-marking or egress IP DSCP marking. It is typically used for boundary interfaces where the next IP hop does not have the ability to translate a received DSCP value into a value with a PHB that its own DiffServ domain can understand.

- 2 The DSCP of the packet may be translated as the packet is received when it enters the boundary interface of the next domain.

Figure 5, “DSCP marking through a domain boundary interface,” (page 40) shows how the DSCP of the packets are mapped to new values.

Figure 5
DSCP marking through a domain boundary interface



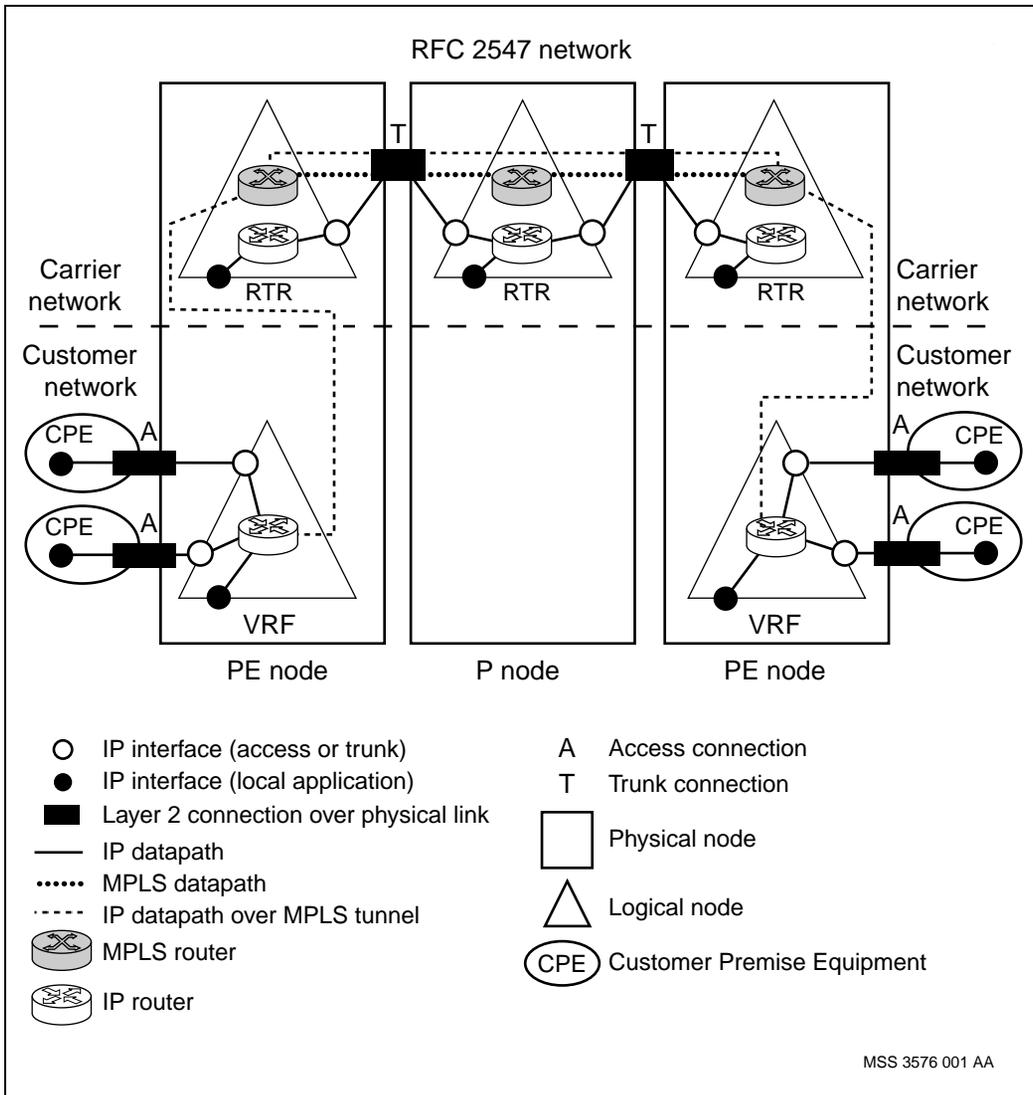
Chapter 4

Layer 3 traffic management in RFC 2547 networks

Use this section to learn about the special RFC 2547 considerations that may affect the planning and implementation of traffic management on your network.

A Layer 3 RFC 2547 network is depicted by Figure 6, “RFC 2547 network,” (page 42).

Figure 6
RFC 2547 network



Layer 3 traffic management on the Multiservice Switch for RFC 2547 networks is DiffServ based, and involves six different data paths described in the sections that follow.

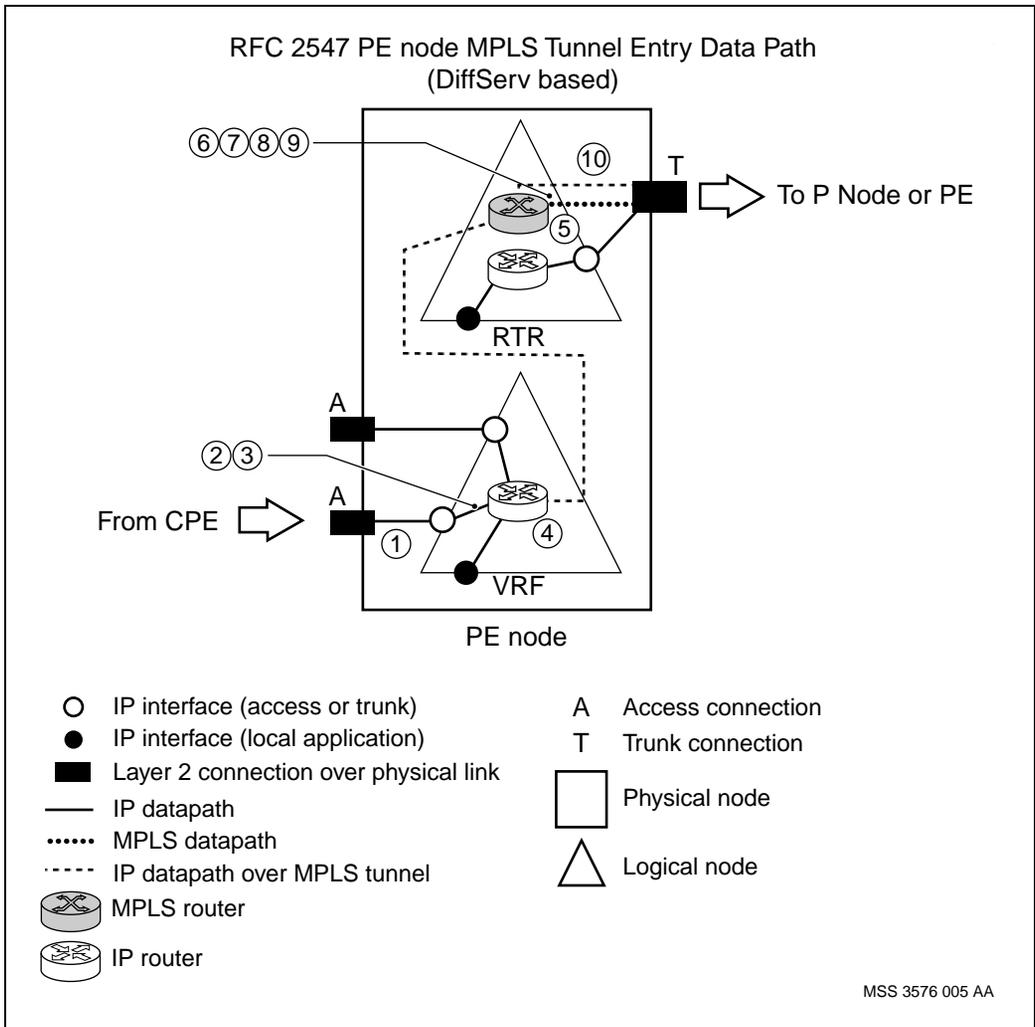
Navigation

- “RFC2547 PE node: MPLS tunnel entry data path” (page 44)
- “RFC 2547 PE node: MPLS tunnel exit data path” (page 50)
- “RFC2547 PE node: IP tandem data path on the Vrf” (page 55)
- “RFC2547 PE node: IP tandem data path on the Rtr” (page 61)
- “RFC2547 P node: MPLS tandem data path” (page 64)
- “RFC 2547 P node: IP tandem data path” (page 67)

RFC2547 PE node: MPLS tunnel entry data path

Table 2, “RFC2547 PE node: MPLS tunnel entry data path,” (page 47) provides information pertaining to the MPLS tunnel entry data path at the PE node. The table column “Index” pertains to the information markers in Figure 7, “RFC2547 PE node: MPLS tunnel entry data path,” (page 45).

Figure 7
RFC2547 PE node: MPLS tunnel entry data path



Legend

Context	Index	Capability
Ingress Access Interface	1	Ingress L2TM
	2	Ingress IP Classification and Marking
	3	Ingress IP Policing
Node	4	IP Routing
	5	MPLS Routing
Egress Trunk Interface	6	MPLS EXP Marking
	7	Select PHB for congestion control
	8	Multi Link congestion control
	9	Single Link Congestion Control
	10	Egress L2TM

Table 2
RFC2547 PE node: MPLS tunnel entry data path

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Access Interface	1	Ingress L2TM	ATM Ethernet Frame Relay PPP	Ingress layer 2 traffic management	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	2	Ingress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior.	Default behavior	none
			MF Map	Map selected fields in IP, TCP, and UDP headers to a PHB value in the Vrf DiffServ domain. Mark the corresponding DSCP value into the IP header.	Optional behavior on PQC12 and PQC6 access FPs.	Optional behavior	Ingress FP requires ipDiffServ.
			DSCP Map	Map IP DSCP field to a PHB value. Mark the corresponding DSCP value into the IP header.	Optional behavior	Optional behavior	Ingress FP requires ipDiffServ
	Link Map	Map the ipCos value of the ingress link to a PHB value in the Vrf DiffServ domain. Mark the corresponding DSCP value into the IP header.	Optional behavior	Optional behavior	Ingress FP requires ipDiffServ		
3	Ingress IP Policing		Map the IP DSCP field to a PHB value in the Vrf DiffServ domain. Use the PHB value to select a meter. Police the IP packet according to the meter configuration and the rate and burst size of the traffic through the meter. If the packet is out of profile with respect to committed information rate and committed burst size then map the PHB to a new PHB and mark the corresponding DSCP value into the IP header. If the packet is out of profile with respect to excess burst size then drop the packet.	Optional behavior on PQC12 and PQC6 access FPs if the ingress media is ATM, Frame Relay, or PPP.	Optional behavior on PQC6 access FPs if the ingress media is ATM, Frame Relay, or PPP.	Ingress FP requires ipPolicing	
(Sheet 1 of 3)							

Table 2 (continued)
RFC2547 PE node: MPLS tunnel entry data path

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Node	4	IP Routing		Route the IP packet toward the next IP hop.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
	5	MPLS Routing		Encapsulate the IP packet with an MPLS service label. Encapsulate the MPLS service label with an MPLS transport label if necessary. Route the MPLS packet toward the next MPLS hop.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Egress Trunk Interface	6	MPLS EXP Marking	Eight way marking	Map the IP DSCP field to a PHB value in the Rtr DiffServ domain. Map the PHB value to an EXP value in the Rtr DiffServ domain. Mark the EXP value into the MPLS service label. Mark the EXP value into MPLS transport label if there is one.	Default behavior on FQM and GQM trunk FPs . Default behavior on PQC12 trunk FPs if the access FP is PQC12. Optional behavior on PQC12 trunk FPs if the access FP is PQC6.	Default behavior on PQC12 access FPs. Optional behavior on PQC6 access FPs.	None
			Four way marking	Map the IP DSCP field to a PHB value in the Vrf DiffServ domain. Map the PHB value to a scheduling class value for four queues (SC4Q) in the Vrf DiffServ domain. Map the SC4Q value to an EXP value in the Rtr DiffServ domain. Mark the EXP value into MPLS service label. Mark the EXP value into MPLS transport label if there is one.	Optional behavior on PQC12 trunk FPs if access FP is PQC6.	Default behavior on PQC6 access FPs.	None
	7	Select PHB for congestion control	DSCP-based in Vrf DiffServ domain	Map the IP DSCP field to a PHB value in the Vrf DiffServ domain.	Not supported.	Optional behavior on PQC6 access FP.	None
			DSCP-based in Rtr DiffServ domain	Map the IP DSCP field to a PHB value in the Rtr DiffServ domain.	Default behavior on FQM and GQM trunk FPs.	Default behavior on PQC6 access FPs.	None

(Sheet 2 of 3)

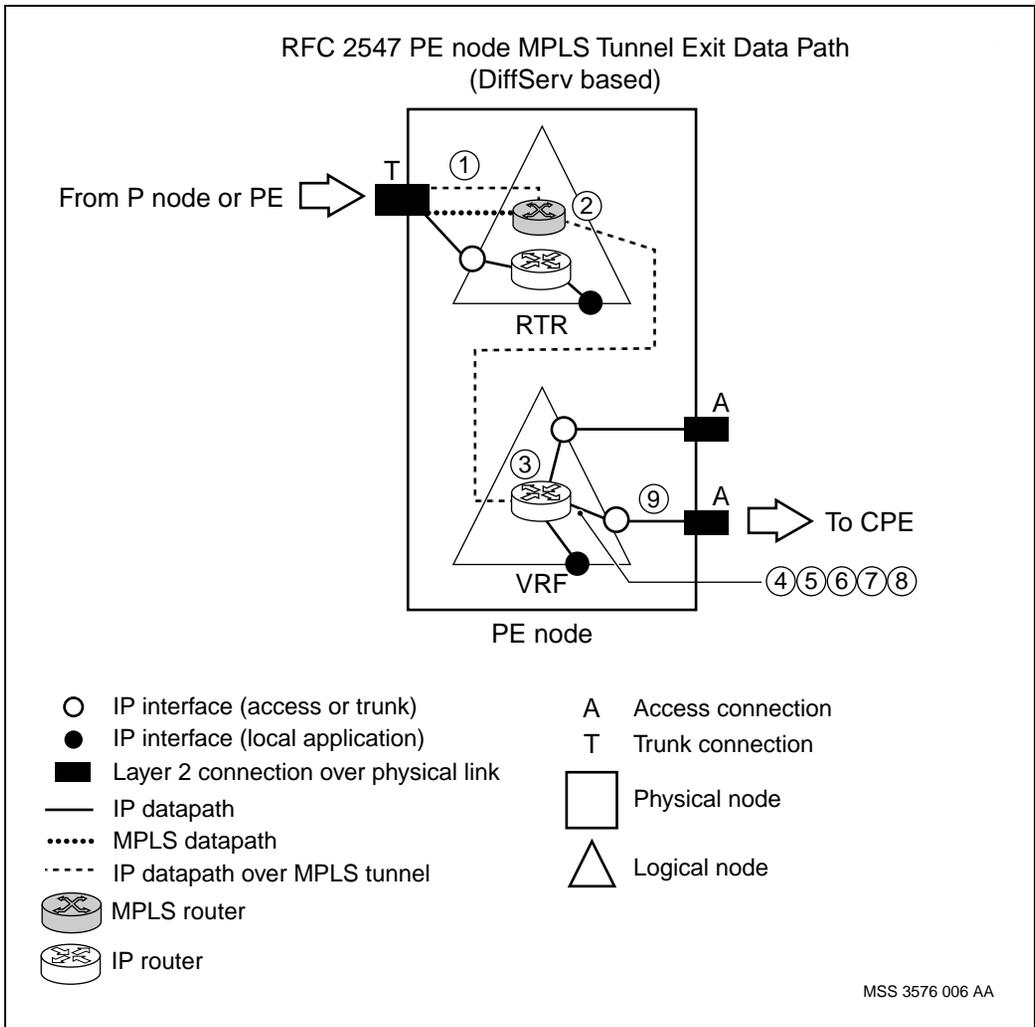
Table 2 (continued)
RFC2547 PE node: MPLS tunnel entry data path

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Trunk Interface (contd)	7 (contd)	Select PHB for congestion control (contd)	EXP-based	Map the MPLS EXP field to a PHB value in the Rtr diffserv domain.	Default behavior on PQC12 trunk FPs.	Not supported.	None
	8	Multi Link congestion control	Select Connection Class	Map the congestion control PHB to a connection class value in the DiffServ domain used to select the PHB. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if the egress media is ATM with VC bundle.	Default behavior if the egress media is ATM with VC bundle.	None
	9	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the DiffServ domain used to select the PHB. Pass this value to egress L2 TM.	Default behavior.	Default behavior.	None
			Select Drop Precedence	Map the congestion control PHB to drop a precedence value in the DiffServ domain used to select the PHB. Pass this value to egress L2 TM.	Default behavior.	Default behavior.	None
	10	Egress L2TM	ATM Ethernet	Egress layer 2 traffic management.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOC3PosAtm FP.							
(Sheet 3 of 3)							

RFC 2547 PE node: MPLS tunnel exit data path

Table 3, “RFC2547 PE node: MPLS tunnel exit data path,” (page 53) provides information pertaining to the MPLS tunnel exit data path at the PE node. The table column “Index” pertains to the information markers in Figure 8, “RFC2547 PE node: MPLS tunnel exit data path,” (page 51).

Figure 8
RFC2547 PE node: MPLS tunnel exit data path



Legend

Context	Index	Capability
Ingress Trunk Interface	1	Ingress L2TM
Node	2	MPLS Routing
	3	IP Routing
Egress Access Interface	4	Egress IP Policing
	5	Select PHB for congestion control
	6	Egress IP Classification and Marking
	7	Multi Link congestion control
	8	Single Link Congestion Control
	9	Egress L2TM

Table 3
RFC2547 PE node: MPLS tunnel exit data path

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Trunk Interface	1	Ingress L2TM	ATM Ethernet	Ingress layer 2 traffic management.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Node	2	MPLS Routing		Remove the MPLS transport label if there is one. Remove the MPLS service label. Forward the packet to the egress access FP.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
	3	IP Routing		Route the IP packet toward the next IP hop.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Egress Access Interface	4	Egress IP Policing		Map the IP DSCP field to a PHB value in the Vrf DiffServ domain. Use the PHB value to select a meter. Police the IP packet according to the meter configuration and the rate and burst size of the traffic through the meter. If the packet is out of profile with respect to committed information rate and committed burst size then map the PHB to a new PHB and mark the corresponding DSCP value into the IP header. If the new and former PHB values are identical then tag the IP packet with "high" drop precedence. If the packet is out of profile with respect to excess burst size then drop the packet.	Optional behavior on PQC12 and PQC6 access FPs if the egress media is Frame Relay or PPP. Cannot be used in combination with egress IP classification and marking with DSCP map.	Optional behavior on PQC6 access FPs if the egress media is Frame Relay or PPP. Cannot be used in combination with egress IP classification and marking with DSCP map.	Egress FP requires ipPolicing
	5	Select PHB for congestion control	DSCP-based	Map the IP DSCP field to a PHB value in the Vrf DiffServ domain.	Default behavior	Default behavior	None
	6	Egress IP Classification and Marking	Preserve DSCP		The IP DSCP field in IP header is not modified.	Default behavior	Default behavior
DSCP Map				Map IP DSCP field to a PHB value. Mark the corresponding DSCP value into the IP header.	Optional behavior. Cannot be used in combination with egress IP Policing.	Optional behavior. Cannot be used in combination with egress IP Policing.	Egress FP requires ipDiffServ
(Sheet 1 of 2)							

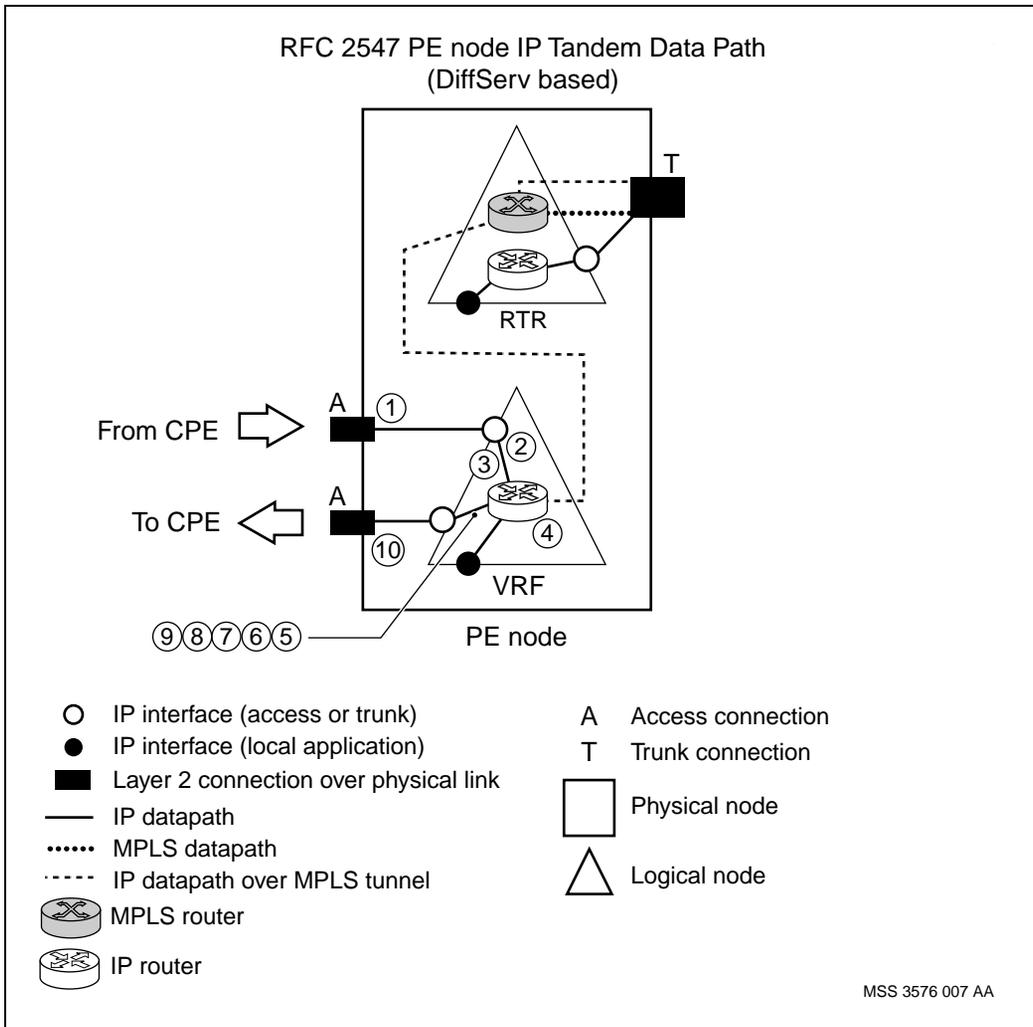
Table 3 (continued)
RFC2547 PE node: MPLS tunnel exit data path

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Access Interface (contd)	7	Multi Link congestion control	Select Connection Class	Map the congestion control PHB to a connection class value in the Vrf DiffServ domain. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if the egress media is ATM or Frame Relay with VC bundle.	Default behavior if the egress media is ATM or Frame Relay with VC bundle.	None
	8	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the Vrf DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Default behavior	None
			Select Drop Precedence	If the packet was tagged with "high" drop precedence at the egress IP policing stage then use that value otherwise map the congestion control PHB to drop a precedence value in the Vrf DiffServ domain. Pass the drop precedence value to egress L2 TM.	Default behavior	Default behavior	None
	9	Egress L2TM	ATM Ethernet Frame Relay PPP	Egress layer 2 traffic management.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Comments: "N/A" indicates not within the scope of L3 traffic management. "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOC3PosAtm FP.							
(Sheet 2 of 2)							

RFC2547 PE node: IP tandem data path on the Vrf

Table 4, “RFC2547 PE Node: IP tandem data path on the Vrf,” (page 58) provides information pertaining to the IP tandem data path on the Vrf, at the PE node. The table column “Index” pertains to the information markers in Figure 9, “RFC2547 PE node: IP tandem data path on the Vrf,” (page 56).

Figure 9
RFC2547 PE node: IP tandem data path on the Vrf



Legend

Context	Index	Capability
Ingress Access Interface	1	Ingress L2TM
	2	Ingress IP Classification and Marking
	3	Ingress IP Policing
Node	4	IP Routing
Egress Access Interface	5	Egress IP Policing
	6	Select PHB for congestion control
	7	Egress IP Classification and Marking
	8	Multi Link congestion control
	9	Single Link Congestion Control
	10	Egress L2TM

Table 4
RFC2547 PE Node: IP tandem data path on the Vrf

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Access Interface	1	Ingress L2TM	ATMEthernet Frame Relay PPP	Ingress layer 2 traffic management.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
	2	Ingress IP Classification and Marking	Preserve DSCP	The IP DSCP field in IP header is not modified.	Default behavior	Default behavior	None
			MF Map	Map selected fields in IP, TCP, and UDP headers to a PHB value in the Vrf DiffServ domain. Mark the corresponding DSCP value into the IP header.	Optional behavior on PQC12 and PQC6 ingress FPs	Optional behavior	Ingress FP requires ipDiffServ
			DSCP Map	Map IP DSCP field to a PHB value. Mark the corresponding DSCP value into IP header.	Optional behavior	Optional behavior	Ingress FP requires ipDiffServ
			Link Map	Map the ipCos value of ingress link to a PHB value in the Vrf DiffServ domain. Mark the corresponding DSCP value into the IP header.	Optional behavior	Optional behavior	Ingress FP requires ipDiffServ
3	Ingress IP Policing		Map the IP DSCP field to a PHB value in the Vrf DiffServ domain. Use the PHB value to select a meter. Police the IP packet according to the meter configuration and the rate and burst size of the traffic through the meter. If the packet is out of profile with respect to committed information rate and committed burst size then map the PHB to a new PHB and mark the corresponding DSCP value into the IP header. If the packet is out of profile with respect to excess burst size then drop the packet.	Optional behavior on PQC12 and PQC6 ingress FPs if the ingress media is ATM, Frame Relay, or PPP	Optional behavior on PQC6 ingress FPs if the ingress media is ATM, Frame Relay, or PPP	Ingress FP requires ipPolicing	
(Sheet 1 of 3)							

Table 4 (continued)
RFC2547 PE Node: IP tandem data path on the Vrf

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Node	4	IP Routing		Route the IP packet toward the next IP hop.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Egress Access Interface	5	Egress IP Policing		Map the IP DSCP field to a PHB value in the Vrf DiffServ domain. Use the PHB value to select a meter. Police the IP packet according to the meter configuration and the rate and burst size of the traffic through the meter. If the packet is out of profile with respect to committed information rate and committed burst size then map the PHB to a new PHB and mark the corresponding DSCP value into the IP header. If the new and former PHB values are identical then tag the IP packet with "high" drop precedence. If the packet is out of profile with respect to excess burst size then drop the packet.	Optional behavior on PQC12 and PQC6 egress FPs if the egress media is Frame Relay. Cannot be used in combination with egress IP classification and marking with DSCP map	Optional behavior on PQC6 egress FPs if the egress media is Frame Relay. Cannot be used in combination with egress IP classification and marking with DSCP map	Egress FP requires ipPolicing
	6	Select PHB for congestion control	DSCP-based	Map the IP DSCP field to a PHB value in the Vrf DiffServ domain.	Default behavior	Default behavior	None
	7	Egress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			DSCP Map	Map IP DSCP field to a PHB value. Mark the corresponding DSCP value into the IP header.	Optional behavior. Cannot be used in combination with egress IP Policing	Optional behavior. Cannot be used in combination with egress IP Policing	Ingress PQC12 and PQC6 FPs require ipDiffServ. Egress FQM and GQM FPs require ipDiffServ
(Sheet 2 of 3)							

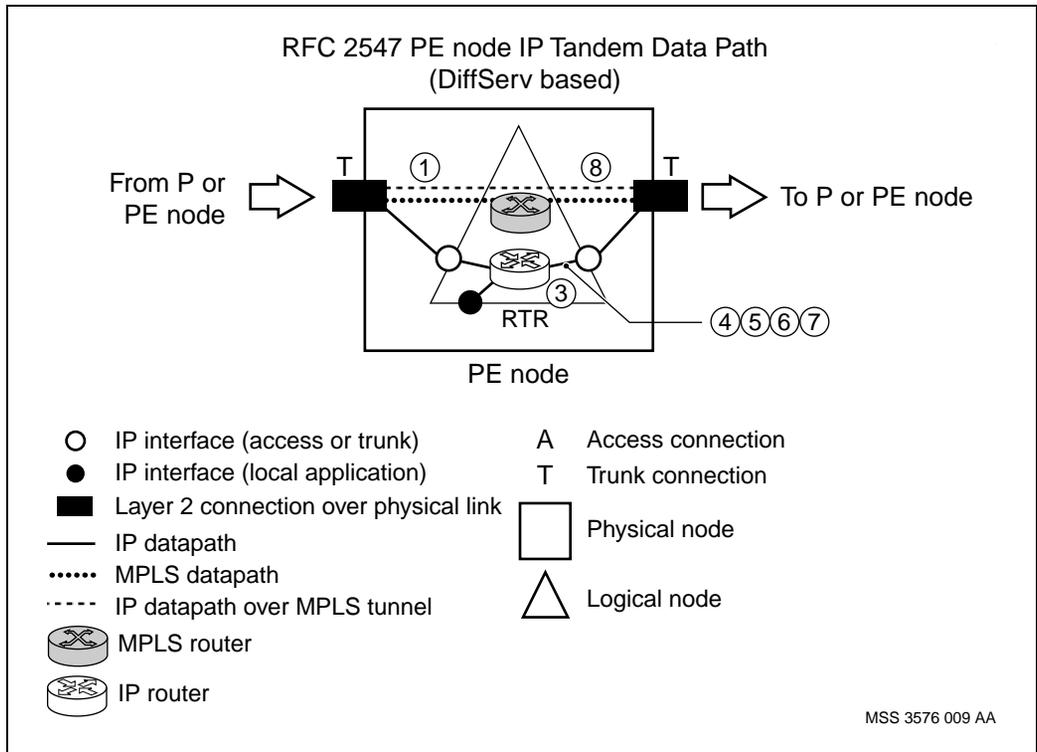
Table 4 (continued)
RFC2547 PE Node: IP tandem data path on the Vrf

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Access Interface (contd)	8	Multi Link congestion control	Select Connection Class	Map the congestion control PHB to a connection class value in the Vrf DiffServ domain. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if the egress media is ATM or FrameRelay with VC bundle	Default behavior if the egress media is ATM or Frame Relay with VC bundle	None
	9	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the Vrf DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Default behavior	None
			Select Drop Precedence	If the packet was tagged with "high" drop precedence at the egress IP policing stage then use that value otherwise map the congestion control PHB to drop a precedence value in the Vrf DiffServ domain. Pass the drop precedence value to egress L2 TM.	Default behavior	Default behavior	None
	10	Egress L2TM	ATMEthernet Frame Relay PPP	Egress layer 2 traffic management.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "QGM" indicates 16pOC3PosAtm FP.							
(Sheet 3 of 3)							

RFC2547 PE node: IP tandem data path on the Rtr

Table 5, “RFC2547 PE node: IP tandem data path on the Rtr,” (page 63) provides information pertaining to the IP tandem data path on the Rtr, at the PE node. The table columns “Index” pertains to the information markers in Figure 10, “RFC2547 PE node: IP tandem data path on the Rtr,” (page 61).

Figure 10
RFC2547 PE node: IP tandem data path on the Rtr



Legend

Context	Index	Capability
Ingress Trunk Interface	1	Ingress L2TM
	2	Ingress IP Classification and Marking
Node	3	IP Routing
Egress Trunk Interface	4	Select PHB for congestion control
	5	Egress IP Classification and Marking
	6	Multi Link Congestion Control
	7	Single Link Congestion Control
	8	Egress L2TM

Table 5
RFC2547 PE node: IP tandem data path on the Rtr

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Trunk Interface	1	Ingress L2TM	ATM Ethernet	ATM Ethernet	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	2	Ingress IP Classification and Marking	Preserve DSCP	The IP DSCP field in IP header is not modified.	Default behavior	Default behavior	None
Node	3	IP Routing		Route the IP packet toward the next IP hop.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
Egress Trunk Interface	4	Select PHB for congestion control	DSCP-based	Map the IP DSCP field to a PHB value in the Rtr DiffServ domain.	Default behavior	Default behavior	None
	5	Egress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
	6	Multi Link Congestion Control	Select Connection Class	Map the congestion control PHB to a connection class value in the Rtr DiffServ domain. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if the egress media is ATM with VC bundle	Default behavior if the egress media is ATM with VC bundle	None
	7	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the Rtr DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Default behavior	None
			Select Drop Precedence	Map the congestion control PHB to drop a precedence value in the Rtr DiffServ domain. Pass this value to egress L2 TM.	Default behavior for all trunk FPs	Default behavior for all trunk FPs	none
	8	Egress L2TM	ATM Ethernet	Egress layer 2 traffic management.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOC3PosAtm FP.							

RFC2547 P node: MPLS tandem data path

Table 6, “RFC2547 P node: MPLS tandem data path,” (page 65) provides information pertaining to the MPLS tandem data path at the P node. The table column “Index” pertains to the information markers in Figure 11, “RFC2547 P node: MPLS tandem data path,” (page 64).

Figure 11
RFC2547 P node: MPLS tandem data path

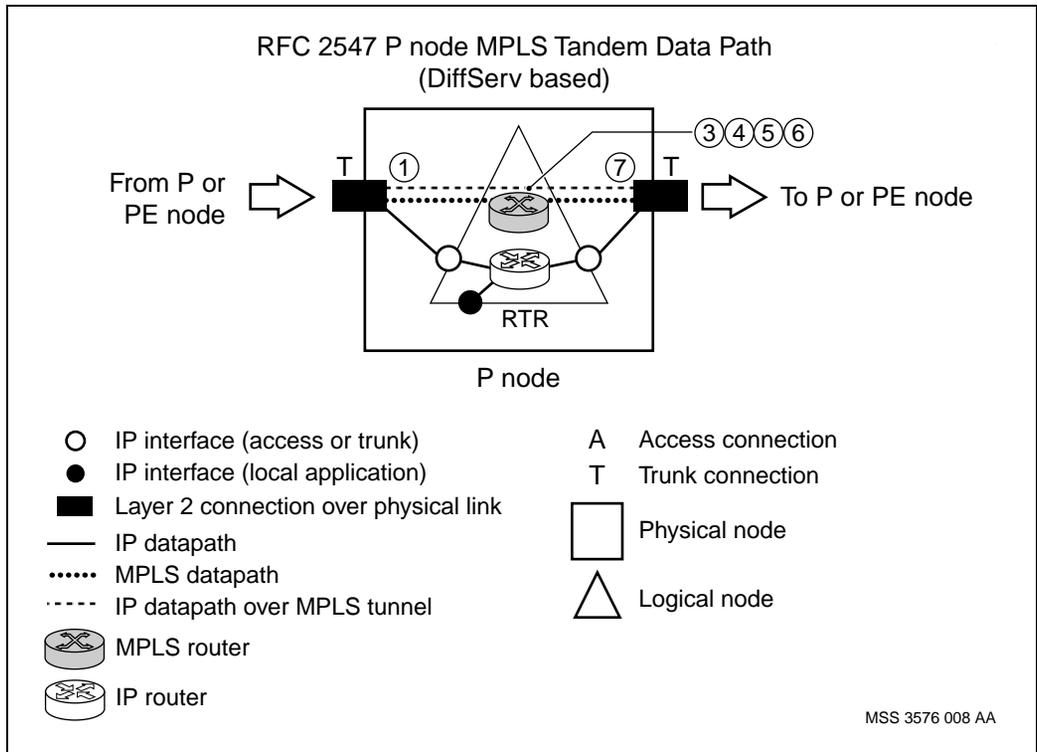


Table 6
RFC2547 P node: MPLS tandem data path

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Trunk Interface	1	Ingress L2TM	ATM Ethernet Frame Relay PPP	Ingress layer 2 traffic management	Not in L3TM scope	Not in L3TM scope	Not in L3TM scope
Node	2	MPLS Routing		Remove the incoming MPLS transport label. Encapsulate the MPLS service label with an MPLS transport label if necessary. Route the MPLS packet toward the next MPLS hop.	Not in L3TM scope	Not in L3TM scope	Not in L3TM scope
Egress Trunk Interface	3	MPLS EXP Marking	Eight way marking	The EXP field in the MPLS service label is not modified. If there is an outgoing MPLS transport label, then set its EXP field equal to the EXP field of the MPLS service label.	Default behavior	Default behavior	None
	4	Select PHB for Congestion Control	EXP-based	Map the MPLS EXP field to a PHB value in the Rtr diffserv domain.	Default behavior	Default behavior	None
	5	Multi Link congestion control	Select Connection Class	Map the congestion control PHB to a connection class value in the Rtr DiffServ domain. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if egress the media is ATM with VC bundle	Default behavior if the egress media is ATM with VC bundle	None
	6	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the Rtr DiffServ domain . Pass this value to egress L2 TM.	Default behavior	Default behavior	None
			Select Drop Precedence	Map the congestion control PHB to drop a precedence value in the Rtr DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Default behavior	None
7	Egress L2TM	ATM Ethernet	Egress layer 2 traffic management.	Not in L3TM scope	Not in L3TM scope	Not in L3TM scope	
(Sheet 1 of 2)							

Table 6 (continued)
RFC2547 P node: MPLS tandem data path

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOC3PosAtm FP.							
(Sheet 2 of 2)							

RFC 2547 P node: IP tandem data path

This data path is identical to the data path described in section “RFC2547 PE node: IP tandem data path on the Rtr” (page 61).

Chapter 5

Layer 3 traffic management in RFC 2764 networks

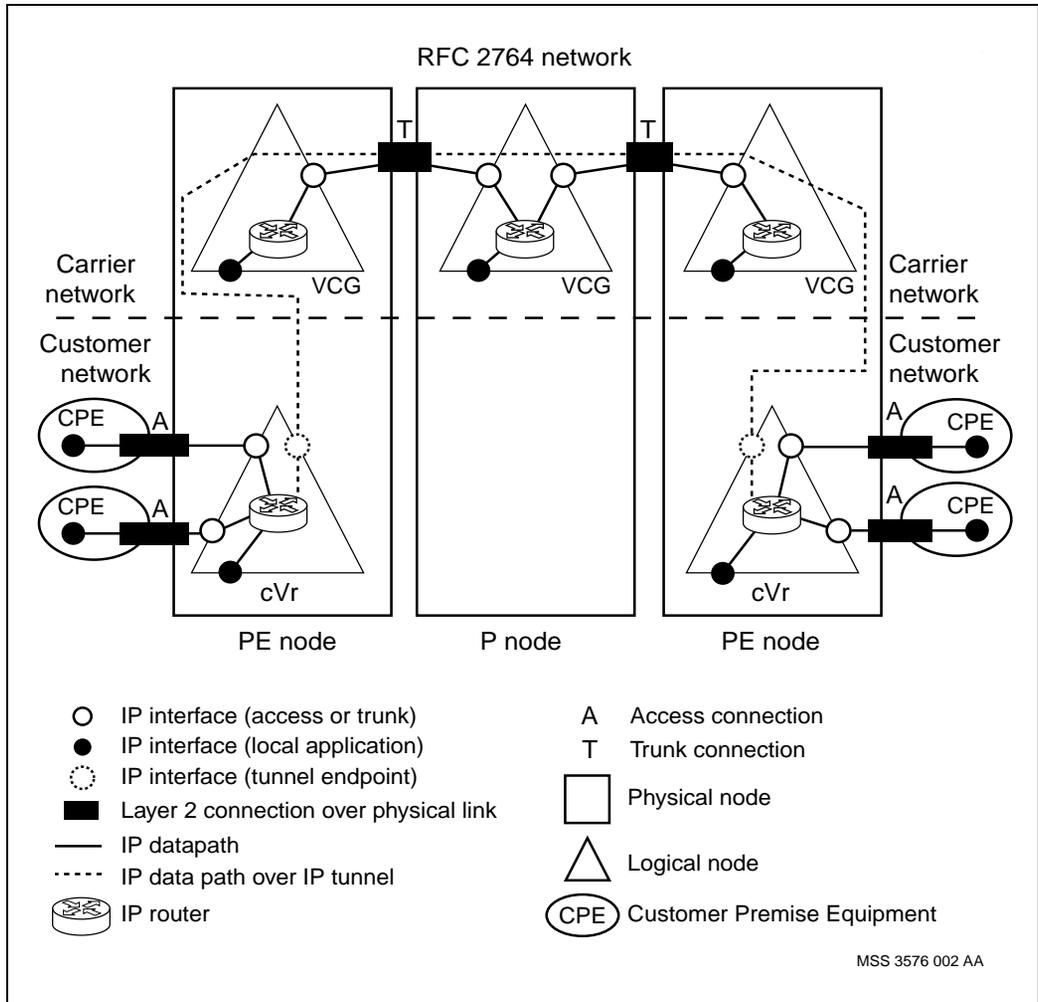
Use this section to learn about the special RFC 2764 considerations that may affect the planning and implementation of traffic management on your network.

A Layer 3 RFC 2764 network is depicted by Figure 12, “RFC 2764 network,” (page 70).

Layer 3 traffic management on the Multiservice Switch for an RFC 2764 PE node can be DiffServ based, Class of Service (CoS) based, or DiffServ/CoS hybrid based.

Layer 3 traffic management on a RFC 2764 P node is the same as Layer 3 traffic management on a VIPR node. See “Layer 3 traffic management in VIPR networks” (page 107) for more information.

Figure 12
RFC 2764 network



Navigation

- “DiffServ based layer 3 traffic management in RFC 2764 networks” (page 71)
- “CoS based layer 3 traffic management in RFC 2764 networks” (page 85)
- “DiffServ/Cos hybrid based Layer 3 traffic management in RFC 2764 networks” (page 96)

DiffServ based layer 3 traffic management in RFC 2764 networks

Layer 3 traffic management on a RFC 2764 PE node is DiffServ based if every virtual router on the node (the VCG and all the customer VRs) has a DiffServDomain subcomponent.

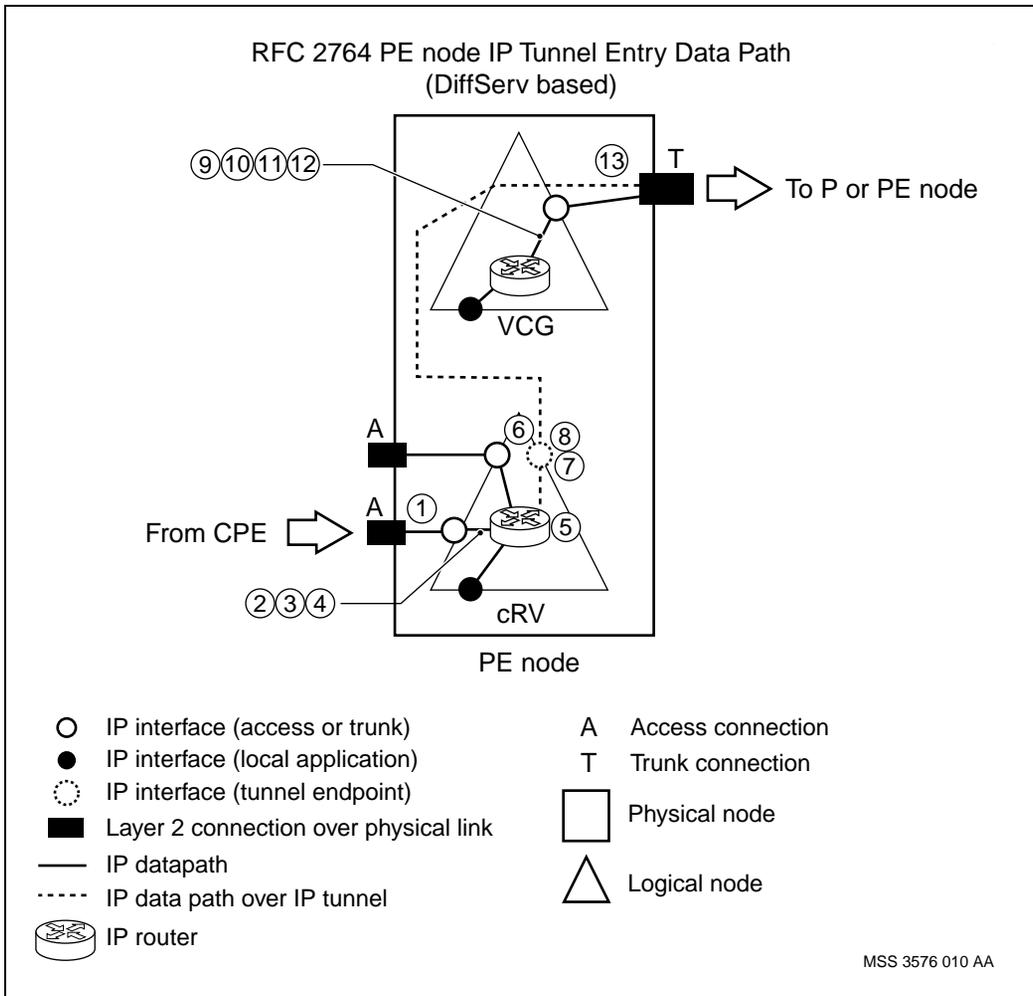
DiffServ based layer 3 traffic management in RFC 2764 networks involves four different data paths.

- RFC 2764 PE node: IP tunnel entry data path
- RFC 2764 PE node: IP tunnel exit data path
- RFC 2764 PE node: IP tandem data path on the cVR
- RFC 2764 PE node: IP tandem data path on the VCG

RFC 2764 PE node: IP tunnel entry data path

Table 7, “RFC 2764 PE node: IP tunnel entry data path, DiffServ based,” (page 74) provides information pertaining to the IP tunnel entry data path at the PE node, for DiffServ based networks. The table column “Index” pertains to the information markers in Figure 13, “RFC 2764 PE node: IP tunnel entry data path, DiffServ based,” (page 72).

Figure 13
RFC 2764 PE node: IP tunnel entry data path, DiffServ based



Legend

Context	Index	Capability
Ingress Access Interface	1	Ingress L2TM
	2	Ingress IP Firewall
	3	Ingress IP Classification and Marking
	4	Ingress IP Policing
Node	5	IP Routing
IP Tunnel Interface	6	Egress IP Classification and Marking
	7	IP Tunnel Entry
	8	Carrier IP DSCP Marking
Node	9	IP Routing
Egress Trunk Interface	10	Select PHB for Congestion Control
	11	Multi Link Congestion Control
	12	Single Link Congestion Control
	13	Egress L2TM

Table 7
RFC 2764 PE node: IP tunnel entry data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Access Interface	1	Ingress L2TM	ATM Ethernet Frame Relay PPP	Ingress layer 2 traffic management	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	2	Ingress IP Firewall		Map the source and destination fields in the IP header to a permit or deny action according to the IP firewall configuration used by the ingress interface. If the action is deny then drop the packet.	Optional behavior	Optional behavior on PQC12 and PQC6 access FPs	Ingress FP requires ipFilter
	3	Ingress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			MF Map	Map selected fields in IP, TCP, and UDP headers to a PHB value in the cVR DiffServ domain according to the DiffServProfile MfMap components used by the ingress interface. Mark the corresponding DSCP value into the IP header.	Optional behavior	Optional behavior	Ingress FP requires ipDiffServ
(Sheet 1 of 4)							

Table 7 (continued)
RFC 2764 PE node: IP tunnel entry data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Access Interface (contd)	3 (contd)	Ingress IP Classification and Marking (contd)	DSCP Map	Map the IP DSCP field to a PHB value in the cVR DiffServ domain according to the DiffServProfile DscpMap component used by the ingress interface. Mark the corresponding DSCP value into the IP header.	Optional behavior	Optional behavior	Ingress FP requires ipDiffServ
			Link Map	Map the ipCos value of the ingress link to a PHB value in the cVR DiffServ domain according to the DiffServProfile LinkMap component used by the ingress interface. Mark the corresponding DSCP value into the IP header.	Optional behavior	Optional behavior	Ingress FP requires ipDiffServ
	4	Ingress IP Policing	Route the IP packet toward the next IP hop.	Optional behavior if the ingress media is ATM, Frame Relay, or PPP	Optional behavior on PQC6 access FPs if the ingress media is ATM, Frame Relay, or PPP	Ingress FP requires ipPolicing	
Node	5	IP Routing		Route the IP packet toward the next IP hop.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
(Sheet 2 of 4)							

Table 7 (continued)
RFC 2764 PE node: IP tunnel entry data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
IP Tunnel Interface	6	Egress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			DSCP Map	Map the IP DSCP field to a PHB value in the cVR DiffServ domain according to the DiffServProfile DscpMap component used by the IP tunnel interface. Mark the corresponding DSCP value into the IP header.	Optional behavior	Optional behavior.	Ingress FP requires ipDiffServ
	7	IP Tunnel Entry		Encapsulate the IP packet with a carrier IP header.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
	8	Carrier IP DSCP Marking	Without Customer to Carrier SLA	Set the value of the carrier IP DSCP field equal to the value of the customer IP DSCP field.	Default behavior	Default behavior	None
			With Customer to Carrier SLA	Map the MPLS EXP field to a PHB value in the Rtr diffserv domain.	Default behavior on PQC12 trunk FPs.	Not supported.	None
Node	9	IP Routing		Route the carrier IP packet toward the next IP hop.	Not in L3 TM scope	Not in L3 TM scope	None
Egress Trunk Interface	10	Select PHB for Congestion Control	DSCP-based	Map the carrier IP DSCP field to a PHB value in the VCG DiffServ domain.	Default behavior.	Default behavior.	None
	11	Multi Link Congestion Control	Select Connection Class	Map the congestion control PHB to a connection class value in the VCG DiffServ domain. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if egress media is ATM with VC bundle	Default behavior if egress media is ATM with VC bundle	None
	12	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the VCG DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Default behavior	None
			Select Drop Precedence	Map the congestion control PHB to a drop precedence value in the VCG DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Default behavior	None

(Sheet 3 of 4)

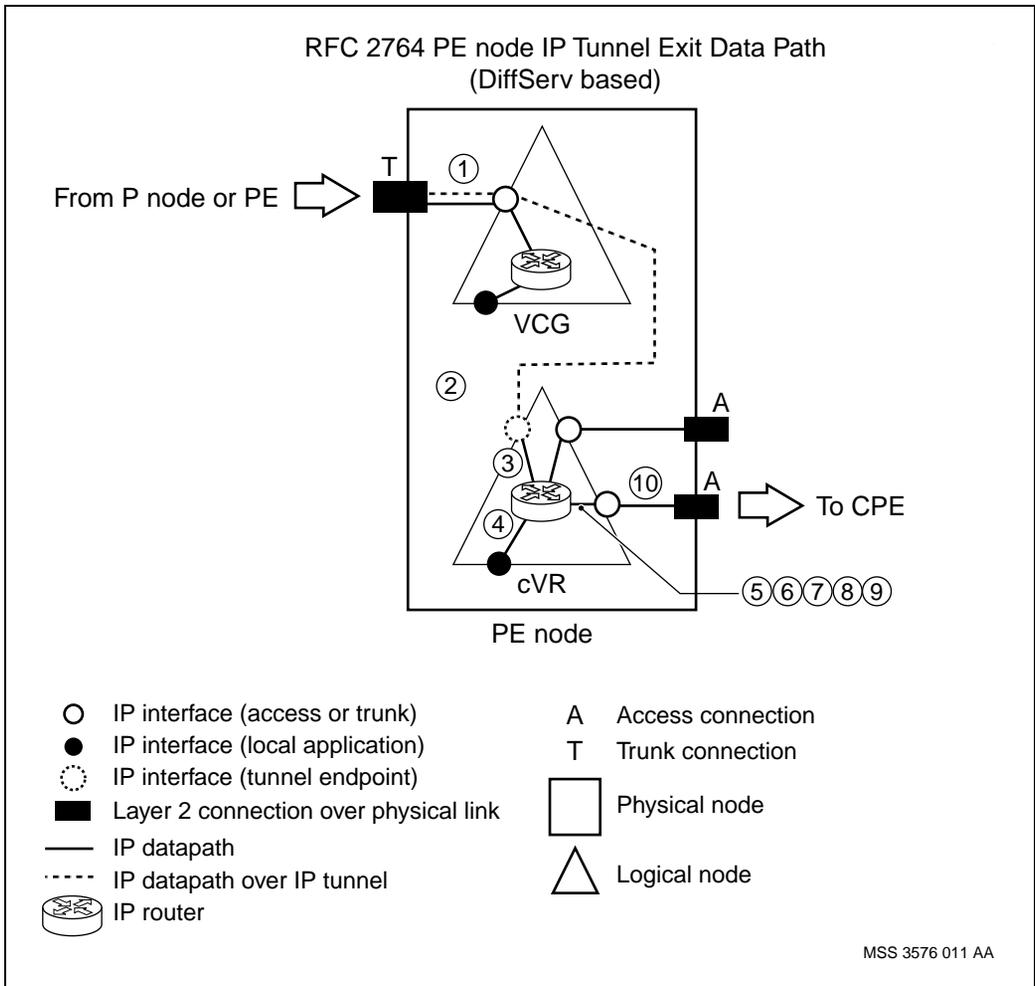
Table 7 (continued)
RFC 2764 PE node: IP tunnel entry data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Trunk Interface (contd)	13	Egress L2TM	ATM	Egress layer 2 traffic management.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOC3PosAtm FP.							
(Sheet 4 of 4)							

RFC 2764 PE node: IP tunnel exit data path, DiffServ based

Table 8, “RFC 2764 PE node: IP tunnel exit data path, DiffServ based,” (page 81) provides information pertaining to the IP tunnel entry data path at the PE node, for DiffServ based networks. The table column “Index” pertains to the information markers in Figure 14, “RFC 2764 PE node: IP tunnel exit data path, DiffServ based,” (page 79).

Figure 14
RFC 2764 PE node: IP tunnel exit data path, DiffServ based



Legend

Context	Index	Capability
Ingress Trunk Interface	1	Ingress L2TM
IP Tunnel Interface	2	IP Tunnel Exit
	3	Ingress IP Classification and Marking
Node	4	IP Routing
Egress Access Interface	5	Egress IP Policing
	6	Select PHB for Congestion Control
	7	Egress IP Classification and Marking
	8	Multi Link Congestion Control
	9	Single Link Congestion Control
	10	Egress L2TM

Table 8
RFC 2764 PE node: IP tunnel exit data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Trunk Interface	1	Ingress L2TM	ATM	Ingress layer 2 traffic management	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
IP Tunnel Interface	2	IP Tunnel Exit	Optimized	Forward the carrier IP packet to the egress access FP. Remove the carrier IP header.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
			Not optimized	Remove the carrier IP header. Forward the customer IP packet to the egress access FP.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
	3	Ingress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			MF Map	Map selected fields in IP, TCP, and UDP headers to a PHB value in the Vrf DiffServ domain according to the DiffServProfile MfMap components used by the IP tunnel interface. Mark the corresponding DSCP value into the IP header.	Optional behavior	Optional behavior	Ingress FP requires ipDiffServ if tunnel is not optimized. Egress FP requires ipDiffServ if tunnel is optimized
			DSCP Map	Map the IP DSCP field to a PHB value according to the DiffServProfile DscpMap component used by the IP tunnel interface. Mark the corresponding DSCP value into the IP header.	Optional behavior	Optional behavior	Ingress FP requires ipDiffServ if tunnel is not optimized. Egress FP requires ipDiffServ if tunnel is optimized
Node	4	IP Routing		Route the IP packet toward the next IP hop.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
(Sheet 1 of 3)							

Table 8 (continued)
RFC 2764 PE node: IP tunnel exit data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Access Interface	5	Egress IP Policing		Map the IP DSCP field to a PHB value in the cVR DiffServ domain. Use the PHB value to select a meter according to the PolicerProfile Meter components used by the egress interface. Police the IP packet according to the meter configuration and the rate and burst size of the traffic through the meter. If the packet is out of profile with respect to committed information rate and committed burst size then map the PHB to a new PHB and mark the corresponding DSCP value into the IP header. If the new and former PHB values are identical then tag the IP packet with "high" drop precedence. If the packet is out of profile with respect to excess burst size then drop the packet.	Optional behavior if egress media is Frame Relay and tunnel is not optimized. Optional behavior if egress media is Frame Relay or PPP and tunnel is optimized. Cannot be used in combination with egress IP classification and marking with DSCP map	Optional behavior on PQC6 access FPs if egress media is Frame Relay and tunnel is not optimized. Optional behavior on PQC6 access FPs if egress media is Frame Relay or PPP and tunnel is optimized. Cannot be used in combination with egress IP classification and marking with DSCP map	Egress FP requires ipPolicer
	6	Select PHB for Congestion Control	DSCP-based in Vrf DiffServ domain	Map the IP DSCP field to a PHB value in the cVR DiffServ domain.	Default behavior	Default behavior	None
	7	Egress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			DSCP Map	Map the IP DSCP field to a PHB value in the cVR DiffServDomain according to the DiffServProfile DscpMap used by the egress interface. Mark the corresponding DSCP value into the IP header.	Optional behavior. Cannot be used in combination with egress IP Policing.	Optional behavior. Cannot be used in combination with egress IP Policing.	Egress FP requires ipDiffServ
(Sheet 2 of 3)							

Table 8 (continued)
RFC 2764 PE node: IP tunnel exit data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Access Interface (contd)	8	Multi Link Congestion Control	Select Connection Class	Map the congestion control PHB to a connection class value in the cVR DiffServ domain. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if egress media is ATM or Frame Relay with VC bundle	Default behavior if egress media is ATM or Frame Relay with VC bundle	None
	9	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the cVR DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Default behavior	None
			Select Drop Precedence	If the packet was tagged with "high" drop precedence at the egress IP policing stage then use that value otherwise map the congestion control PHB to a drop precedence value in the cVR DiffServ domain. Pass the drop precedence value to egress L2 TM.	Default behavior	Default behavior	None
	10	Egress L2TM	ATM Ethernet Frame Relay PPP	Egress layer 2 traffic management.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOC3PosAtm FP.							
(Sheet 3 of 3)							

RFC 2764 PE node: IP tandem data path on the cVR, DiffServ based

This data path is identical to the data path presented in section “VIPR node: IP tandem data path, DiffSev based” (page 109). See that section for more information.

RFC 2764 PE node: IP tandem data path on the VCG, DiffServ based

This data path is identical to the data path presented in section “VIPR node: IP tandem data path, DiffSev based” (page 109). See that section for more information.

CoS based layer 3 traffic management in RFC 2764 networks

Layer 3 traffic management on a RFC 2764 PE node is CoS based if every virtual router on the node (the VCG and all the customer VRs) does not have a DiffServDomain subcomponent. Layer 3 traffic management on a RFC 2764 P node is CoS based if the VCG virtual router does not have a DiffServDomain subcomponent.

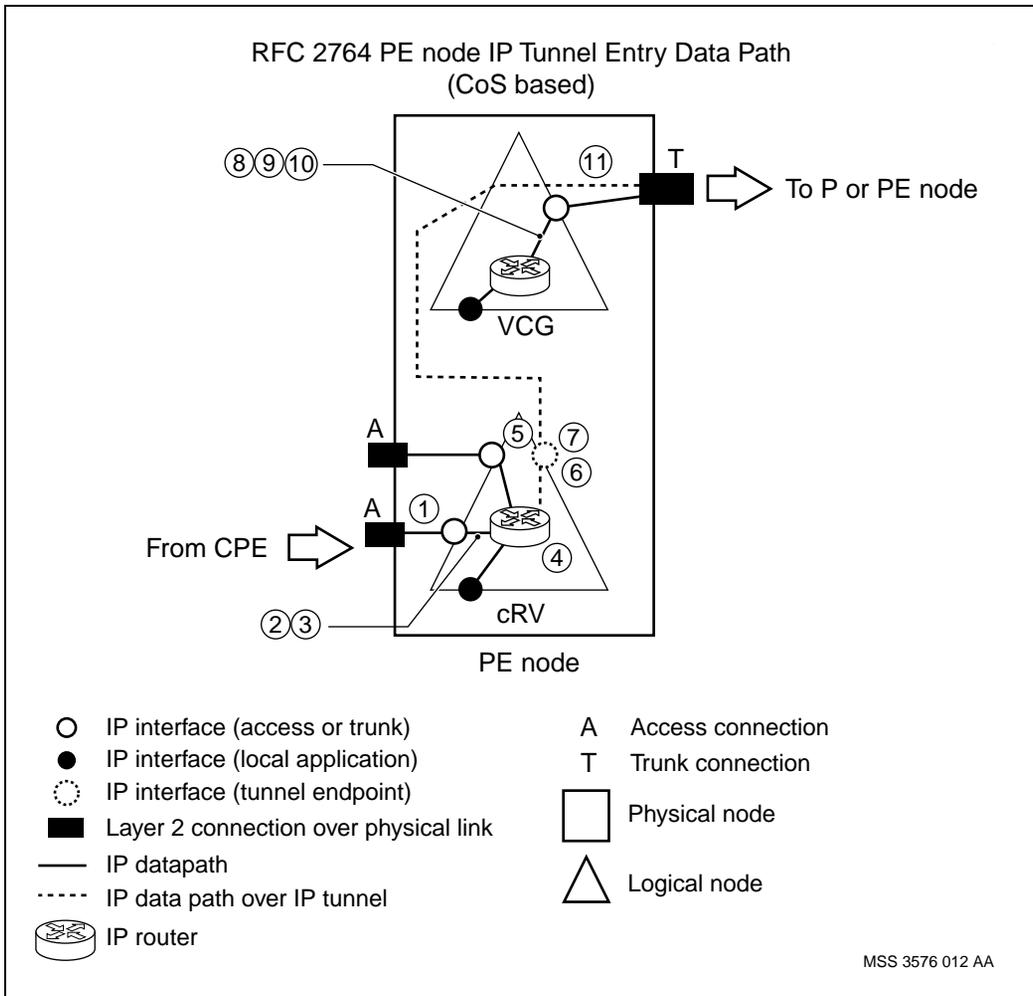
CoS based layer 3 traffic management in RFC 2764 networks involves four different data paths.

- RFC 2764 PE node: IP tunnel entry data path
- RFC 2764 PE node: IP tunnel exit data path
- RFC 2764 PE node: IP tandem data path on the cVR
- RFC 2764 PE node: IP tandem data path on the VCG

RFC 2764 PE node: IP tunnel entry data path, CoS based

Table 9, “RFC 2764 PE node: IP tunnel entry data path, CoS based,” (page 88) provides information pertaining to the IP tunnel entry data path at the PE node, for CoS based layer 3 traffic management in RFC 2764 networks. The table column “Index” pertains to the information markers in Figure 15, “RFC 2764 PE node: IP tunnel entry data path, CoS based,” (page 86).

Figure 15
RFC 2764 PE node: IP tunnel entry data path, CoS based



Legend

Context	Index	Capability
Ingress Access Interface	1	Ingress L2TM
	2	Ingress IP Firewall
	3	Ingress IP Classification
Node	4	IP Routing
IP Tunnel Interface	5	Customer IP DSCP Marking
	6	IP Tunnel Entry
	7	Carrier IP DSCP Marking
Node	8	IP Routing
Egress Trunk Interface	9	Multi Link Congestion Control
	10	Single Link Congestion Control
	11	Egress L2TM

Table 9
RFC 2764 PE node: IP tunnel entry data path, CoS based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Access Interface	1	Ingress L2TM	ATM Ethernet Frame Relay PPP	Ingress layer 2 traffic management	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	2	Ingress IP Firewall		Map the source and destination fields in the IP header to a permit or deny action according to the IP firewall configuration used by the ingress interface. If the action is deny then drop the packet.	Optional behavior	Optional behavior on PQC12 and PQC6 access FPs	Ingress FP requires ipFilter
	3	Ingress IP Classification	MF Map	Map selected fields in IP, TCP, and UDP headers to a CoS value according to the CosPolicyGroup Policy components used by the ingress interface. Tag the IP packet with this value.	Optional behavior	Optional behavior	Ingress FP requires ipCos
			DSCP Map	Map the IP DSCP field to a CoS value according to the CosPolicyGroup Policy components used by the ingress interface. Tag the IP packet with this value.	Optional behavior	Optional behavior	Ingress FP requires ipCos
			Link Map	Identify the ipCos value of the ingress link. Tag the IP packet with this value.	Default behavior	Default behavior	None
Node	4	IP Routing		Route the IP packet toward the next IP hop.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
IP Tunnel Interface	5	Customer IP DSCP Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			Modify DSCP	Map the tagged CoS value to a DSCP value according to the CosPolicyGroup Ect components used by the IP tunnel interface. Mark the DSCP value into the IP header.	Optional behavior	Optional behavior.	Ingress FP requires ipCos
	6	IP Tunnel Entry		Encapsulate the IP packet with a carrier IP header.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
(Sheet 1 of 2)							

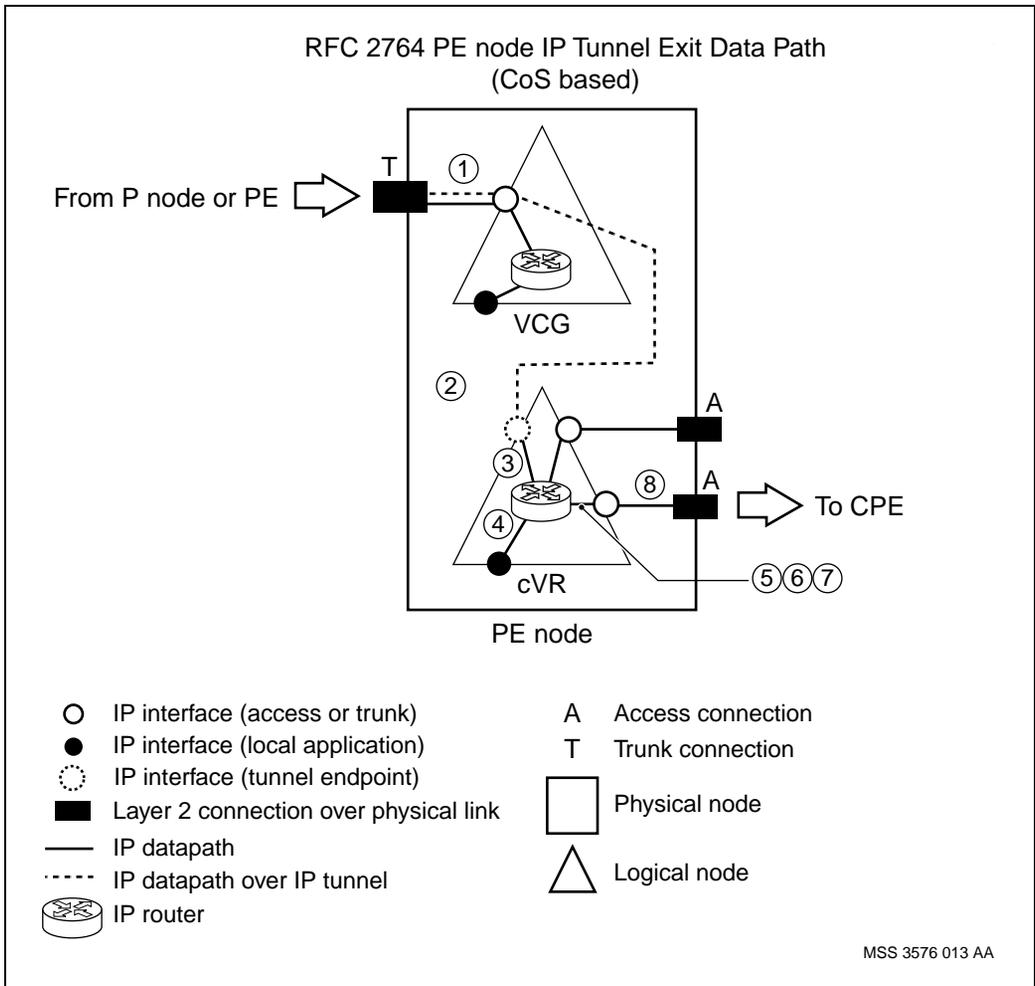
Table 9 (continued)
RFC 2764 PE node: IP tunnel entry data path, CoS based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
IP Tunnel Interface (contd)	7	Carrier IP DSCP Marking		Set the value of the carrier IP DSCP field equal to the value of the customer IP DSCP field.	Default behavior	Default behavior	None
Node	8	IP Routing		Route the carrier IP packet toward the next IP hop.	Not in L3 TM scope	Not in L3 TM scope	None
Egress Trunk Interface	9	Multi Link Congestion Control	Select Connection Class	Use the tagged CoS value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if the egress media is ATM with VC bundle	Default behavior if the egress media is ATM with VC bundle	None
			10	Single Link Congestion Control	Select Scheduling Class	Pass scheduling class value 0 to egress L2 TM.	Default behavior
	Select Drop Precedence	Map the tagged CoS value to an emission priority value according to the CosPolicyGroup Ect components used by the egress interface. Map the emission priority value to a scheduling class value. Pass the scheduling class value to egress L2 TM.			Default behavior	Default behavior	None
		Pass the drop precedence derived at ingress L2 TM to egress L2 TM.		Default behavior.	Default behavior	None	
	11	Egress L2TM	ATM	Map the tagged CoS value to a drop precedence value according to the CosPolicyGroup Ict components used at the ingress interface. Pass the drop precedence value to egress L2 TM.	Optional behavior	Optional behavior	Ingress FP requires ipCos
Egress layer 2 traffic management.				Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope	
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1).							
(Sheet 2 of 2)							

RFC 2764 PE node: IP tunnel exit data path, CoS based

Table 10, “RFC 2764 PE node: IP tunnel exit data path, CoS based,” (page 93) provides information pertaining to the IP tunnel entry data path at the PE node, for CoS based layer 3 traffic management in RFC 2764 networks. The table column “Index” pertains to the information markers in Figure 16, “RFC 2764 PE node: IP tunnel exit data path, CoS based,” (page 91).

Figure 16
RFC 2764 PE node: IP tunnel exit data path, CoS based



Legend

Context	Index	Capability
Ingress Trunk Interface	1	Ingress L2TM
IP tunnel interface	2	IP Tunnel Exit
	3	Ingress IP Classification
Node	4	IP Routing
Egress Access Interface	5	Customer IP DSCP Marking
	6	Multi Link Congestion Control
	7	Single Link Congestion Control
Egress L2TM	8	ATM Ethernet Frame Relay PPP

Table 10
RFC 2764 PE node: IP tunnel exit data path, CoS based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Trunk Interface	1	Ingress L2TM	ATM	Ingress layer 2 traffic management	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
IP tunnel interface	2	IP Tunnel Exit	Optimized	Forward the carrier IP packet to the egress access FP. Remove the carrier IP header.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
			Not optimized	Remove the carrier IP header. Forward the customer IP packet to the egress access FP.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
	3	Ingress IP Classification	MF Map	Map selected fields in IP, TCP, and UDP headers to a CoS value according to the CosPolicyGroup Policy components used by the IP tunnel interface. Tag the IP packet with this value.	Optional behavior	Optional behavior	Ingress FP requires ipCos if tunnel is not optimized. Egress FP requires ipCos if tunnel is optimized
			DSCP Map	Map the IP DSCP field to a CoS value according to the CosPolicyGroup Policy components used by the IP tunnel interface. Tag the IP packet with this value.	Optional behavior	Optional behavior	Ingress FP requires ipCos if tunnel is not optimized. Egress FP requires ipCos if tunnel is optimized
			Link Map	Identify the ipCos value of the ingress link. Tag the IP packet with this value.	Default behavior	Default behavior	None
Node	4	IP Routing		Route the IP packet toward the next IP hop.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
(Sheet 1 of 3)							

Table 10 (continued)
RFC 2764 PE node: IP tunnel exit data path, CoS based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Access Interface	5	Customer IP DSCP Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			Modify DSCP	Map the tagged CoS value to a DSCP value according to the CosPolicyGroup Ect components used at the egress interface. Mark the DSCP value into the IP header.	Optional behavior	Optional behavior.	Ingress FP requires ipCos if tunnel is not optimized. Egress FP requires ipCos if tunnel is optimized.
	6	Multi Link Congestion Control	Select Connection Class	Use the tagged CoS value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if the egress media is ATM or Frame Relay with VC bundle	Default behavior if the egress media is ATM or Frame Relay with VC bundle	None
	7	Single Link Congestion Control	Select Scheduling Class	Pass scheduling class value 0 to egress L2 TM.	Default behavior	Default behavior	None
				Map the tagged CoS value to an emission priority value according to the CosPolicyGroup Ect components used by the egress interface. Map the emission priority value to a scheduling class value. Pass the scheduling class value to egress L2 TM.	Optional behavior	Optional behavior	Ingress FP requires ipCos if tunnel is not optimized. Egress FP requires ipCos if tunnel is optimized.
	(Sheet 2 of 3)						

Table 10 (continued)
RFC 2764 PE node: IP tunnel exit data path, CoS based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Access Interface (contd)	7 (contd)	Single Link Congestion Control (contd)	Select Drop Precedence	Pass the drop precedence derived by ingress L2 TM to egress L2 TM.	Default behavior	Default behavior	None
				Map the tagged CoS value to a drop precedence value according to the CosPolicyGroup lct components used by the IP tunnel interface. Pass the drop precedence value to egress L2 TM.	Optional behavior	Optional behavior	Ingress FP requires ipCos if tunnel is not optimized. Egress FP requires ipCos if tunnel is optimized
Egress L2TM	8	ATM Ethernet Frame Relay PPP	Egress layer 2 traffic management	Route the carrier IP packet toward the next IP hop.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1).							
(Sheet 3 of 3)							

RFC 2764 PE node: IP tandem data path on the cVR, CoS based

This data path is identical to the data path presented in section “CoS based layer 3 traffic management in VIPR networks” (page 116). See that section for more information.

RFC 2764 PE node: IP tandem data path on the VCG, CoS based

This data path is identical to the data path presented in section “CoS based layer 3 traffic management in VIPR networks” (page 116). See that section for more information.

DiffServ/Cos hybrid based Layer 3 traffic management in RFC 2764 networks

Layer 3 traffic management on a RFC 2764 PE node is DiffServ/CoS hybrid based if the VCG virtual router has a DiffServDomain subcomponent but the customer virtual routers do not. This option is relevant to RFC 2764 networks that are undergoing migration from CoS based layer 3 traffic management to DiffServ based layer 3 traffic management

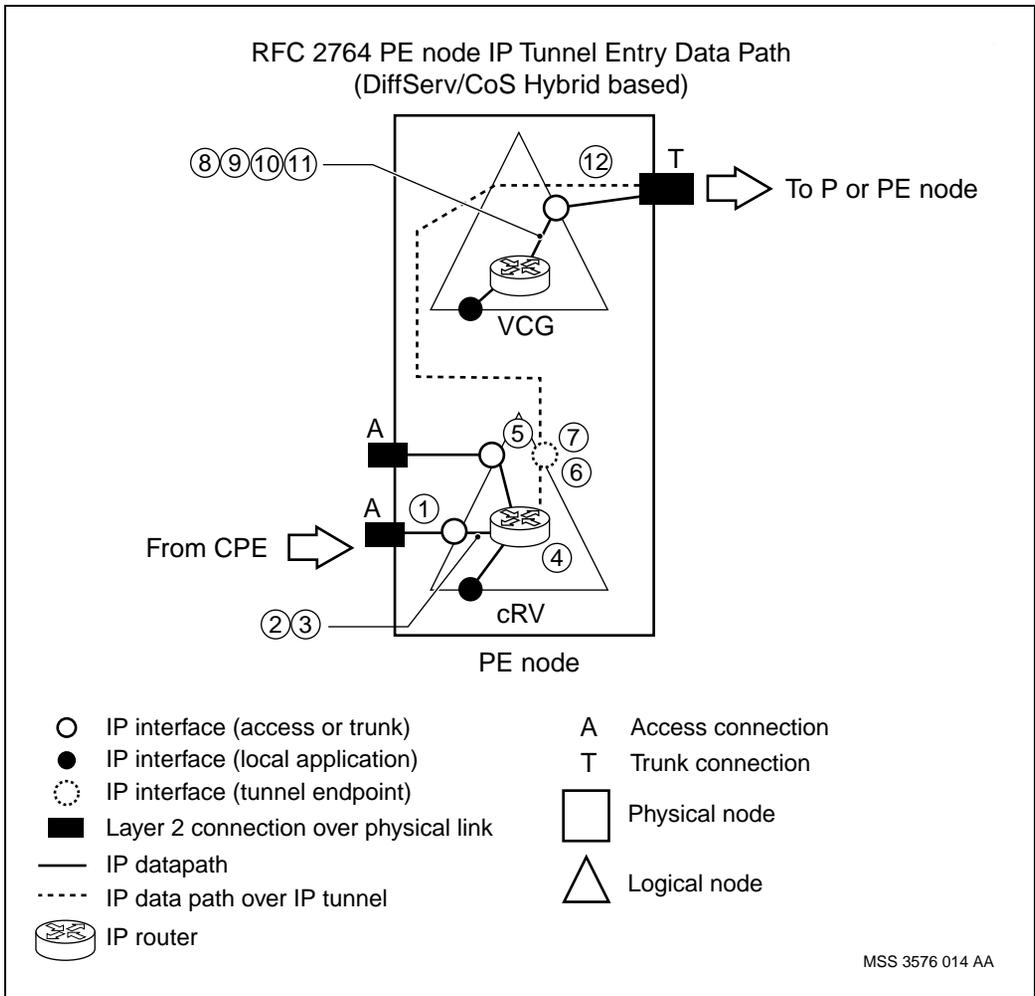
DiffServ/Cos hybrid based layer 3 traffic management in RFC 2764 networks involves four different data paths.

- RFC 2764 PE node: IP tunnel entry data path
- RFC 2764 PE node: IP tunnel exit data path
- RFC 2764 PE node: IP tandem data path on the cVR
- RFC 2764 PE node: IP tandem data path on the VCG

RFC 2764 PE node: IP tunnel entry data path, DiffServ/CoS hybrid based

Table 11, “RFC 2764 PE node: IP tunnel entry data path, DiffServ/CoS hybrid based,” (page 99) provides information pertaining to the IP tunnel entry data path at the PE node, for DiffServ/CoS hybrid based networks. The table column “Index” pertains to the information markers in Figure 17, “RFC 2764 PE node: IP tunnel entry data path, DiffServ/CoS hybrid based,” (page 97).

Figure 17
RFC 2764 PE node: IP tunnel entry data path, DiffServ/CoS hybrid based



Legend

Context	Index	Capability
Ingress Access Interface	1	Ingress L2TM
	2	Ingress IP Firewall
	3	Ingress IP Classification
Node	4	IP Routing
IP Tunnel Interface	5	Customer IP DSCP Marking
	6	IP Tunnel Entry
	7	Carrier IP DSCP Marking
Node	8	IP Routing
Egress Trunk Interface	9	Select PHB for Congestion Control
	10	Multi Link Congestion Control
	11	Single Link Congestion Control
	12	Egress L2TM

Table 11
RFC 2764 PE node: IP tunnel entry data path, DiffServ/CoS hybrid based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Access Interface	1	Ingress L2TM	ATM Ethernet Frame Relay PPP	Ingress layer 2 traffic management	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	2	Ingress IP Firewall		Map the source and destination fields in the IP header to a permit or deny action according to the IP firewall configuration used by the ingress interface. If the action is deny then drop the packet.	Optional behavior	Optional behavior on PQC12 and PQC6 ingress FPs	Ingress FP requires ipFilter
	3	Ingress IP Classification	MF Map	Map selected fields in IP, TCP, and UDP headers to a CoS value according to the CosPolicyGroup Policy components used by the ingress interface. Tag the IP packet with this value.	Optional behavior	Optional behavior	Ingress FP requires ipCos
			DSCP Map	Map the IP DSCP field to a CoS value according to the CosPolicyGroup Policy components used by the ingress interface. Tag the IP packet with this value.	Optional behavior	Optional behavior	Ingress FP requires ipCos
			Link Map	Identify the ipCos value of the ingress link. Tag the IP packet with this value.	Default behavior	Default behavior	None
Node	4	IP Routing		Route the IP packet toward the next IP hop.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
IP Tunnel Interface	5	Customer IP DSCP Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			Modify DSCP	Map the tagged CoS value to a DSCP value according to the CosPolicyGroup Ect components used by the IP tunnel interface. Mark the DSCP value into the IP header.	Optional behavior	Optional behavior.	Ingress FP requires ipCos
	6	IP Tunnel Entry		Encapsulate the IP packet with a carrier IP header.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
(Sheet 1 of 2)							

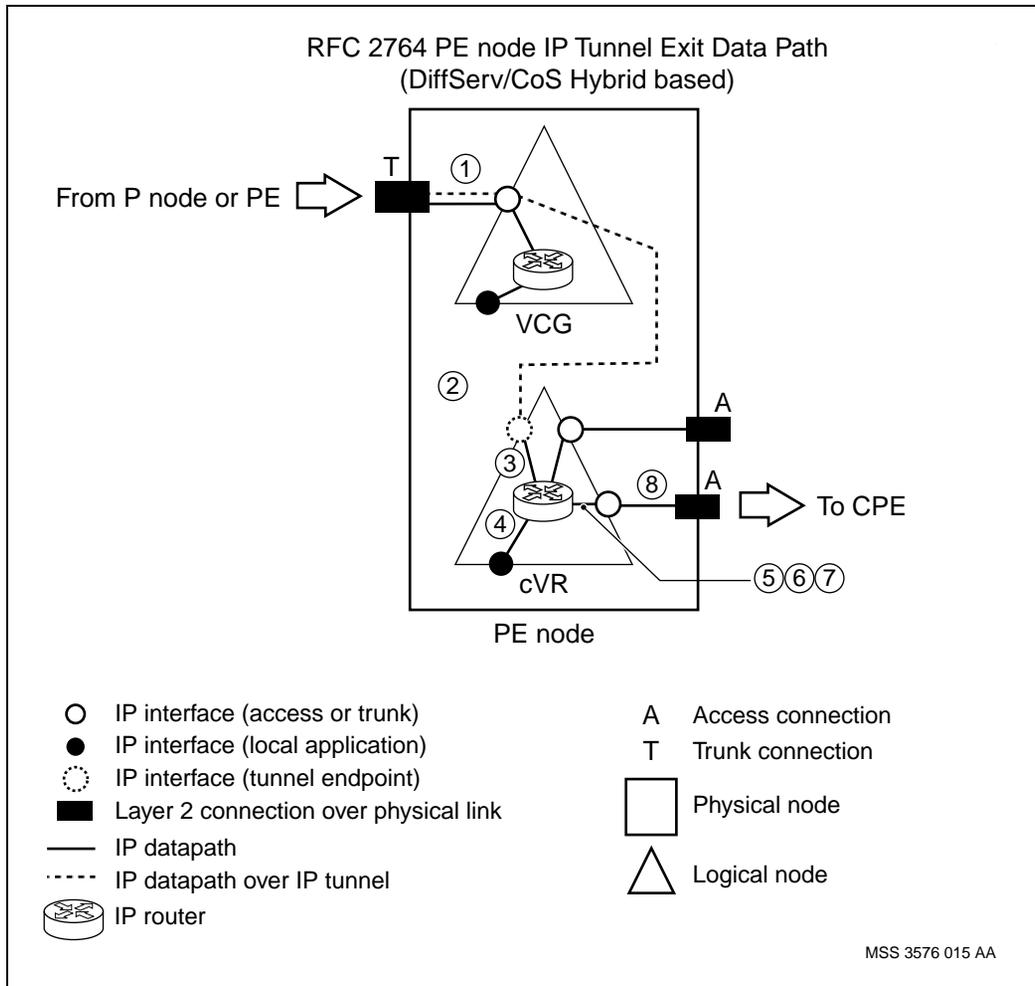
Table 11 (continued)
RFC 2764 PE node: IP tunnel entry data path, DiffServ/CoS hybrid based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
IP Tunnel Interface (contd)	7	Carrier IP DSCP Marking	Without Customer to Carrier SLA	If the value of customer IP DSCP field matches the phbRoutingSource attribute value of the VCG DiffServ domain then set the value of the carrier IP DSCP field equal to the value of the customer IP DSCP field; otherwise set the value of the carrier IP DSCP field to 0.	Default behavior	Default behavior	None
			With Customer to Carrier SLA	Map the tagged CoS value to a PHB value in the VCG DiffServ domain using the VpnCosMap component. Mark the corresponding DSCP value into the carrier IP header.	Optional behavior.	Optional behavior.	Ingress FP requires ipDiffServ.
Node	8	IP Routing		Route the carrier IP packet toward the next IP hop.	Not in L3 TM scope	Not in L3 TM scope	None
Egress Trunk Interface	9	Select PHB for Congestion Control	DSCP-based	Map the carrier IP DSCP field to a PHB value in the VCG DiffServ domain.	Default behavior	Default behavior	None
	10	Multi Link Congestion Control	Select Connection Class	Map the congestion control PHB to a connection class value in the VCG DiffServ domain. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if egress media is ATM with VC bundle	Default behavior if egress media is ATM with VC bundle	None
	11	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the VCG DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Default behavior	None
			Select Drop Precedence	Pass the drop precedence derived at ingress L2 TM to egress L2 TM.	Default behavior.	Default behavior	None
	12	Egress L2TM	ATM	Egress layer 2 traffic management.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1).							
(Sheet 2 of 2)							

RFC 2764 PE node: IP tunnel exit data path, DiffServ/CoS hybrid based

Table 12, “RFC 2764 PE node: IP tunnel exit data path, DiffServ/CoS hybrid based,” (page 104) provides information pertaining to the IP tunnel exit data path at the PE node, for DiffServ/CoS hybrid based networks. The table column “Index” pertains to the information markers in Figure 18, “RFC 2764 PE node: IP tunnel exit data path, DiffServ/CoS hybrid based,” (page 102).

Figure 18
RFC 2764 PE node: IP tunnel exit data path, DiffServ/CoS hybrid based



Legend

Context	Index	Capability
Ingress Trunk Interface	1	Ingress L2TM
IP tunnel interface	2	IP Tunnel Exit
	3	Ingress IP Classification
Node	4	IP Routing
Egress Access Interface	5	Customer IP DSCP Marking
	6	Multi Link Congestion Control
	7	Single Link Congestion Control
Egress L2TM	8	ATM Ethernet Frame Relay PPP

Table 12
RFC 2764 PE node: IP tunnel exit data path, DiffServ/CoS hybrid based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Trunk Interface	1	Ingress L2TM	ATM	Ingress layer 2 traffic management	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
IP tunnel interface	2	IP Tunnel Exit	Optimized	Forward the carrier IP packet to the egress access FP. Remove the carrier IP header.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
			Not optimized	Remove the carrier IP header. Forward the customer IP packet to the egress access FP.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
	3	Ingress IP Classification	MF Map	Map selected fields in IP, TCP, and UDP headers to a CoS value according to the CosPolicyGroup Policy components used by the IP tunnel interface. Tag the IP packet with this value.	Optional behavior	Optional behavior	Ingress FP requires ipCos if tunnel is not optimized. Egress FP requires ipCos if tunnel is optimized
			DSCP Map	Map the IP DSCP field to a CoS value according to the CosPolicyGroup Policy components used by the IP tunnel interface. Tag the IP packet with this value.	Optional behavior	Optional behavior	Ingress FP requires ipCos if tunnel is not optimized. Egress FP requires ipCos if tunnel is optimized
			Link Map	Map the carrier IP DSCP field to a PHB value in the VCG DiffServ domain. Map the PHB to a connection class value in the VCG DiffServ domain. Tag the IP packet with a CoS value equal to the connection class value.	Default behavior	Default behavior	None
Node	4	IP Routing		Route the IP packet toward the next IP hop.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
(Sheet 1 of 3)							

Table 12 (continued)

RFC 2764 PE node: IP tunnel exit data path, DiffServ/CoS hybrid based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Access Interface	5	Customer IP DSCP Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			Modify DSCP	Map the tagged CoS value to a DSCP value according to the CosPolicyGroup Ect components used at the egress interface. Mark the DSCP value into the IP header.	Optional behavior	Optional behavior.	Ingress FP requires ipCos if tunnel is not optimized. Egress FP requires ipCos if tunnel is optimized.
	6	Multi Link Congestion Control	Select Connection Class	Use the tagged CoS value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if the egress media is ATM or Frame Relay with VC bundle	Default behavior if the egress media is ATM or Frame Relay with VC bundle	None
	7	Single Link Congestion Control	Select Scheduling Class	Pass scheduling class value 0 to egress L2 TM.	Default behavior	Default behavior	None
				Map the tagged CoS value to an emission priority value according to the CosPolicyGroup Ect components used by the egress interface. Map the emission priority value to a scheduling class value. Pass the scheduling class value to egress L2 TM.	Optional behavior	Optional behavior	Ingress FP requires ipCos if tunnel is not optimized. Egress FP requires ipCos if tunnel is optimized.

(Sheet 2 of 3)

Table 12 (continued)**RFC 2764 PE node: IP tunnel exit data path, DiffServ/CoS hybrid based**

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Access Interface (contd)	7 (contd)	Single Link Congestion Control (contd)	Select Drop Precedence	Map the carrier IP DSCP field to a PHB value in the VCG DiffServ domain. Map the PHB to a drop precedence value in the VCG DiffServ domain. Pass the drop precedence value to egress L2 TM.	Default behavior	Default behavior	None
				Map the tagged CoS value to a drop precedence value according to the CosPolicyGroup Ict components used by the IP tunnel interface. Pass the drop precedence value to egress L2 TM.	Optional behavior	Optional behavior	Ingress FP requires ipCos if tunnel is not optimized. Egress FP requires ipCos if tunnel is optimized
Egress L2TM	8	ATM Ethernet Frame Relay PPP		Egress layer 2 traffic management.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1).							
(Sheet 3 of 3)							

RFC 2764 PE node: IP tandem data path on the cVR, DiffServ/CoS hybrid based

This data path is identical to the data path presented in section “CoS based layer 3 traffic management in VIPR networks” (page 116). See that section for more information.

RFC 2764 PE node: IP tandem data path on the VCG, DiffServ/CoS hybrid based

This data path is identical to the data path presented in section “VIPR node: IP tandem data path, DiffServ based” (page 109). See that section for more information.

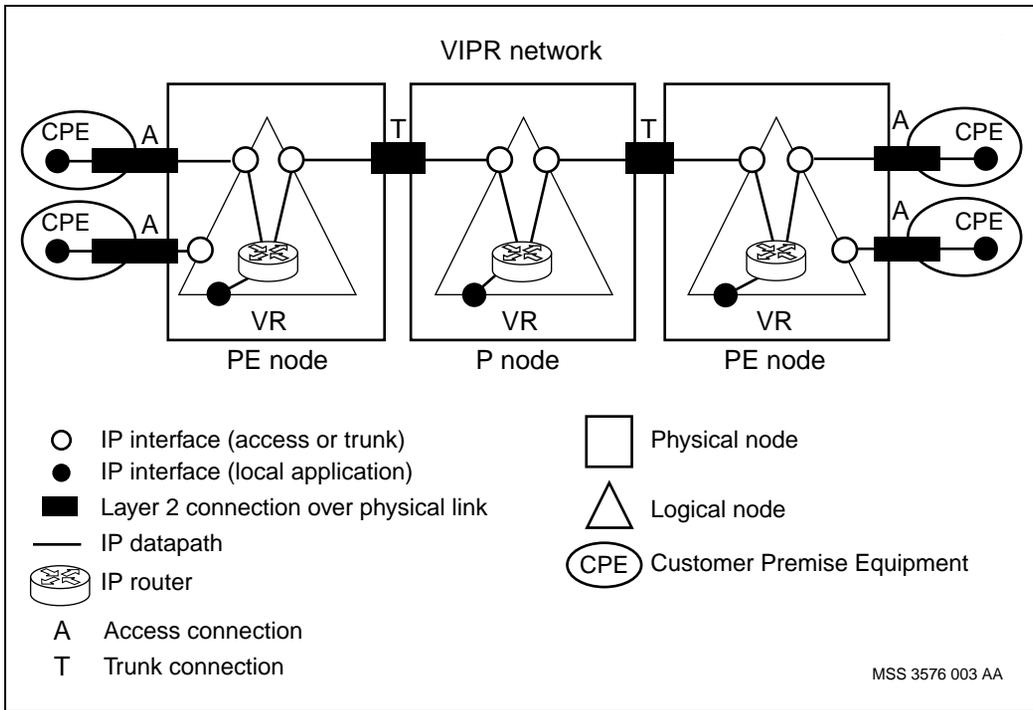
Chapter 6

Layer 3 traffic management in VIPR networks

Use this section to learn about the special VIPR considerations that may affect the planning and implementation of traffic management on your network.

An example of a VIPR network is depicted by Figure 19, “VIPR network,” (page 108).

Figure 19
VIPR network



Layer 3 traffic management on the Multiservice Switch for VIPR networks can be DiffServ based or Class of Service (CoS) based.

Navigation

- “DiffServ based layer 3 traffic management in VIPR networks” (page 109)
- “CoS based layer 3 traffic management in VIPR networks” (page 116)

DiffServ based layer 3 traffic management in VIPR networks

DiffServ based layer 3 traffic management in VIPR networks involves one data path.

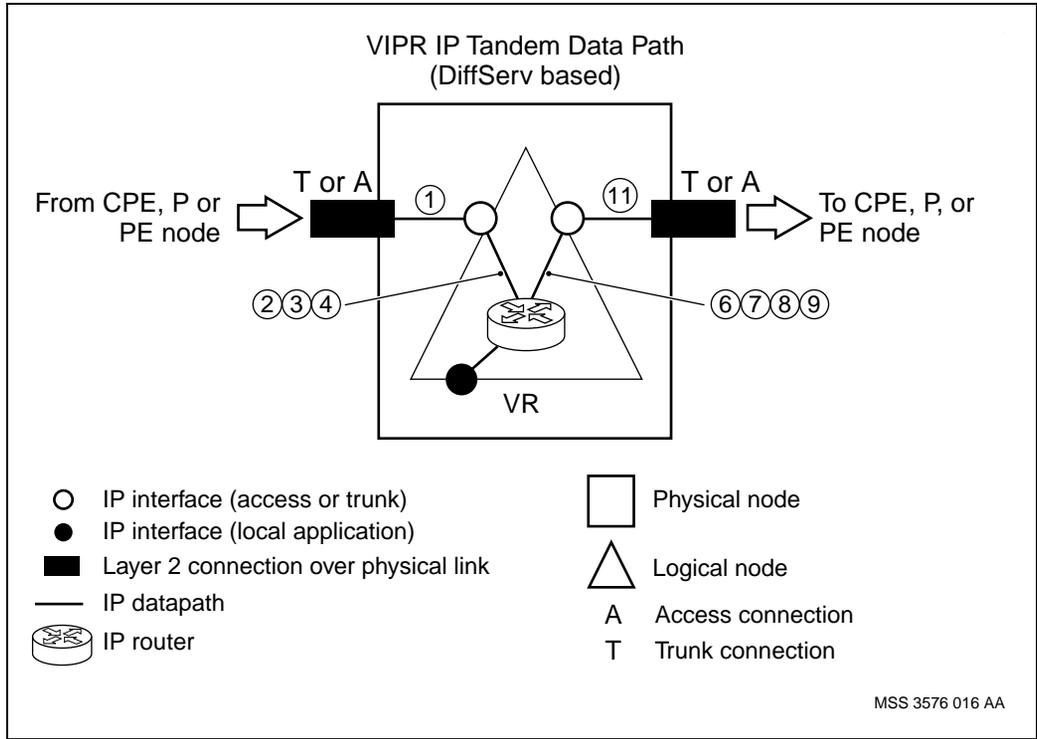
- VIPR node: IP tandem data path

More traffic management options are supported on this data path if the virtual router `vpnMode` attribute is set to `customer`.

VIPR node: IP tandem data path, DiffSev based

Table 13, “VIPR node: IP tandem data path, DiffSev based,” (page 112) provides information pertaining to the IP tunnel entry data path at the PE node, for DiffServ based networks. The table column “Index” pertains to the information markers in Figure 20, “VIPR node: IP tandem data path, DiffSev based,” (page 110).

Figure 20
VIPR node: IP tandem data path, DiffServ based



Legend

Context	Index	Capability
Ingress Interface	1	Ingress L2TM
	2	Ingress IP Firewall
	3	Ingress IP Classification and Marking
	4	Ingress IP Policing
Node	5	IP Routing
Egress Interface	6	Egress IP Policing
	7	Select PHB for Congestion Control
	8	Egress IP Classification and Marking
	9	Multi Link Congestion Control
	10	Single Link Congestion Control
	11	Egress L2TM

Table 13
VIPR node: IP tandem data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Interface	1	Ingress L2TM	ATM Ethernet Frame Relay PPP	Ingress layer 2 traffic management	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	2	Ingress IP Firewall		Map the source and destination fields in the IP header to a permit or deny action according to the IP firewall configuration used by the ingress interface. If the action is deny then drop the packet.	Optional behavior on PQC12 and PQC6 ingress access FPs	Optional behavior on PQC12 and PQC6 ingress access FP.	Ingress FP requires ipFilter
	3	Ingress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			MF Map	Map selected fields in IP, TCP, and UDP headers to a PHB value in the Vr DiffServ domain according to the DiffServProfile MfMap components used by the ingress interface. Mark the corresponding DSCP value into the IP header.	Optional behavior on PQC12 and PQC6 ingress access FPs	Optional behavior on ingress access FPs	Ingress FP requires ipDiffServ
	(Sheet 1 of 4)						

Table 13 (continued)
VIPR node: IP tandem data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Interface (contd)	3 (contd)	Ingress IP Classification and Marking (contd)	DSCP Map	Map the IP DSCP field to a PHB value in the Vr DiffServ domain according to the DiffServProfile DscpMap component used by the ingress interface. Mark the corresponding DSCP value into the IP header.	Optional behavior on ingress access FPs	Optional behavior on ingress access FPs	Ingress FP requires ipDiffServ
			Link Map	Map the ipCos value of the ingress link to a PHB value in the Vr DiffServ domain according to the DiffServProfile LinkMap component used by the ingress interface. Mark the corresponding DSCP value into the IP header.	Optional behavior on ingress access FPs	Optional behavior on ingress access FPs	Ingress FP requires ipDiffServ
	4	Ingress IP Policing		Map the IP DSCP field to a PHB value in the Vr DiffServ domain. Use the PHB value to select a meter according to the PolicerProfile Meter components used by the ingress interface. Police the IP packet according to the meter configuration and the rate and burst size of the traffic through the meter. If the packet is out of profile with respect to committed information rate and committed burst size then map the PHB to a new PHB and mark the corresponding DSCP value into the IP header. If the packet is out of profile with respect to excess burst size then drop the packet.	Optional behavior on PQC12 and PQC6 ingress access FPs if the ingress media is ATM, Frame Relay, or PPP	Optional behavior on PQC6 ingress access FPs if the ingress media is ATM, Frame Relay, or PPP	Ingress FP requires ipPolicing
Node	5	IP Routing		Route the IP packet toward the next IP hop.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
(Sheet 2 of 4)							

Table 13 (continued)
VIPR node: IP tandem data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Interface	6	Egress IP Policing		Map the IP DSCP field to a PHB value in the Vr DiffServ domain. Use the PHB value to select a meter according to the PolicerProfile Meter components used by the egress interface. Police the IP packet according to the meter configuration and the rate and burst size of the traffic through the meter. If the packet is out of profile with respect to committed information rate and committed burst size then map the PHB to a new PHB and mark the corresponding DSCP value into the IP header. If the new and former PHB values are identical then tag the IP packet with "high" drop precedence. If the packet is out of profile with respect to excess burst size then drop the packet.	Optional behavior on PQC12 and PQC6 egress access FPs if the egress media is Frame Relay. Cannot be used in combination with egress IP classification and marking with DSCP map	Optional behavior on PQC6 egress access FPs if the egress media is Frame Relay. Cannot be used in combination with egress IP classification and marking with DSCP map	Egress FP requires ipPolicing
	7	Select PHB for Congestion Control	DSCP-based	Map the IP DSCP field to a PHB value in the Vr DiffServ domain	Default behavior	Default behavior	None
	8	Egress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			DSCP Map	Map the IP DSCP field to a PHB value according to the DiffServProfile DscpMap component used by the egress interface. Mark the corresponding DSCP value into the IP header.	Optional behavior on egress access FPs. Cannot be used in combination with egress IP Policing	Optional behavior on egress access FPs. Cannot be used in combination with egress IP Policing	Ingress PQC12 and PQC6 FPs require ipDiffServ. Egress FQM and GQM FPs require ipDiffServ
	9	Multi Link Congestion Control	Select Connection Class	Map the congestion control PHB to a connection class value in the Vr DiffServ domain. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if the egress media is ATM or Frame Relay with VC bundle	Default behavior if the egress media is ATM or Frame Relay with VC bundle	None.

(Sheet 3 of 4)

Table 13 (continued)
VIPR node: IP tandem data path, DiffServ based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Interface (contd)	10	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the Vr DiffServ domain. Pass this value to egress L2 TM.	Default behavior.	Default behavior.	None
			Select Drop Precedence	If the ingress access interface is configured to retain the drop precedence derived at ingress L2 TM then use that value. If the packet was tagged with "high" drop precedence at the egress IP policing stage then use that value. In any other case map the congestion control PHB to a drop precedence value in the Vr DiffServ domain. Pass the drop precedence value to egress L2 TM.	Default behavior.	Default behavior.	None
	11	Egress L2TM	ATM Ethernet Frame Relay PPP Virtual Media	Egress layer 2 traffic management.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOc3PosAtm FP.							
(Sheet 4 of 4)							

CoS based layer 3 traffic management in VIPR networks

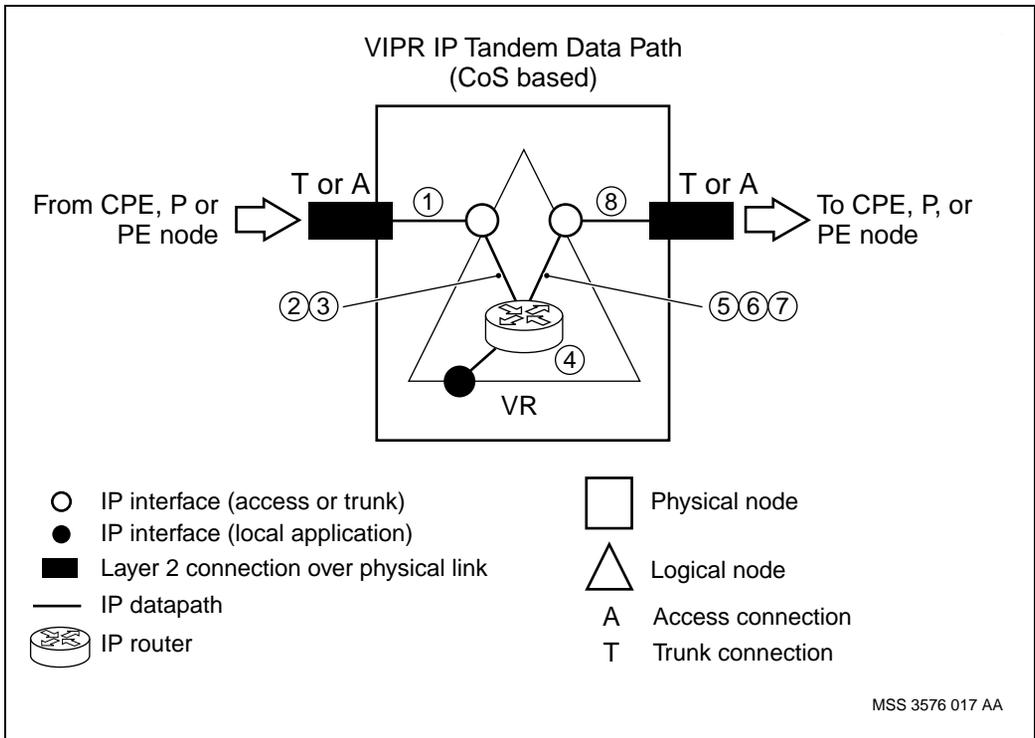
Cos based layer 3 traffic management in VIPR networks involves one data path, which is described in the section that follows.

More traffic management options are supported on this data path if the virtual router `vpnMode` attribute is set to `customer`.

VIPR node: IP tandem data path, CoS based

Table 14, “VIPR node: IP tandem data path, CoS based,” (page 119) provides information pertaining to the IP tunnel entry data path at the PE node, for CoS based networks. The table column “Index” pertains to the information markers in Figure 21, “VIPR node: IP tandem data path, CoS based,” (page 117).

Figure 21
VIPR node: IP tandem data path, CoS based



Legend

Context	Index	Capability
Ingress Interface	1	Ingress L2TM
	2	Ingress IP Firewall
	3	Ingress IP Classification
Node	4	IP Routing
Egress Interface	5	Egress IP Marking
	6	Multi Link Congestion Control
	7	Single Link Congestion Control
	8	Egress L2TM

Table 14
VIPR node: IP tandem data path, CoS based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Interface	1	Ingress L2TM	ATM Ethernet Frame Relay PPP	Ingress layer 2 traffic management	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	2	Ingress IP Firewall		Map the source and destination fields in the IP header to a permit or deny action according to the IP firewall configuration used by the ingress interface. If the action is deny then drop the packet.	Optional behavior on iPQC12 and PQC6 ingress access FPs	Optional behavior on iPQC12 and PQC6 ingress access FPs	Ingress FP requires ipFilter
	3	Ingress IP Classification	MF Map	Map selected fields in IP, TCP, and UDP headers to a CoS value according to the CosPolicyGroup Policy components used by the ingress interface. Tag the IP packet with this value.	Optional behavior on PQC12 and PQC6 ingress access FPs	Optional behavior on ingress access FPs	Ingress FP requires ipCos
			DSCP Map	Map the IP DSCP field to a CoS value according to the CosPolicyGroup Policy components used by the ingress interface. Tag the IP packet with this value.	Optional behavior on PQC12 and PQC6 ingress access FPs	Optional behavior on ingress access FPs	Ingress FP requires ipCos
			Link Map	Identify the ipCos value of the ingress link. Tag the IP packet with this value.	Optional behavior	Optional behavior	Ingress FP requires ipDiffServ
	Node	4	IP Routing		Route the IP packet toward the next IP hop.	Not in L3TM scope	Not in L3TM scope
(Sheet 1 of 3)							

Table 14 (continued)
VIPR node: IP tandem data path, CoS based

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Interface	5	Egress IP Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Default behavior	None
			Modify DSCP	Map the tagged CoS value to a DSCP value according to the CosPolicyGroup Ect components used by the egress interface. Mark the DSCP value into the IP header.	Optional behavior on PQC12 and PQC6 egress access FPs if the ingress FP is PQC12 or PQC6	Optional behavior on egress access FPs	Ingress FP requires ipCos
	6	Multi Link Congestion Control	Select Connection Class	Use the tagged CoS value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if the egress media is ATM or Frame Relay with VC bundle	Default behavior if the egress media is ATM or Frame Relay with VC bundle	None
7	Single Link Congestion Control	Select Scheduling Class		Pass scheduling class value 0 to egress L2 TM.	Default behavior except on FQM and GQM egress FPs if the ingress FP is FQM or GQM	Default behavior	None
				Map the tagged CoS value to a scheduling class value based according to the following relation: CoS 0,1,2,3 maps to scheduling class 0,2,4,5. Pass this value to egress L2 TM.	Default behavior on FQM and GQM egress FPs if the ingress FP is FQM or GQM	Not supported.	None
			Select Drop Precedence	Pass the drop precedence derived by ingress L2 TM to egress L2 TM.	Default behavior	Default behavior	None
(Sheet 2 of 3)							

Table 14 (continued)**VIPR node: IP tandem data path, CoS based**

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Egress Interface (contd)	7 (contd)	Single Link Congestion Control (contd)	Select Drop Precedence (contd)	Map the tagged CoS value to a drop precedence value according to the CosPolicyGroup lct components used by the ingress interface. Pass the drop precedence value to egress L2 TM.	Optional behavior if the ingress FP is PQC12 or PQC6	Optional behavior	None
	8	Egress L2TM	ATM Ethernet Frame Relay PPP Virtual Media	Egress layer 2 traffic management.	Not in L3 TM scope	Not in L3 TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOc3PosAtm FP.							
(Sheet 3 of 3)							

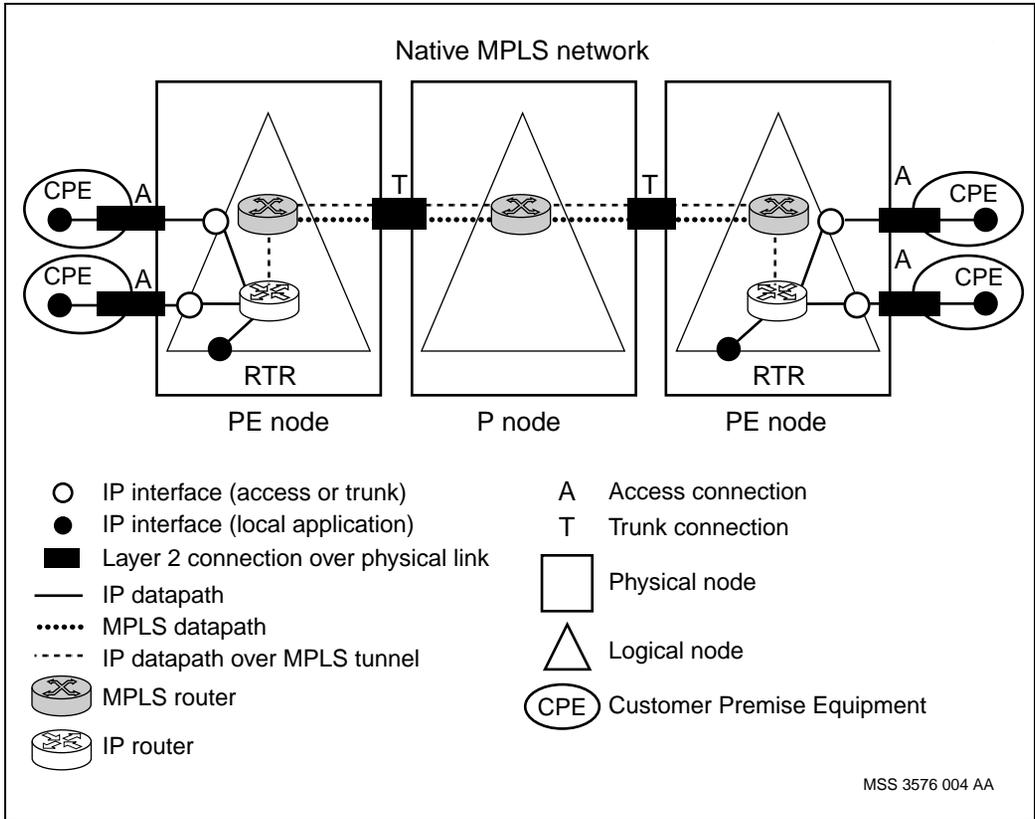
Chapter 7

Layer 3 traffic management in native MPLS networks

Use this section to learn about the special native MPLS considerations that may affect the planning and implementation of traffic management on your network.

A Layer 3 native MPLS network is depicted by Figure 22, “Native MPLS network,” (page 124).

Figure 22
Native MPLS network



Layer 3 traffic management on the Multiservice Switch for native MPLS networks is DiffServ based, and supports the native MPLS PE node (label edge router) only. The native MPLS P node (label switched router) is deployed on other vendor equipment.

DiffServ based layer 3 traffic management in native MPLS networks involves three different data path, as described in the sections that follow.

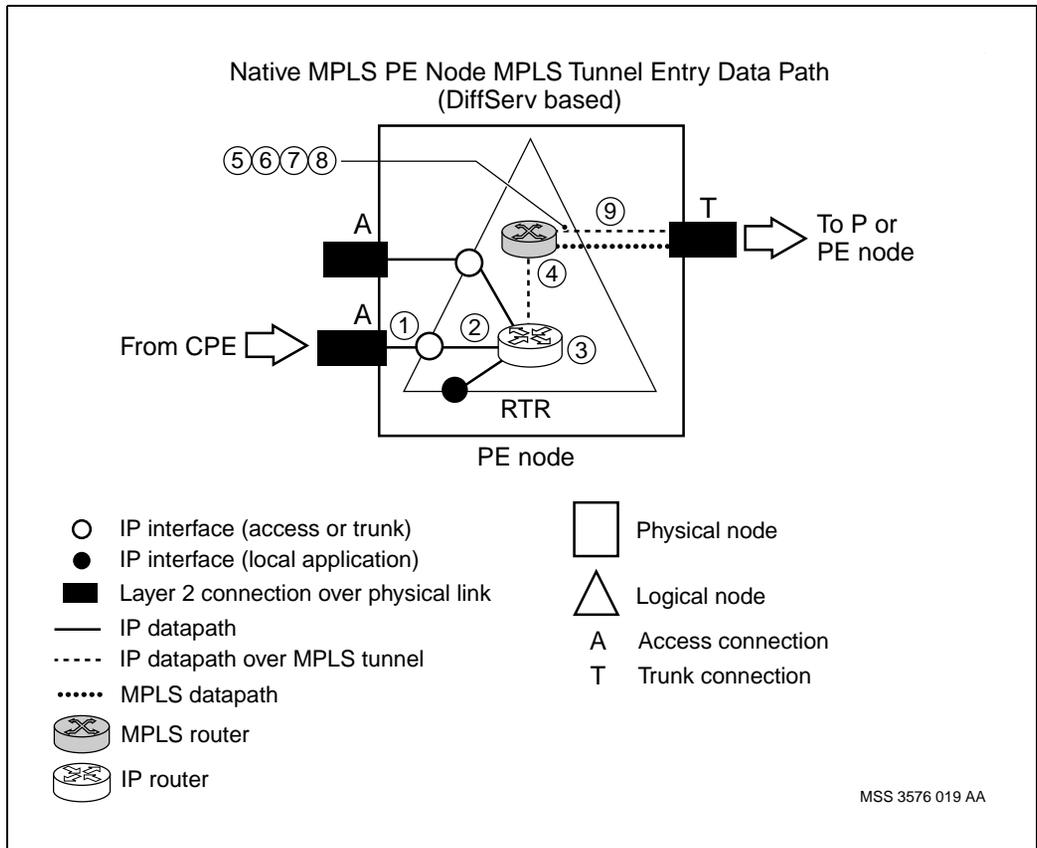
Navigation

- “Native MPLS PE node: IP tunnel entry data path” (page 126)
- “Native MPLS PE node: MPLS tunnel exit data path” (page 130)
- “Native MPLS PE node: IP tandem data path on the Rtr” (page 133)

Native MPLS PE node: IP tunnel entry data path

Table 15, “Native MPLS PE node: IP tunnel entry data path,” (page 128) provides information pertaining to the MPLS tunnel entry data path at the PE node. The table column “Index” pertains to the information markers in Figure 23, “Native MPLS PE node: IP tunnel entry data path,” (page 126).

Figure 23
Native MPLS PE node: IP tunnel entry data path



Legend

Context	Index	Capability
Ingress Access Interface	1	Ingress L2TM
	2	Ingress IP Classification and Marking
Node	3	IP Routing
	4	MPLS Routing
Egress Trunk Interface	5	MPLS EXP Marking
	6	Select PHB for Congestion Control
	7	Multi Link Congestion Control
	8	Single Link Congestion Control
	9	Egress L2TM

Table 15
Native MPLS PE node: IP tunnel entry data path

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Access Interface	1	Ingress L2TM	ATM	Ingress layer 2 traffic management	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	2	Ingress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior.	Not supported	none
Node	3	IP Routing		Route the IP packet toward the next IP hop.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	4	MPLS Routing		Encapsulate the IP packet with an MPLS transport label. Route the MPLS packet toward the next MPLS hop.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
Egress Trunk Interface	5	MPLS EXP Marking	Eight way marking	Map the IP DSCP field to a PHB value in the Rtr DiffServ domain. Map the PHB value to an EXP value in the Rtr DiffServ domain. Mark the EXP value into the MPLS transport label.	Default behavior	Not supported	None
	6	Select PHB for Congestion Control	DSCP-based	Map the IP DSCP field to a PHB value in the Rtr DiffServ domain.	Default behavior.	Not supported	None
	7	Multi Link Congestion Control	Select Connection Class	Map the congestion control PHB to a connection class value in the Rtr DiffServ domain. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Not supported	Not supported	None
	8	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the Rtr DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Not supported	None
			Select Drop Precedence	Map the congestion control PHB to a drop precedence value in the Rtr DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Not supported	None
9	Egress L2TM	Ethernet	Egress layer 2 traffic management.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope	
(Sheet 1 of 2)							

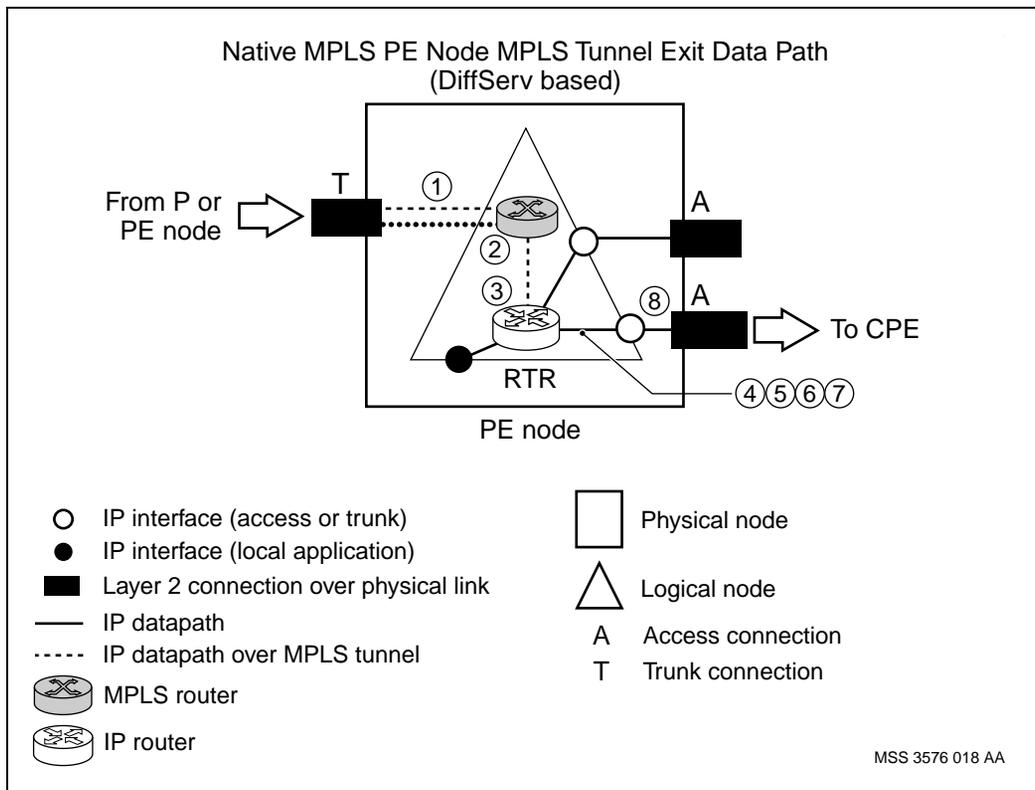
Table 15 (continued)
Native MPLS PE node: IP tunnel entry data path

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOc3PosAtm FP.							
(Sheet 2 of 2)							

Native MPLS PE node: MPLS tunnel exit data path

Table 16, “Native MPLS PE node: MPLS tunnel exit data path,” (page 132) provides information pertaining to the MPLS tunnel exit data path at the PE node. The table column “Index” pertains to the information markers in Figure 24, “Native MPLS PE node: MPLS tunnel exit data path,” (page 130).

Figure 24
Native MPLS PE node: MPLS tunnel exit data path



Legend

Context	Index	Capability
Ingress Access Interface	1	Ingress L2TM
Node	2	MPLS Routing
	3	IP Routing
Egress Access Interface	4	Select PHB for Congestion Control
	5	Egress IP Classification and Marking
	6	Multi Link Congestion Control
	7	Single Link Congestion Control
	8	Egress L2TM

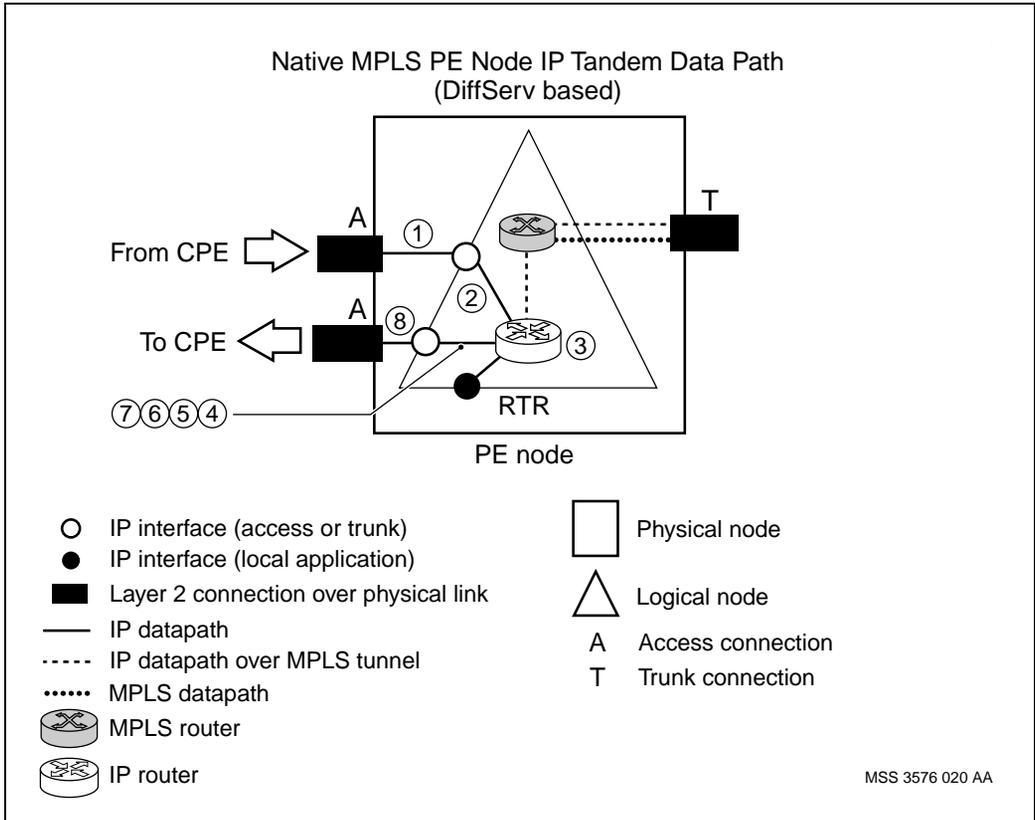
Table 16
Native MPLS PE node: MPLS tunnel exit data path

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Access Interface	1	Ingress L2TM	Ethernet	Ingress layer 2 traffic management	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
Node	2	MPLS Routing		Remove the MPLS transport label.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	3	IP Routing		Route the IP packet toward the next IP hop.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
Egress Access Interface	4	Select PHB for Congestion Control	DSCP-based in Vrf DiffServ domain	Map the IP DSCP field to a PHB value in the Rtr DiffServ domain.	Default behavior	Not supported	None
	5	Egress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior.	Not supported	None
	6	Multi Link Congestion Control	Select Connection Class	Map the congestion control PHB to a connection class value in the Rtr DiffServ domain. Use the connection class value to select one of up to four layer 2 connections to the next IP hop.	Default behavior if egress media is ATM with VC bundle	Not supported	None
	7	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the Rtr DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Not supported	None
			Select Drop Precedence	Map the congestion control PHB to a drop precedence value in the Rtr DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Not supported	None
	8	Egress L2TM	ATM	Egress layer 2 traffic management.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOc3PosAtm FP.							

Native MPLS PE node: IP tandem data path on the Rtr

Table 17, “Native MPLS PE node: IP tandem data path on the Rtr,” (page 135) provides information pertaining to the IP tandem data path on the Rtr, at the PE node. The table column “Index” pertains to the information markers in Figure 25, “Native MPLS PE node: IP tandem data path on the Rtr,” (page 133).

Figure 25
Native MPLS PE node: IP tandem data path on the Rtr



Legend

Context	Index	Capability
Ingress Interface	1	Ingress L2TM
	2	Ingress IP Classification and Marking
Node	3	IP Routing
Egress Interface	4	Select PHB for Congestion Control
	5	Egress IP Classification and Marking
	6	Multi Link Congestion Control
	7	Single Link Congestion Control
	8	Egress L2TM

Table 17
Native MPLS PE node: IP tandem data path on the Rtr

Context	Index	Capability	Option	Behavior	20k/15K	7K	Feature List
Ingress Interface	1	Ingress L2TM	ATM Ethernet	Ingress layer 2 traffic management	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
	2	Ingress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Not supported	None
Node	3	IP Routing		Route the IP packet toward the next IP hop.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
Egress Interface	4	Select PHB for Congestion Control	DSCP-based	Map the IP DSCP field to a PHB value in the Rtr DiffServ domain.	Default behavior	Not supported	None
	5	Egress IP Classification and Marking	Preserve DSCP	The IP DSCP field in the IP header is not modified.	Default behavior	Not supported	None
	6	Multi Link Congestion Control	Select Connection Class	Map the congestion control PHB to a scheduling class value in the Rtr DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Not supported	None
	7	Single Link Congestion Control	Select Scheduling Class	Map the congestion control PHB to a scheduling class value in the Rtr DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Not supported	None
			Select Drop Precedence	Map the congestion control PHB to a drop precedence value in the Rtr DiffServ domain. Pass this value to egress L2 TM.	Default behavior	Not supported	None
	8	Egress L2TM	ATM Ethernet	Egress layer 2 traffic management.	Not in L3TM scope	Not in L3TM scope	Not in L3 TM scope
Comments: "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2). "PQC12" indicates PQC version 12 (also known as PQC12v1). "FQM" indicates 4pGigE FP. "GQM" indicates 16pOc3PosAtm FP.							

Chapter 8

Layer 3 traffic management for local packets

Use this section to learn about Layer 3 traffic management for packets that are generated at a local source or are terminated at a local destination.

Navigation

- DiffServ based Layer 3 traffic management for IP packets generated at a local source
- DiffServ based Layer 3 traffic management for IP packets terminated at a local destination
- Cos based Layer 3 traffic management for IP packets generated at a local source
- Cos based Layer 3 traffic management for IP packets terminated at a local destination

DiffServ based Layer 3 traffic management for IP packets generated at a local source

Ingress Layer 2 and Layer 3 traffic management functions are not performed on IP packets that are generated at local source on the *VirtualRouter (Vr)*, *VpnRouteForwarder*, (*Vrf*) or *Router (Rtr)*.

The DSCP field of a locally generated IP packet used for network routing control is initialized to the DSCP value that corresponds to the *phbRoutingSource (phbR)* attribute of the *DifferentiatedServicesDomain* component. The default value of this attribute is CS6 (48). The DSCP field of any other locally generated packet is initialized to the DSCP value that

corresponds to the *phbGeneralSource (phbG)* attribute of the *DifferentiatedServicesDomain* component. The default value of this attribute is DF (0).

For a locally generated ICMP echo request packet, the initial value of the DSCP field is overridden by the *-tos* or *-dscp* option of the *ping* command. For a locally generated ICMP echo reply packet, the initial value of the DSCP field is overridden by the value contained in the IP header of the received request packet.

After IP routing, each locally generated IP packet is treated according to the value of its DSCP field, just as if it had been an IP packet received at an ingress interface.

DiffServ based Layer 3 traffic management for IP packets terminated at a local destination

Egress Layer 2 and Layer 3 traffic management functions are not performed on IP packets that are terminated at a local destination on the *VirtualRouter (Vr)*, *VpnRouteForwarder (Vrf)*, or *Router (Rtr)*.

Cos based Layer 3 traffic management for IP packets generated at a local source

Ingress Layer 2 and Layer 3 traffic management functions are not performed on IP packets that are generated at a local source on the *VirtualRouter (Vr)*.

The tagged Class Of Service (CoS) value for a locally generated IP packet used for network routing control is initialized to 3. For any other locally generated IP packet, the tagged CoS value is initialized to 0. For a locally generated ICMP echo request packet, the initial tagged CoS value is overridden by the *-cos* option of the “ping” command. For a locally generated ICMP echo reply packet, the initial tagged CoS value is replaced with the tagged CoS value of the received request packet.

The tagged drop precedence (DP) value for a locally generated IP packet used for network routing control is initialized to low. For any other locally generated IP packet, the tagged DP value is initialized to “medium”.

The DSCP field of a locally generated packet used for network control is initialized to 48 (CS6). For any other locally generated IP packet, the DSCP field is initialized to 0 (DF). For a locally generated ICMP echo request packet, the initial value of the DSCP field is overridden by the -tos or -dscp option of the “ping” command. For a locally generated ICMP echo reply packet, the initial value of the DSCP field is replaced by the value contained in the IP header of the received request packet.

After IP routing, each locally generated IP packet is treated according to its tagged CoS and drop precedence values, just as if it had been an IP packet received at an ingress interface.

Cos based Layer 3 traffic management for IP packets terminated at a local destination

Egress Layer 2 and Layer 3 traffic management functions are not performed on IP packets that are terminated at a local destination the *VirtualRouter (Vr)*.

Chapter 9

Layer 2 to Layer 3 interactions

Use this section to learn about how Layer 3 traffic management interacts with Layer 2 traffic management.

Navigation

- Layer 3 interactions with Layer 2 overview
- Layer 3 interactions with ATM
- Layer 3 interactions with Ethernet
- Layer 3 interactions with Frame Relay (FR)
- Layer 3 interactions with Point-to-Point Protocol (PPP)

Layer 3 interactions with Layer 2 overview

IP and MPLS packets are forwarded on Layer 2 connections between nodes. The supported media for the Layer 2 connections can vary depending on network topology.

Each Layer 2 media offers different traffic management capabilities and capabilities can vary depending on functional processor (FP) type.

Layer 3 interactions with ATM

Table 18, “Interactions with ATM at ingress,” (page 142) explains what Layer 2 treatments are applied before a packet is treated at Layer 3. Table 19, “Interactions with ATM at egress,” (page 143) explains what Layer 2 treatments are applied after a packet is treated at Layer 3.

Table 18
Interactions with ATM at ingress

Capability	Option	GQM	PQC
Ingress L2 Policing	Police ATM cells according to the rate and burst size of the traffic on the ingress connection. If the cell is out-of-profile with respect to the policer configuration for the connection then remark the ATM cell loss priority (CLP) bit to 1 or drop the cell. The receive packet-wise discard option should be enabled to ensure that if one L2 cell gets dropped for being out-of-profile then the entire L3 packet is dropped.	Optional behavior	Optional behavior
Ingress DP tagging	Map the ATM cell loss priority (CLP) bit to a discard priority (DP) value. Pass the DP value to L3 TM. The default mapping for CBR and rtVBR is CLP 0, 1 to DP low, high, respectively. The default mapping for nrtVBR is CLP 0, 1 to DP medium, high, respectively. The default mapping for UBR is CLP 0 and CLP 1 to DP high.	Default behavior	Default behavior
Comments: "PQC" indicates PQC12 or PQC6. "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2) "PQC12 indicates PQC12 version 12 (also known as PQC12v1) "GQM" indicates 16pOC3PosAtm FP.			

Table 19
Interactions with ATM at egress

Capability	Option	GQM	PQC
Egress VCC Selection	Use the connection class value assigned by L3 TM to select one of up to four outgoing Layer 2 connections.	Default behavior with LLC encapsulation and VC bundle	Default behavior with LLC encapsulation and VC bundle
Egress CLP bit marking	Map the drop precedence (DP) value assigned by L3 TM to 0 or 1. Mark the ATM cell loss priority (CLP) bit with this value. The default mapping for CBR, rtVBR, nrtVBR is DP low and medium to CLP 0, DP high to CLP 1. The default mapping for UBR is DP low, medium, high to CLP 1.	Default behavior	Default behavior on some data paths (see comments below).
	Mark the ATM cell loss priority (CLP) bit to 0 regardless of the drop precedence (DP) value assigned by L3 TM.	Not supported	Default behavior on some data paths (see comments below).
Egress Queue Selection	Map the ATM service category of the egress connection to a L2 emission priority (EP) value. Use the EP value to select an outgoing queue. The default mapping is CBR to EP 0, rtVbr to EP 1, rtVbr to EP4, UBR to EP 7	Default behavior	Default behavior
Egress Drop or Forward	Use the drop precedence assigned by L3 TM to determine whether a cell should be enqueued or dropped. The transmit packet-wise discard option should be enabled to ensure that if one L2 cell gets dropped due to queue congestion then the entire L3 packet is dropped.	Default behavior	Default behavior on FPs with AQM
	Use the ATM CLP bit to determine whether cell should be enqueued or dropped. The transmit packet-wise discard option should be enabled to ensure that if one L2 cell gets dropped due to queue congestion then the entire L3 packet is dropped.	Not supported	Default behavior on FPs with APC
Egress Scheduling Algorithm	Strict Priority	Default behavior for EP 0 and EP 1	Default behavior for EP 0 and EP 1
	Weighted Fair	Default behavior for EP 2 to EP 7	Default behavior for EP 2 to EP 7
(Sheet 1 of 2)			

Table 19 (continued)
Interactions with ATM at egress

Capability	Option	GQM	PQC
Egress Shaping	Restrict the egress traffic to a specified maximum rate and maximum burst size.	Optional behavior	Optional behavior
Comments: “PQC” indicates PQC12 or PQC6. “PQC6” indicates PQC version 6 (also known as PQC2 or PQC6v2) “PQC12” indicates PQC12 version 12 (also known as PQC12v1) “GQM” indicates 16pOC3PosAtm FP. The ATM CLP bit is always marked to 0 on the following data paths: RFC2764 network on IP optimized tunnel exit data path with PQC12 or PQC6 access FP RFC2547 network on MPLS tunnel entry data path with PQC6 trunk FP RFC2547 network on MPLS tunnel exit data path with PQC12 or PQC6 access FP VIPR network with bridge termination encapsulation type on egress PQC12 or PQC6FP			
(Sheet 2 of 2)			

Interaction of ATM end-to-end loopback and connection class

ATM OAM cells used to manage the loopback have a discard priority that is not necessarily lower than the IP or MPLS data traffic that is transmitted over the ATM connection and the ATM connection gets congested then the OAM cells can get dropped.

If ATM end-to-end loopback is enabled then locking one ATM connection at one end will allow the traffic to move automatically to the ATM connection that is configured with the "next best" ipCos value. When this happens the cycle repeats until all ATM connections are down. Everything is fine if the ATM connections are never congested.

If ATM end-to-end loopback is disabled and the ATM connection that carries IP control packets is locked at one end only, then the L3 connections over the ATM connection will be lost. The ATM connection must be manually locked at both ends for traffic to move to the ATM connection that is configured with the "next best" ipCos value.

Layer 3 interactions with Ethernet

Table 20, “Interactions with Ethernet at ingress,” (page 145) explains what Layer 2 treatments are applied before a packet is treated at Layer 3. Table 21, “Interactions with Ethernet at egress,” (page 146) explains what Layer 2 treatments are applied after a packet is treated at Layer 3.

Table 20
Interactions with Ethernet at ingress

Capability	Option	GQM	PQC
Ingress DP tagging	Pass the drop precedence (DP) value "high" to L3 TM.	Default behavior	Default behavior
Comments: "PQC" indicates PQC12 or PQC6. "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2) "PQC12 indicates PQC12 version 12 (also known as PQC12v1) "FQM" indicates 4pGigE FP.			

Table 21
Interactions with Ethernet at egress

Capability	Option	GQM	PQC
Egress Queue Selection	Map the scheduling class for eight queues (SC8Q) value assigned by L3 TM a L2 emission priority (EP) value. Use the EP value to select an outgoing queue. The default mapping is SC 0,1,2,3,4,5,6,7 to EP 7,6,5,4,3,2,1,0.	Default behavior	Not supported
	Map the scheduling class for four queues (SC4Q) value assigned by L3 TM a L2 emission priority (EP) value. Use the EP value to select an outgoing queue.	Not supported	Default behavior
	Map the scheduling class for two queues (SC2Q) value assigned by L3 TM a L2 emission priority (EP) value. Use the EP value to select an outgoing queue.	Not supported	Not supported
Egress Drop or Forward	Use the drop precedence assigned by L3 TM to determine whether a frame should be enqueued or dropped.	Default behavior	Default behavior
Egress Scheduling Algorithm	Strict Priority	Default behavior for EP 0 and 1	Not supported
	Weighted Fair	Default behavior for EP 2 to EP 7	Default behavior
Comments: "PQC" indicates PQC12 or PQC6. "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2) "PQC12 indicates PQC12 version 12 (also known as PQC12v1) "FQM" indicates 4pGigE FP.			

Layer 3 interactions with Frame Relay (FR)

Table 22, “Interactions with Frame Relay (FR) at ingress,” (page 147) explains what Layer 2 treatments are applied before a packet is treated at Layer 3. Table 23, “,” (page 148) explains what Layer 2 treatments are applied after a packet is treated at Layer 3.

Table 22
Interactions with Frame Relay (FR) at ingress

Capability	Option	GQM	PQC
Ingress L2 Policing	Police FR frames according to the rate and burst size of the traffic on the ingress connection. If the frame is out-of-profile with respect to the policer configuration for the connection then remark the FR discard eligible (DE) bit to 1 or drop the frame.	Optional behavior	Optional behavior
Ingress DP tagging	Map the FR discard eligible (DE) bit to a discard priority (DP) value. Pass the DP value to L3 TM.	Default behavior	Default behavior
Comments: “PQC” indicates PQC12 or PQC6. “PQC6” indicates PQC version 6 (also known as PQC2 or PQC6v2) “PQC12” indicates PQC12 version 12 (also known as PQC12v1)			

Table 23
Interactions with Frame Relay (FR) at egress

Capability	Option	GQM	PQC
Egress DLCI Selection	Use the connection class value assigned by L3 TM to select one of up to four outgoing Layer 2 connections.	Default behavior with VC bundle	Default behavior with VC bundle
Egress DEbit marking	Map the drop precedence (DP) value assigned by L3 TM to 0 or 1. Mark the FR discard eligible (DE) bit with this value.	Default behavior	Default behavior
Egress Queue Selection	Map the scheduling class for four queues (SC4Q) value assigned by L3 TM a L2 emission priority (EP) value. Use the EP value to select an outgoing queue.	Default behavior	Not supported
	Map the scheduling class for two queues (SC2Q) value assigned by L3 TM a L2 emission priority (EP) value. Use the EP value to select an outgoing queue.	Not supported	Default behavior
Egress Drop or Forward	Use the drop precedence assigned by L3 TM to determine whether a frame should be enqueued or dropped.	Default behavior	Default behavior
Egress Scheduling Algorithm	Strict Priority	Not supported	Default Behavior
	Weighted Fair	Default behavior	Not supported
Comments: "PQC" indicates PQC12 or PQC6. "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2) "PQC12" indicates PQC12 version 12 (also known as PQC12v1)			

Layer 3 interactions with Point-to-Point Protocol (PPP)

Table 24, “Interactions with Point-to-Point Protocol (PPP) at ingress,” (page 149) explains what Layer 2 treatments are applied before a packet is treated at Layer 3. Table 25, “Interactions with Point-to-Point Protocol (PPP) at egress,” (page 150) explains what Layer 2 treatments are applied after a packet is treated at Layer 3.

Table 24
Interactions with Point-to-Point Protocol (PPP) at ingress

Capability	Option	GQM	PQC
Ingress DP tagging	Pass the drop precedence (DP) value "high" to L3 TM.	Default behavior	Default behavior
Comments: "PQC" indicates PQC12 or PQC6. "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2) "PQC12" indicates PQC12 version 12 (also known as PQC12v1)			

Table 25
Interactions with Point-to-Point Protocol (PPP) at egress

Capability	Option	GQM	PQC
Egress Queue Selection	Map the scheduling class for four queues (SC4Q) value assigned by L3 TM a L2 emission priority (EP) value. Use the EP value to select an outgoing queue.	Default behavior	Not supported
	Map the scheduling class for two queues (SC2Q) value assigned by L3 TM a L2 emission priority (EP) value. Use the EP value to select an outgoing queue.	Not supported	Default behavior
Egress Scheduling Algorithm	Strict Priority	Not supported	Default Behavior
	Weighted Fair	Default behavior	Not supported
Comments: "PQC" indicates PQC12 or PQC6. "PQC6" indicates PQC version 6 (also known as PQC2 or PQC6v2) "PQC12" indicates PQC12 version 12 (also known as PQC12v1)			

Chapter 10

IP classification and marking

Use this section to learn about the types of IP packet classification and marking methods available on a Nortel Networks Multiservice Switch.

Navigation

- “Ingress IP packet classification and marking” (page 151)
- “Egress IP packet classification and marking” (page 154)

Ingress IP packet classification and marking

The DSCP field of IP packets that are received at an interface is modified according to ingress classification and marking policies that are configured at that interface. The PHB assigned to the packet is determined by the DSCP value of the IP packet after ingress classification and marking. If the DSCP matches a PHB included in the Differentiated Service domain, then that PHB is applied; otherwise the PHB for best effort ("df" of default forwarding) is applied. Downstream devices may also identify and categorize traffic based on the new DSCP value.

Differentiated Services code point (DSCP) classification

The DSCP classification is performed using provisioned policy information, that is, using the *DscpMap* component under *DiffServProfile*. This provisioned information is a simple mapping that serves to associate the DSCP in an IP packet with a PHB index.

The attribute *assignedPhb* specifies the PHB assigned to the IP packet when the packet DSCP matches one of DSCPs in the list. This attribute is a vector index.

Multi-field (MF) classification

The Multi-field classification is provisioned using the *MultiFieldMap* component under *DiffServProfile*. This provisioned information includes a set of rules that serve to associate the information in the IP and Layer 4 headers of an IP packet with a PHB index.

The *assignedPhb* attribute specifies the PHB assigned to the IP packet when the packet field matches one of the set of rules. The attribute can have one of several values.

A rule includes a PHB index, and the following information:

- an IP address and prefix length (number of bits of address which are significant)
- layer 4 protocol (one of ICMP, TCP, or UDP)
- a single or range of TCP/UDP ports

MF classification is only performed on a packet if the IP header protocol field is one of ICMP, TCP, or UDP. The IP address and port information is only used if the protocol field in the header is either TCP or UDP.

For ICMP, a single PHB value is applied to all packets regardless of the addressing information in the packet. Therefore, the IP address and port range provisioned in the policies are ignored when the protocol is provisioned as ICMP.

When provisioning MF classification, the port value of *any* indicates that any possible value in the source and destination TCP/UDP port fields is a match. Matches with any IP address are indicated by setting the prefix length to 0, or setting the IP address to the value of 0.0.0.0.

When performing MF classification, IP DiffServ first compares the destination IP address and destination TCP/UDP port information from the packet against the provisioned MF classification rules. If no match is found then the source IP address and source TCP/UDP port information is searched for. This approach gives higher precedence to the destination IP address and port over the source IP address and port.

Most protocols follow a client-server paradigm. In most cases, the server IP address and/or TCP/UDP port is known. In other cases, there can be many client IP addresses with a random port value that is dynamically established during initial connection setup. Therefore, it is not always possible to specify a MF classification rule based on client addressing. Checking the destination addressing and then source addressing allows a single flow classification rule to handle both traffic to and from a server.

When several MF classification rules apply to a packet, the more specific rule is used. For example, if one MF classification rule applies to a range of IP addresses and another rule applies to a specific IP address within that range, then the one which is specific to the IP address is used. Precedence is always given to the MF classification rule that applies to the destination address over the source address. For example, an MF classification rule that applies to any IP address (in other words the destination address) has precedence over the one which applies to the source address, even if the source address is a 32-bit exact match.

IP DiffServ gives precedence to the IP address over the TCP/UDP port value. For instance, if one rule specifies a range of IP addresses, which includes the IP address from the packet, and another rule specifies the exact TCP/UDP port value from the packet, the first rule is matched.

Multifield classification is supported on PQC and SBIC interfaces only.

Link classification

For the Link classification, the DSCP field of the packet received on the *Vr Ip IpPort* or *Rtr Vrf Interface* component is set according to the DSCP provisioned under the *LinkMap assignedPhb* attribute. This attribute specifies the PHB assigned to the IP packet when the packet Dscp matches one of the connection class values in the list. This attribute is a vector index.

Packet marking

The DSCP field of the IP packet is marked according to on the PHB assigned to the packet by any of the ingress classification methods listed above. If no classification method is configured at the ingress interface then the DSCP field of the packet is preserved.

Egress IP packet classification and marking

The DSCP field of IP packets that are transmitted at an interface is modified according to egress classification and marking policies that are configured at that interface. The PHB assigned to the packet is determined by the DSCP value of the IP packet prior to egress classification and marking. In other words, egress classification and marking does not alter the PHB assigned to packet, but downstream devices may identify and categorize traffic based on the new DSCP value.

Differentiated Services code point (DSCP) classification

The DSCP classification is performed using provisioned policy information, that is, using the *DscpMap* component under *DiffServProfile*. This provisioned information is a simple mapping that serves to associate the DSCP in an IP packet with a PHB index.

The attribute *assignedPhb* specifies the PHB assigned to the IP packet when the packet DSCP matches one of DSCPs in the list. This attribute is a vector index.

Packet marking

The DSCP field of the IP packet is marked according to on the PHB assigned to the packet by any of the egress classification methods listed above. If no classification method is configured at the egress interface then the DSCP field of the packet is preserved.

Chapter 11

IP firewalls

Use this section to learn about conceptual information pertaining to the implementation of Nortel Networks Multiservice Switch IP flow filters.

For information about configuring IP flow filters, see NN10600-591 *Nortel Networks Multiservice Switch 7400/15000/20000 Layer 3 Traffic Management Configuration*.

Navigation

- Overview of IP flow filters
- IP flow filter definition

Overview of IP flow filters

IP flow filters help you to create a secure network by allowing you to decide which IP packet flows will be permitted or denied entry to the network. Permit or deny actions are based on the most specific match to the IP packet flow's source IP address, destination IP address, or both the source and destination IP addresses.

If permitted, packets are processed and forwarded to their destination. If denied, packets are dropped. If packets do not match any flow of the assigned filter, they are dropped.

When the source IP address, destination IP address, or both the source and destination IP addresses match more than one of the specified flows under a filter, the match with the most precise source address is given precedence. If

the source address specification is selected, and the destination address matches more than one of the specified flows under a filter, the match with the most precise destination address is given precedence.

Permit and deny actions for IP packet flows based on source address are illustrated in Figure 26, “Filtering based on source address,” (page 156).

Permit and deny actions for IP packet flows based on destination address are illustrated in Figure 27, “Filtering based on destination address,” (page 157).

Figure 26
Filtering based on source address

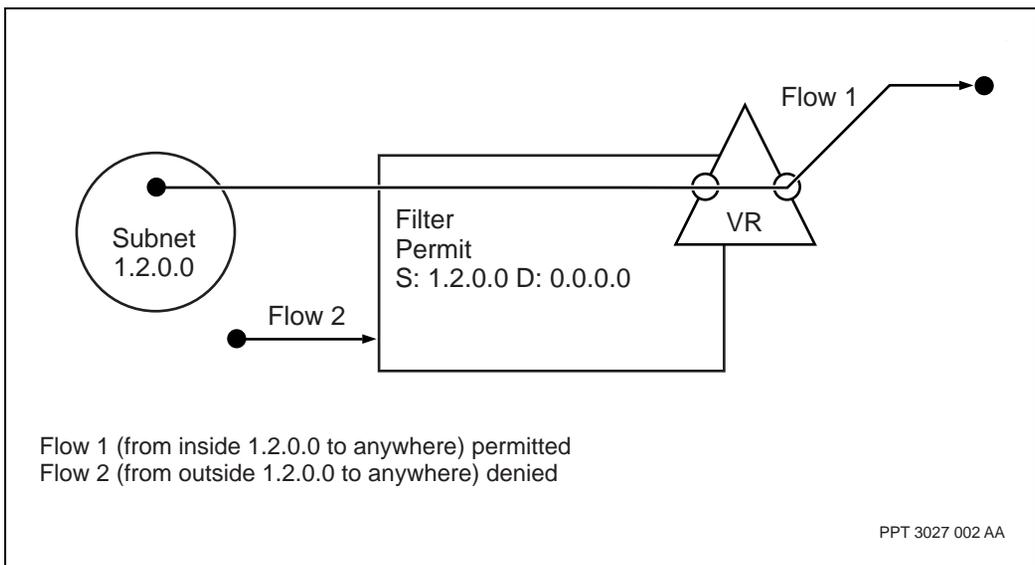
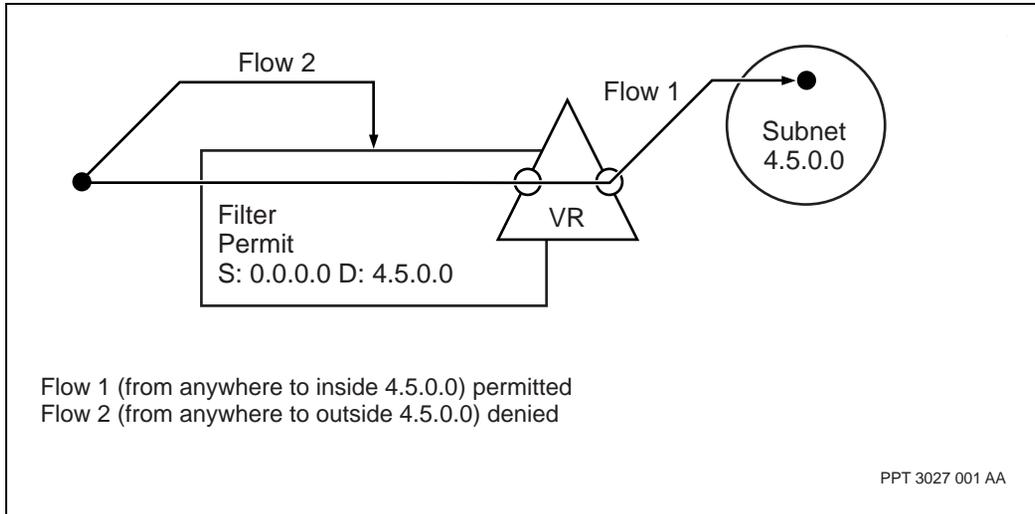


Figure 27
Filtering based on destination address



You can configure IP flow filters under the *Ip* subcomponent of a VR. After the flow filter is created for a VR, you can assign the filter to any of the protocol ports on that VR. To assign a flow filter to a protocol port on a VR, link the filter to the *IpPort* component under the protocol port. To assign a flow filter to all of the protocol ports on a VR, link the filter to the *Ip* component under the VR. The filter assignment for a protocol port will override the filter assignment for a VR. A filter created under one VR cannot be assigned to another VR or to the protocol ports of another VR.

IP flow filtering actions will not be applied to any protocol port that is linked to IP tunnel media or to any protocol port assigned to a VR that is linked to IP tunnel media. Since no filtering actions are applied to protocol ports linked to IP tunnel media, all IP packets are forwarded in this situation.

IP flow filters are only supported on PQC2-based cards.

IP flow filter definition

The *Filter* subcomponent of the *Ip* component allows you to add, delete, and modify IP flow filters. A filter is assigned to a router by linking the corresponding *Filter* subcomponent to the *Ip* component under the same

router. A filter is assigned to a protocol port by linking the corresponding *Filter* subcomponent to the *IpPort* subcomponent of the desired *Protocol Port* component, under the same router. If no filter is assigned, all packets are permitted entry to the network.

A filter associated with a protocol port is applied to all ingress IP traffic, but not egress traffic, at that port. Filtering is performed prior to making the forwarding decision that determines the egress port where the packet will be transmitted. If no filter component is associated with a port (because there is neither an assignment to the *Vr Ip* component, nor an assignment to the *Vr Pp IpPort* component) all IP packet flows are admitted at that port.

The *FilterFlow* subcomponent of the *Filter* subcomponent is used to define the filter action (whether an IP flow is permitted or denied) that is required for an individual IP flow. A *Filter* subcomponent must have at least one *FilterFlow* subcomponent whose filter action is permit.

Two filter flow subcomponents should not be configured under one filter component or subcomponent that have equivalent source and destination address specifications.

IP flow filters can block IP packets that are used by routing protocols. When configuring IP flow filters, it is important to configure IP flow filters that permit or deny the IP control packets.

Chapter 12

IP policing

Use this section to learn about conceptual information pertaining to the implementation of Multiservice Switch IP policing services.

For information about configuring IP policing, see NN10600-801 *Nortel Networks Multiservice Switch 7400/15000/20000 IP Configuration Management*.

Navigation

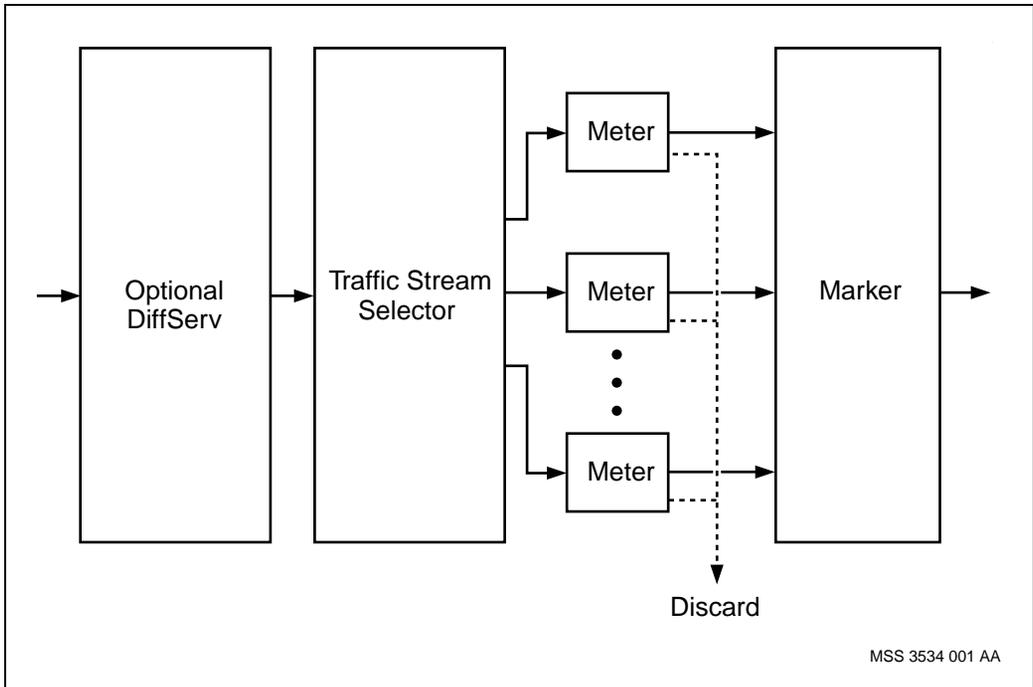
- “Overview of IP policing services” (page 159)
- “Policing services benefits” (page 161)
- “Policing process” (page 161)
- “Traffic meter and policing algorithm” (page 162)
- “Policer types” (page 166)

Overview of IP policing services

This feature manages incoming and outgoing traffic rate on individual IP router’s access interfaces, based on IP packet information. This allows a network provider to monitor traffic flow coming into the interface and either discard the packet or increase drop precedence for packets violating the service agreement.

IP traffic policing can be applied to the traffic received or to the traffic transmitted on the interface. Different traffic profiles can be applied in an ingress and egress direction. Figure 28, “Ingress/Egress Traffic Processing,” (page 160) indicates how IP Policing fits with other stages of packet processing for ingress and egress direction.

Figure 28
Ingress/Egress Traffic Processing



This feature is closely related to the DiffServ feature, in that, when activated, it forms an integrated part of the IP Differentiated Services Traffic Management process.

Policing services benefits

The Multiservice Switch IP policer enhances and complements the IP VPN and DiffServ offerings. The policer enables the service provider to measure customer traffic volume, and enforce compliance to the defined traffic contract. The service provider benefits in several ways:

- The service provider is given the ability of offering new and enhanced types of IP services with different service level agreements (SLAs).
- The policer enables the provider to restrict the amount of traffic admitted to the network from specific customers or classes, which results in overall performance improvement and fairness to all network users. The excess traffic can be tagged as low priority and forwarded as best effort or discarded, at user option.
- The policer provides comprehensive per Meter DiffServ class statistics, which can be used for monitoring SLAs, and for network engineering purposes.
- Traffic policing can be applied in ingress and egress directions. The egress policing mechanism provides an additional tool for IP VPN deployment, where traffic from multiple access points may be directed to an egress link. Furthermore, the combination of egress policing and scheduling enables the provider to offer bandwidth guarantees to the various traffic classes on the provider-to-customer access link.

Policing process

There are four components that make up the policing process:

- “Traffic stream selection” (page 161)
- “Traffic metering” (page 162)
- “Policing actions” (page 162)
- “Packet marking” (page 162)

Traffic stream selection

Traffic stream is selected based on the PHB associated with the packet. Packets PHB is derived from the packets DS code point according to the routers DiffServ domain. Based on the PHB, packet is steered to the appropriate traffic meter.

Note: Prior to applying the policing function, a packet's DS code point can be modified by the DiffServ feature.

Traffic metering

Traffic meter measures the traffic stream properties against the traffic rates specified by the traffic profile associated with the Meter. The main purpose of the meter is to identify in-profile and out-of-profile packets. There can be up to two levels of conformance specified per traffic meter.

Policing actions

Traffic that is identified as an out-of-profile can be subjected to different actions. It can be dropped or it can be assigned a new PHB. All in-profile traffic is forwarded.

Packet marking

Packets which were not discarded by the Policer are directed to the Marker. The Marker simply marks the packets DS code point based on the current PHB of the packet according to the routers DiffServ domain.

Note: The egress packet marking functionality, implemented as part of the DiffServ feature, is disabled when the egress policing is enabled on an interface.

Traffic meter and policing algorithm

The Multiservice Switch IP traffic policer applies the Frame Relay traffic management algorithm. This algorithm uses the parameters described in the following section.

Traffic meter and policing algorithm parameters

The Frame Relay traffic management algorithm uses following parameters:

- Committed Information Rate (CIR) - This is the rate (in bits/s) at which the network agrees to transfer information under normal condition. The rate is measured over the measurement interval T.
- Committed Burst Size (Bc) - The maximum amount of data (in bits) that a network agrees to transfer under normal conditions over a measurement interval T.

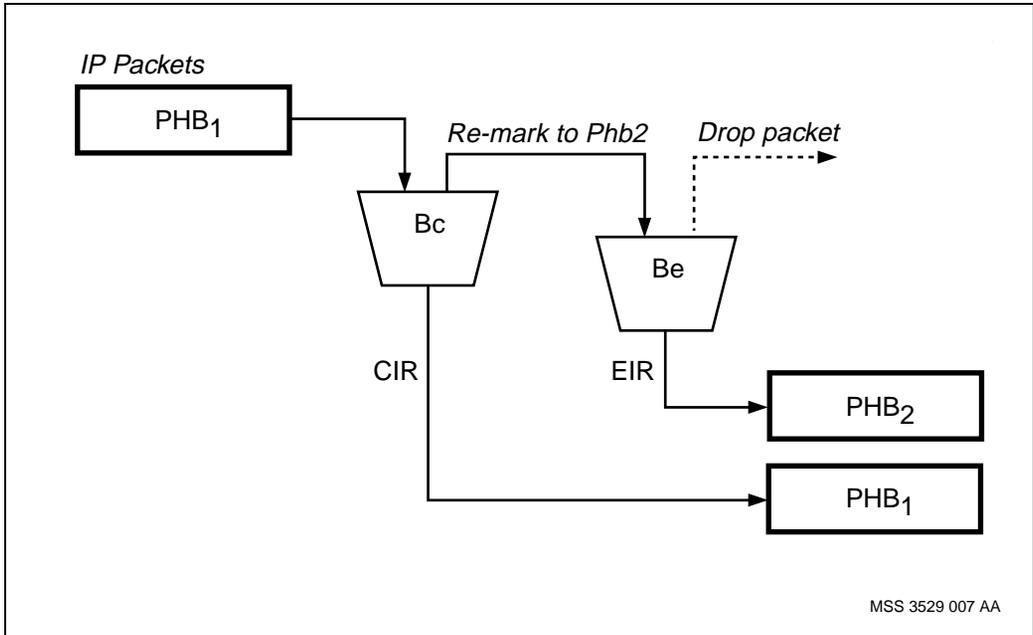
- **Excess Information Rate (EIR)** - This is the rate (in bits/s) at which the network will attempt to transfer information. The rate is measured over the measurement interval T.
- **Excess Burst Size (Be)** - The maximum amount of data (in bits) that the network will attempt to deliver over a measurement interval T.
- **Measurement interval (T)** - The time interval over which rates and burst sizes are measured. It is computed as $T=Bc/CIR$ or if $CIR=0$ it has to be specified explicitly.

Leaky bucket mechanism

The Frame Relay traffic management algorithm utilizes the leaky bucket mechanism. Two leaky buckets are used. The first bucket controls CIR and the second bucket controls EIR.

The general operation of the Multiservice Switch IP traffic policer utilizing the leaky buckets is illustrated by Figure 29, “Leaky bucket mechanism,” (page 164).

Figure 29
Leaky bucket mechanism



The bit count limit for the CIR bucket is set to Committed Burst Size (B_c) and the bit count limit for the EIR bucket is set to Excess Burst Size (B_e). The initial bit count for each bucket is set to 0.

When a packet arrives at the CIR bucket and the bucket is not filled (the current bit count in the bucket is less than the bit count limit), the packet adds n bits ($n = \text{packet size} * 8$, which includes IP header and a payload, but not L2 header) to the bit count and the packet is forwarded. If the CIR bucket was full at the time when the packet arrived, then the packet is optionally assigned a new PHB and is directed to the EIR bucket.

When a packet arrives at the EIR bucket and the bucket is not filled (the current bit count in the bucket is less than the bit count limit), the packet adds n bits ($n = \text{packet size} * 8$, which includes IP header and a payload, but not L2 header) to the bit count and the packet is forwarded. If the EIR bucket was full at the time when a packet arrived, the packet is dropped.

The buckets are leaking at the CIR and EIR rates respectively, that is, the bit count in the buckets is decremented at the CIR and EIR rate.

The user configured parameters are provided to specify the traffic profile as well as algorithm operations. Depending on the selected configuration one or both rates can be enforced and different actions can be taken for out-of-profile packets.

Egress policer drop precedence

The egress IP policer provides SLA rate enforcement like the ingress IP policer. It can also be used to provide minimum bandwidth guarantees on functional processors where absolute priority based scheduling systems are implemented, as discussed later in “EF minimum rate guaranteed policer” (page 171).

A PHB for forwarding IP packets is associated with a traffic class and a drop precedence according to the provisioned attributes of the *PerHopBehavior* subcomponents of the *DiffServDomain* component. A PHB is assigned to an IP packet based on the DSCP field of the IP packet header after ingress IP classification, after ingress IP policing, and after egress IP policing, as applicable, according to the data path configuration.

Normally, the drop precedence assigned to an IP packet forwarded at an egress interface is determined by the last PHB assigned to the packet, after those traffic management functions are performed. However, if egress IP policing is configured and the following two conditions are true, then the drop precedence assigned to the packet is 'high'.

- the packet contributes to the committed information rate of the egress meter
- the out-of-profile PHB assigned to the packet by the egress IP policer is identical to the PHB of the packet prior to egress IP policing

Policing statistics

Policing statistics related to a policer operation are collected when the policing is activated on an interface. The statistics are available under interface in a *MeterStatistics* component. The instance of the *MeterStatistics* component corresponds to the *Meter* component instance under *PolicerProfile* associated with the interface. Meter statistics spooling is controlled by setting the spooling attribute in the *IngressServices* or *EgressServices* components under interface.

Policer types

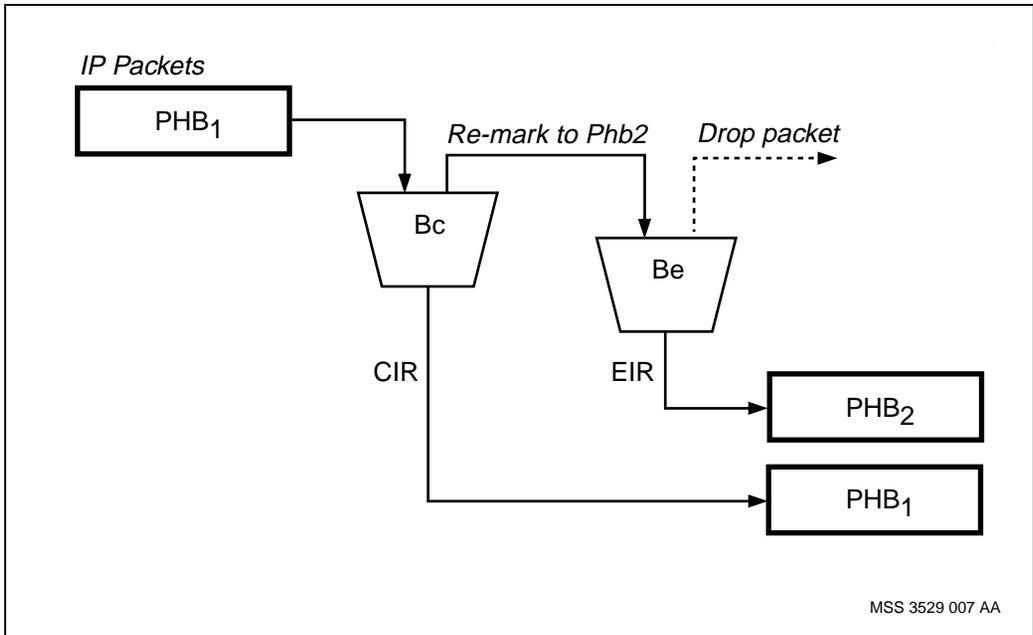
By adjusting policer parameters, different policer types and applications can be configured. The most common are described below. For provisioning procedures and examples of these policer types, refer to NN10600-591 *Nortel Networks Multiservice Switch 7400/15000/20000 Layer 3 Traffic Management Configuration*. For the parameter settings for these policer types, refer to Table 26, “Policer parameters,” (page 173).

- “Dual rate policer” (page 167)
- “Single CIR policer with all excess packets dropped” (page 168)
- “Single CIR policer with all excess forwarded” (page 169)
- “Single EIR rate policer” (page 170)
- “Color blind and color aware policer” (page 170)
- “Combined traffic class policer” (page 171)
- “EF minimum rate guaranteed policer” (page 171)
- “Policer as a DiffServ statistics collector” (page 172)

Dual rate policer

The dual rate policer enforces both rates. First packets are directed to the CIR bucket, conforming traffic is forwarded, excess traffic is re-marked and directed to the EIR bucket. At EIR bucket non-conforming traffic is discarded and conforming traffic is forwarded.

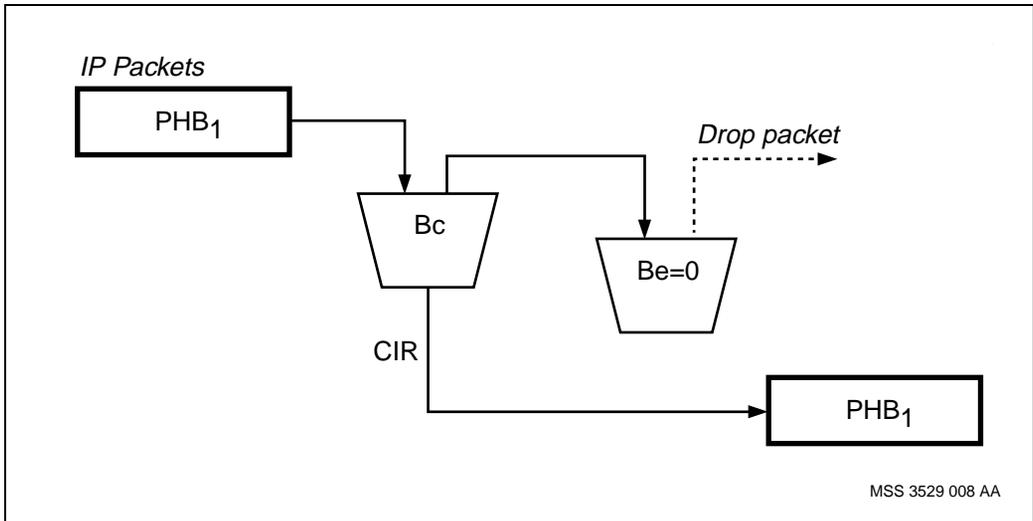
Figure 30
Dual rate policer



Single CIR policer with all excess packets dropped

Single CIR rate policer enforces only CIR rates. All packets are directed to the CIR bucket. Out-of-profile packets are directed to the EIR bucket. Since EIR bucket's Be is zero, all excess packets are dropped.

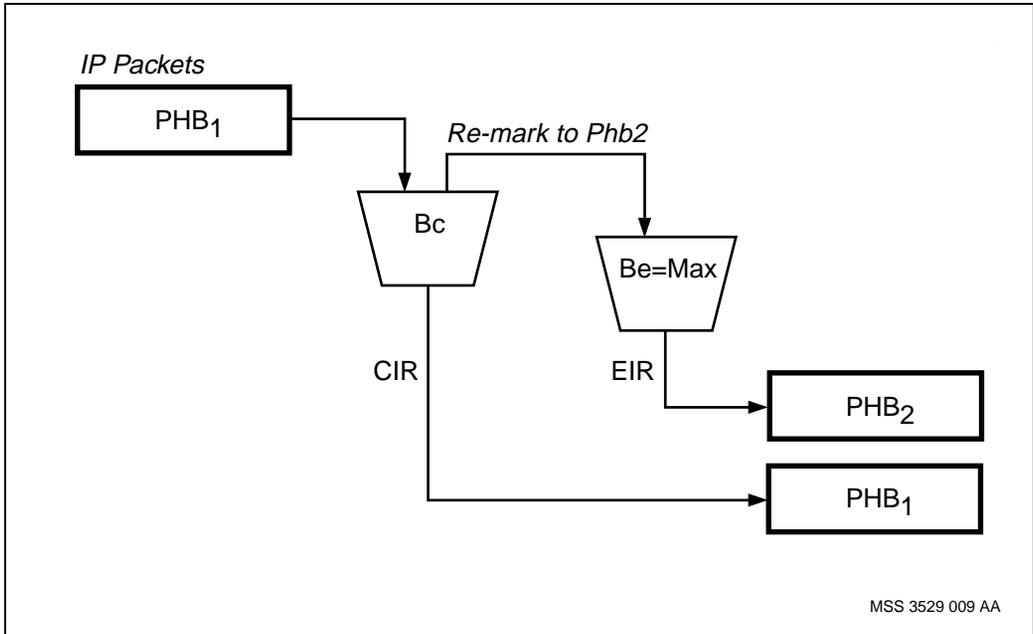
Figure 31
Single CIR policer with all excess dropped



Single CIR policer with all excess forwarded

Single CIR rate policer enforces only CIR rates. All packets are directed to the CIR bucket. Out-of-profile packets are directed to the EIR bucket. Since EIR bucket's Be is set to the maximum value, all excess packets are forwarded.

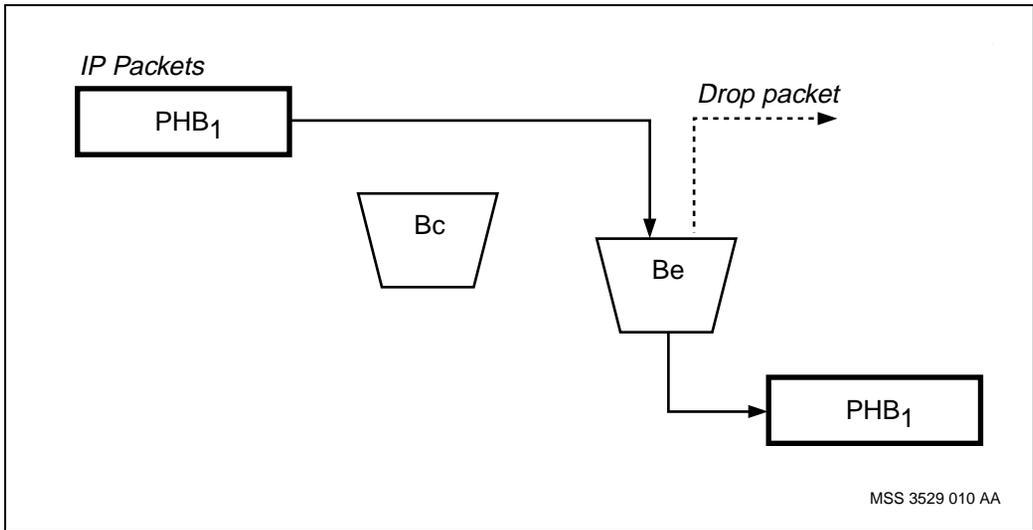
Figure 32
Single CIR policer with all excess forwarded



Single EIR rate policer

Single EIR rate policer enforces only EIR rate. All traffic is directed to the EIR bucket. Out of profile packets are dropped.

Figure 33
Single EIR rate



Color blind and color aware policer

In general the policer operates in a color aware mode. In a color aware mode the PHB of a packet is determined based on the content of a packet by one of the classification methods, and the meter/policer functionality is applied based on that PHB.

In a color blind mode, it is assumed that all packets arriving on an interface belong to the same PHB, and the policer sets the PHB of the packet based on the conformance to the traffic profile.

To achieve a color blind mode, all packets arriving at the interface are classified to the same PHB by using the link mode classification method provided by the DiffServ feature.

For the color blind mode, policer is usually setup in a dual rate mode, however any of the single rate types can be used as well.

Combined traffic class policer

The policing functionality is totally independent of the forwarding treatment. Packets belonging to different traffic classes can be policed by the same policer. One of the applications is a virtual pipe when all packets are directed to the same policer, but when forwarded they will be serviced by different queues.

EF minimum rate guaranteed policer

As mentioned in the section “Egress policer drop precedence” (page 165), the egress IP policer can be used to provide minimum bandwidth guarantees for IP packets that are not forwarded on absolute priority queues. The following example demonstrates this functionality.

Assume IP packets with various DSCP values are transmitted at an ATM egress interface and only those packets that are assigned PHB=EF are transmitted on the CBR link that uses an absolute priority queue. While the volume of EF traffic is greater than the line rate of the egress interface, the absolute priority queue will never be emptied and all the other queues will be starved. Non-EF traffic and excess EF traffic will not get forwarded under these conditions.

Starvation of the other queues can be prevented by using the egress IP policer to limit the EF traffic to maximum committed information rate that is less than the line rate of the egress interface. If the egress IP policer is configured so that the excess burst size for EF traffic is 0, then all EF traffic in excess of the CIR is dropped. More efficient use of the absolute priority queue is possible if the excess burst size for EF is set to the maximum possible value. Then, all EF traffic above the EF CIR rate is assigned ‘high’ discard priority and is forwarded, rather than dropped, when there is unused bandwidth available.

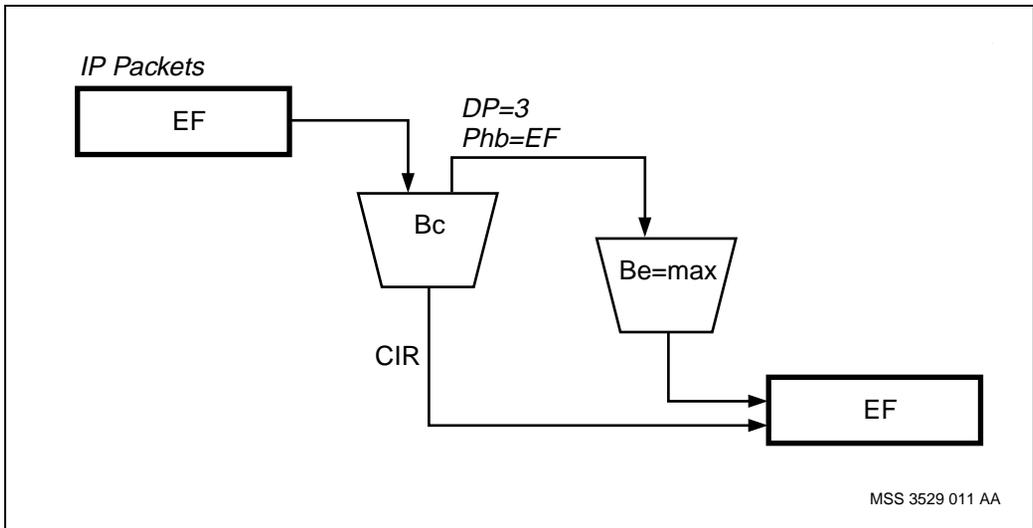
The row “EF minimum rate guaranteed policer” (page 174) of Table 26, “Policer parameters,” (page 173) shows the parameter settings for this functionality.

Figure 34, “EF minimum rate guaranteed,” (page 172) illustrates operation of a minimum bandwidth guarantee policer for EF PHB. The traffic with the code point EF is directed to the CIR bucket. The conforming traffic is forwarded with the EF DS code point and the drop precedence assigned according to the *DiffServDomain* for EF PHB. All non-conforming traffic is

forwarded to the EIR bucket. Since the excess burst size (Be) is set to maximum, all excess traffic is forwarded with the EF DS code, but the drop precedence of the packet is set to the highest value.

Note: This functionality is only applicable for egress policer. It does not apply to ingress policer.

Figure 34
EF minimum rate guaranteed



Policer as a DiffServ statistics collector

The special application for the policing functionality on an interface is to collect traffic statistics on a per Phb or set of Phbs basis, without rate enforcement. This is achieved by configuring the *Meter* component for each Phb or set of Phbs, and setting the allowed burst sizes for the *Meter* to the maximum value. The row “Per Phb statistics parameters” (page 175) of Table 26, “Policer parameters,” (page 173) shows the Per Phb statistics parameter settings.

Table 26
Policer parameters

Policer type	Parameter name	Parameter value
Dual rate policer		
Meter Parameters	committedInformationRate	non zero
	committedBurstSize	non zero
	excessBurstSize	non zero
	measurementInterval	n/a
Phb Parameters	initialBucket	CIR
	outOfProfilePhb	newPhb
Single CIR policer with all excess packets dropped		
Meter Parameters	committedInformationRate	non zero
	committedBurstSize	non zero
	excessBurstSize	zero
	measurementInterval	n/a
Phb Parameters	initialBucket	CIR
	outOfProfilePhb	n/a
Single CIR policer with all excess packets forwarded		
Meter Parameters	committedInformationRate	non zero
	committedBurstSize	non zero
	excessBurstSize	maximum allowed value (50000000)
	measurementInterval	n/a
Phb Parameters	initialBucket	CIR
	outOfProfilePhb	newPhb
(Sheet 1 of 3)		

Table 26 (continued)
Policer parameters

Policer type	Parameter name	Parameter value
Single EIR rate policer		
Meter Parameters	committedInformationRate	n/a
	committedBurstSize	n/a
	excessBurstSize	non zero
	measurementInterval	non zero
Phb Parameters	initialBucket	EIR
	outOfProfilePhb	n/a
EF minimum rate guaranteed policer		
Meter Parameters	committedInformationRate	guaranteedRate
	committedBurstSize	guaranteedBurst
	excessBurstSize	50000000
	measurementInterval	n/a
EF Phb Parameters	initialBucket	CIR
	outOfProfilePhb	EF
(Sheet 2 of 3)		

Table 26 (continued)
Policer parameters

Policer type	Parameter name	Parameter value
Per Phb statistics parameters		
Meter Parameters	committedInformationRate	50000000
	committedBurstSize	50000000
	excessBurstSize	50000000
	measurementInterval	n/a
Phb Parameters	initialBucket	CIR
	outOfProfilePhb	default
(Sheet 3 of 3)		

Chapter 13

IP Class of Service (CoS)

Use this section to learn about conceptual information pertaining to the implementation of Nortel Networks Multiservice Switch IP Class of Service (CoS).

For information about configuring IP CoS, see NN10600-801 *Nortel Networks Multiservice Switch 7400/15000/20000 IP Configuration Management*.

Navigation

- “Overview of IP CoS” (page 178)
- “Packet classification at ingress” (page 180)
- “Packet treatment at egress” (page 186)
- “Frame relay DTE class-based forwarding” (page 189)
- “IP-optimized DLCI class-based forwarding” (page 192)
- “ATM MPE class-based forwarding” (page 193)
- “Gigabit Ethernet class-based forwarding” (page 194)
- “Point-to-point protocol class-based forwarding” (page 198)
- “IP CoS over virtual media” (page 198)

Overview of IP CoS

IP CoS is one of two ways differentiated services can be deployed by Nortel Networks Multiservice Switch for IP. Differentiated Services offers several advantages over IP CoS and should be considered as the best choice for differentiated services on Multiservice Switch 15000 and Multiservice Switch 20000 nodes.

Multiservice Switch nodes support differentiation of IP traffic for different levels of service. It can examine layer 2, layer 3, and layer 4 parameters to classify IP packets. Once classified, the node has the option of marking the DiffServ codepoint (DSCP) and forwarding the packets using different qualities of service (QoS) over media that support multiple QoS.

For information on which function processors support IP CoS, see NN10600-551 *Nortel Networks Multiservice Switch 7400/15000/20000 FP Configuration Reference*.

The table “IP CoS support on access media” (page 179) summarizes support for IP CoS over different access media. Layer 3 and layer 4 classification on PQC-based FPs is only on PQC2-based FPs.

Table 27
IP CoS support on access media

	Frame relay		ATM MPE	PPP	Ethernet		
	FR DTE	IP-optimized DLCI			MS3 (GigE)	PQC	SBIC
Packet classification at ingress							
Layer 2 classification (VC, protocol port)	yes	yes	yes	yes	yes	yes	yes
Layer 3 classification ¹ (DSCP-based)	yes	yes	yes	yes	no	yes	yes
Layer 3/4 classification ¹ (flow-based)	yes	yes	yes	yes	no	yes	yes
Packet marking at egress							
DSCP marking ¹ (modify DSCP field)	yes	yes	yes	yes	no	yes	yes
Class-based packet forwarding at egress							
CoS to VC mapping (select virtual connection)	yes	yes	yes	no	no	no	no
CoS to DP mapping ² (apply drop precedence)	yes	yes	yes	no ³	yes	yes	no
CoS to EP mapping ² (apply emission priority)	yes	yes	no ⁴	yes	yes	yes	no
<p>¹ Feature ipCos is required in the feature list of the ingress FP. Ip CoS capability is based on the ingress media, not the egress media.</p> <p>² Feature ipCos is required in the feature list of the ingress FP if CoS to DP mapping or CoS to EP mapping is required and is to be functional.</p> <p>³ CoS to DP mapping on PPP is supported on SBIC FPs.</p> <p>⁴ EP is determined by CoS to VC mapping.</p>							

You can also configure IP CoS functionality on both IP tunneling protocol ports and virtual media protocol ports. For more information about configuring IP CoS functionality on IP tunneling protocol ports and virtual media protocol ports, see NN10600-801 *Nortel Networks Multiservice Switch 7400/15000/20000 IP Configuration Management*.

Packet classification at ingress

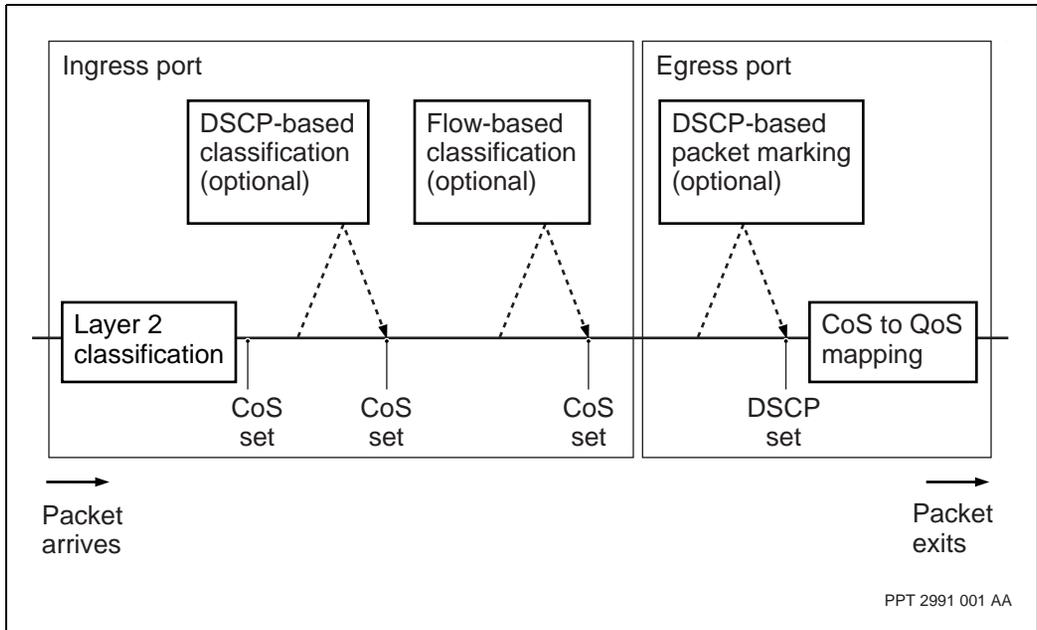
IP CoS classifies IP packets at different points in the transmission path, based on one or more of the following criteria:

- incoming media link
- DiffServ codepoint (DSCP)
- IP address (source or destination), layer 4 (transport) protocol, TCP or UDP port numbers

See the figure “IP CoS assignment on transmission path” (page 181). Each classification supersedes any preceding classification. For more information on how IP CoS manages packet classification, see the following sections:

- “Layer 2 classification” (page 181)
- “DSCP-based classification” (page 182)
- “Flow-based classification” (page 184)
- “IP CoS policies” (page 185)

Figure 35
IP CoS assignment on transmission path



Layer 2 classification

IP CoS performs layer 2 classification on each packet as it enters the Nortel Networks Multiservice Switch node. The packet's first assigned CoS value is based on the layer 2 link on which the packet arrives. Classification can be VC-based or port-based.

For access media that support VCs (frame relay and ATM MPE), the CoS value assigned to packets corresponds to the CoS value associated with the incoming connection through the appropriate *ipcos* attribute. See "Frame relay DTE class-based forwarding" (page 189) and "ATM MPE class-based forwarding" (page 193) for more information.

For other access media capable of bearing IP traffic (PPP and Ethernet), the CoS value assigned to packets arriving on a particular port corresponds to the value configured under the *Vr Pp IpPort ipCoS* attribute.

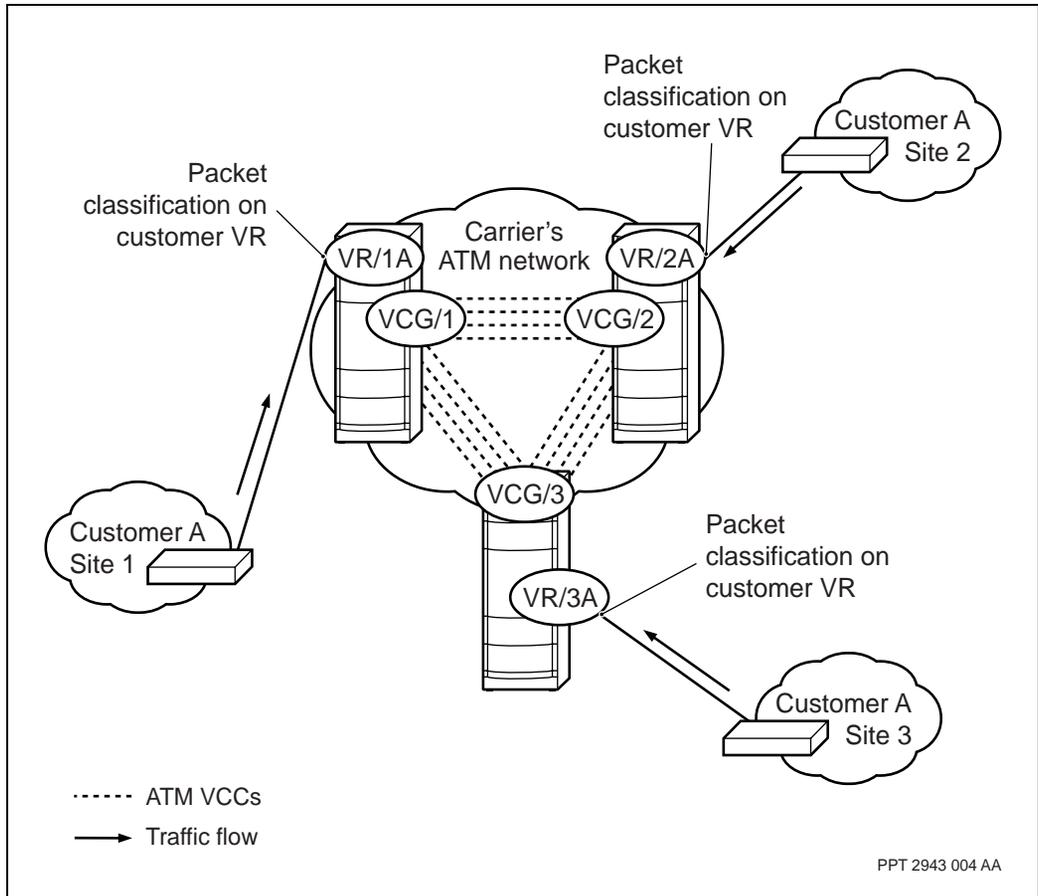
For layer 2 classification on ATM MPE, CQC-based ATM FPs require the assistance of an ILS Forwarder card. For more information, see NN10600-551 *Nortel Networks Multiservice Switch 7400/15000/20000 FP Configuration Reference*.

Layer 2 classification is supported for any egress access and core media even if the ipcos feature is not on the feature list of a logical processor type (LPT).

DSCP-based classification

Carriers can configure IP CoS policies to classify packets based on the DiffServ codepoint (DSCP) in the IP packet header (layer 3 classification) and flow identification parameters (layer 4 classification). Once the ingress port has classified a packet, carriers can configure the customer VR to set the DSCP field of the outgoing IP packet to a value configured from the packet's CoS value. See the figure "Packet classification at ingress" (page 183).

Figure 36
Packet classification at ingress



The IP packet header contains an 8-bit type of service (ToS) byte. The six most significant bits of the ToS field represent the DiffServ codepoint (DSCP). The DSCP field is used to specify a quality of service (QoS) for the packet that can affect the packet delay, throughput, and reliability.

You can enable a virtual router to classify IP packets by configuring a list of ToS byte values in attribute *Vr Ip Pg Policy TosMap tos*. Attribute *Vr Ip Pg EgressCosTreatment tosMask* determines which bits of the ToS byte are

examined. If the ToS value of an incoming packet matches any of the values in the configured list, the packet is assigned the CoS value specified for that IP CoS policy.

It is recommended you keep attribute *Vr Ip Pg EgressCosTreatment tosMask* set to its default value of 0xFC, which represents the DSCP bits.

DSCP-based classification on CQC-based ATM FPs requires the assistance of an ILS Forwarder card. For more information, see NN10600-551 *Nortel Networks Multiservice Switch 7400/15000/20000 FP Configuration Reference*.

Flow-based classification

Flow analysis identifies the flow of IP packets between a source and destination. If enabled, IP CoS can use a flow identification policy to determine the CoS value assigned to packets in the traffic flow.

You can configure IP CoS to distinguish IP traffic flow based on any combination of the following parameters:

- IP address (source or destination) or range of IP addresses
- layer 4 protocol (TCP, UDP or ICMP)
- TCP/UDP port or range of ports

You can specify IP addresses and port numbers as a single value, as a range of values, or as no value (matches all values). IP CoS tries to match the destination IP address and port number first. If there is no match, IP CoS tries to match the source IP address and port number. If the IP address and port number for an incoming packet matches any one of the values in the configured list, IP CoS assigns the CoS value specified for that IP CoS policy.

IP CoS does not classify ICMP packets using IP addresses. If you configure IP CoS to match ICMP packets, IP CoS assigns the same CoS to all ICMP packets, regardless of IP source or destination address. Flow identification also excludes fragmented IP packets or IP packets with options.

For flow-based classification on ATM MPE, CQC-based ATM FPs require the assistance of an ILS Forwarder card. For more information, see NN10600-551 *Nortel Networks Multiservice Switch 7400/15000/20000 FP Configuration Reference*

Transport protocols using static port assignment

Static port assignments simplify IP flow identification. Most protocols fall into the static port assignment protocol category (examples include HTTP, SMTP, and Telnet). The client-server flow is easily identified based on information supplied by the user, typically the client's IP address and port number.

Port numbers in the transport (TCP/UDP) protocol port identify the ends of the logical connections. Some of these ports refer to specific protocols and are often referred to as well-known ports. Protocols running on well-known ports include FTP, Telnet, and HTTP.

IP CoS policies

IP CoS policies determine how packets are classified and the treatment they receive according to their classification. Policy groups contain sets of policies that you can assign to a virtual router or protocol port.

Each policy has an assigned CoS value that references a packet treatment profile under the parent policy group. When there is a policy match, Nortel Networks Multiservice Switch nodes classify the packet with the referenced CoS value.

Each policy has a set of criteria for applying DSCP mapping (specified under the *Vr Ip Pg Policy TosMap* subcomponent) and flow identification (specified under the *Vr Ip Pg Policy IpAddrLayer4Flow* subcomponent). The actions for a match with policy criteria are defined for the policy group through the *Vr Ip Pg ingressCosTreatment* and *Vr Ip Pg egressCosTreatment* components.

You can assign the policy group to the virtual router by setting attribute *Vr Ip cosPolicyAssignment*. If you assign a policy group to a virtual router, all of the policies defined for that policy group are applied to each protocol port under the virtual router.

Alternatively, you can assign policy groups on a port by port basis by setting attribute *Vr Pp IpPort cosPolicyAssignment*. If you assign a policy group to a protocol port, its policies are applied on that particular protocol port only and override any policy groups assigned to the parent virtual router.

Packet treatment at egress

See the following sections for information on how IP CoS can process a packet after it has been classified:

- “Packet marking” (page 186)
- “Class-based packet forwarding” (page 187)

Packet marking

Once a packet has been classified, you can configure IP CoS to mark the packet so that the system forwards information about the assigned CoS value to subsequent nodes without the nodes having to repeat the flow classification process.

Enable packet marking by setting attribute *Vr Ip Pg EgressCosTreatment setTosByte* to yes. When enabled, Nortel Networks Multiservice Switch nodes mark the packet by setting the ToS byte in the IP header, which is based on the CoS value assigned to the packet, the initial value of the packet’s ToS byte, and attribute *Vr Ip Pg EgressCosTreatment tos* (new ToS). Attribute *Vr Ip Pg EgressCosTreatment tosMask* determines which bits of the packet’s ToS byte are marked. It is recommended you keep attribute *Vr Ip Pg EgressCosTreatment tosMask* set to its default value of 0xFC, which represents the DSCP bits.

PQC-based FPs use the default ToS mask of 0xFC regardless of the value of attribute *Vr Ip Pg EgressCosTreatment tosMask*.

The ToS byte is set as follows:

```
[(initial ToS)&(~tosMask)] | [(new ToS)&(tosMask)]
```

Example

Using the default value of 0xFC for *tosMask*, the ToS byte is set as follows:

```
(initial ToS & 00000011) | (new ToS & 11111100)
```

Class-based packet forwarding

IP CoS allows IP traffic differentiation into four separate classes of service. You can map each CoS value to a specific set of QoS parameters for a given media type. In this way, you can control the service that packets receive when they are forwarded out of the Nortel Networks Multiservice Switch node, based on their CoS classification.

For more information, see the following sections:

- “Drop precedence” (page 187)
- “Scheduling class” (page 188)

Drop precedence

IP CoS allows you to configure the drop precedence of a packet. The drop precedence determines a packet’s importance when being forwarded. It is used to determine whether or not an IP packet should be dropped to reduce traffic load during traffic congestion.

You can associate a drop precedence with a CoS value by setting attribute *Vr Ip Pg IngressCosTreatment discardPriority*. There are four drop precedence settings: unchanged, low, medium, and high. A value of low indicates a low drop precedence, meaning that the packet is less likely to be dropped than packets with a different drop precedence. A value of high indicates a high drop precedence, meaning that the packet is more likely to be dropped than packets with a different drop precedence. If you assign a *discardPriority* of unchanged, the drop precedence of IP packets is not modified by IP CoS but is instead determined by ingress layer 2 media.

Drop precedence is supported on all WAN media (ATM MPE, frame relay, and PPP) on all applicable SBIC-based and most PQC2-based FPs. It is not supported on PPP with PQC2-based FPs and Ethernet on SBIC-based FPs.

Scheduling class

IP CoS lets you configure the scheduling class of a packet. The scheduling class is used to determine when a packet is scheduled to be transmitted at its egress interface relative to other packets.

Scheduling class values are derived from the *Vr Ip Pg EgressCosTreatment emissionPriority* attribute used by the IP port where the packet is transmitted. The emission priority attribute of the *EgressCosTreatment* whose instance value matches the CoS index assigned to the packet is used.

The *emissionPriority* attribute of the *EgressCosTreatment* component is mapped to a scheduling class value according to Table 28, “EgressCosTreatment emissionPriority to scheduling class mapping,” (page 188).

Table 28
EgressCosTreatment emissionPriority to scheduling class mapping

Provisioned emissionPriority value	Operational Scheduling Class value		
	8 queues	4 queues	2 queues
1	5	3	1
2	4	2	0
3	2	1	0
4	0	0	0
5	0	0	0
6	0	0	0
7	0	0	0
8	0	0	0

A numerically lower value of the *emissionPriority* (higher value of scheduled class) indicates the packet is less likely to be delayed. A numerically higher value (lower value of scheduled class) indicates the packet is more likely to be delayed. The actual scheduling behavior of packets depends on the scheduling mechanism and number of queues at the egress interface.

Frame relay DTE class-based forwarding

Nortel Networks Multiservice Switch systems frame relay DTE (FrDte) supports CoS to QoS mapping over multiple DLCIs or a single DLCI.

For more information, see the following sections:

- “CoS to QoS mapping over multiple DLCIs” (page 189)
- “CoS to QoS mapping over a single DLCI” (page 190)

CoS to QoS mapping over multiple DLCIs

For CoS to QoS mapping over multiple DLCIs, IP CoS lets you create up to four DLCIs to each next hop router, with each DLCI having separate QoS parameters and a different CoS index.

All dynamic DLCIs under the same *FrDte* component share the same assigned CoS index, which is set in the *FrDte DynamicDlciDefaults ipcos* attribute, unless configuration on a static DLCI overrides it. For full service differentiation, use static DLCIs and assign a unique CoS value to each DLCI using its *FrDte StaticDlci ipcos* attribute.

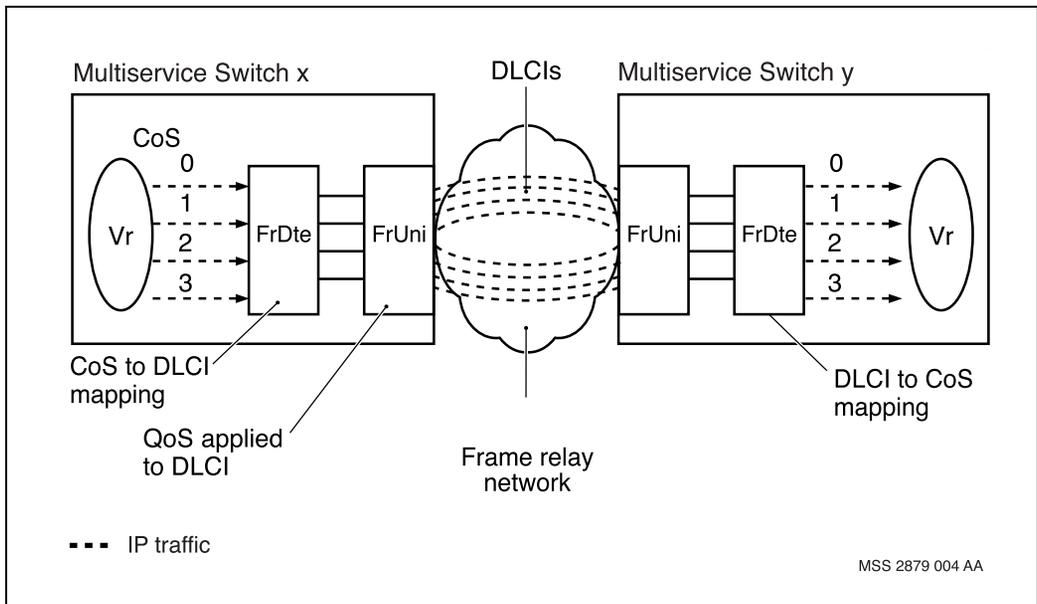
The drop precedence assigned to an IP packet arriving on a frame relay connection is determined by the discard eligibility (DE) bit to drop precedence mapping configured on the FRUNI. This can be overridden by the CoS to drop precedence mapping configured on the CoS policy group used by the ingress protocol port.

When a packet goes out of a frame relay DTE interface where there are multiple DLCIs available for the same IP next hop, the node forwards the packet over the DLCI whose associated *ipcos* attribute matches the packet's CoS value. See the figure “CoS to QoS mapping on multiple DLCIs” (page 190).

If the CoS assigned to the packet does not match an operational DLCI (it exists and is operationally active) with the same CoS value, the packet is transmitted through an operational DLCI with the next lowest CoS value. If no operational DLCIs with a lower CoS value exist, the packet is transmitted through an operational DLCI with the next highest CoS value.

The remote end of each DLCI controls the QoS characteristics for that DLCI. To obtain a particular set of characteristics within a Nortel Networks Multiservice Switch node, route the DLCI through a virtual framer to a FR UNI, where you can set the QoS parameters. For more information, see NN10600-900 *Nortel Networks Multiservice Switch 7400/15000/20000 Frame Relay Technology Fundamentals*.

Figure 37
CoS to QoS mapping on multiple DLCIs



CoS to QoS mapping over a single DLCI

Nortel Networks Multiservice Switch system IP CoS allows flows of IP packets that are transmitted over a single frame relay DTE DLCI to be differentiated by QoS characteristics. The differentiation is done by setting the emission priority and drop precedence of packets in certain IP flows based on their assigned CoS.

The drop precedence assigned to an IP packet arriving on a frame relay connection is determined by the discard eligibility (DE) bit to drop precedence mapping configured on the FRUNI. This can be overridden by the CoS to drop precedence mapping configured on the CoS policy group used by the ingress protocol port.

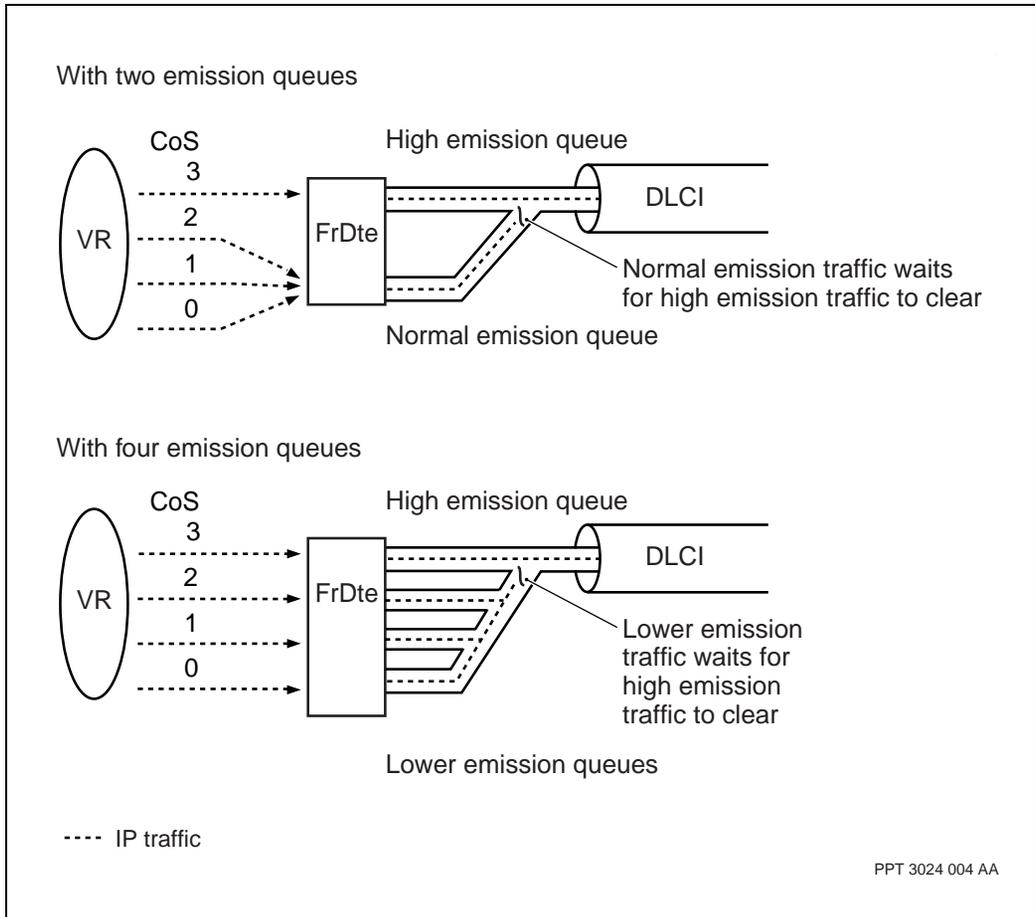
Multiservice Switch nodes use the policy group configured on the egress protocol port for emission priority mapping. When a packet goes out of a frame relay DTE interface on a non-channelized FP, the node assigns an emission priority to the packet based on the packet's CoS value. For CoS to emission priority mapping over a single DLCI, you can specify up to four different emission priorities, one for each CoS value.

CoS to emission priority mapping is not supported over *VirtualFramer* components (a logical connection). Instead, the mapping must be over either *Framer* (a hairpin connection) or *Dconn* (a direct connection) components.

Emission priority mapping based on CoS over a single frame relay DTE DLCI is not supported on channelized SBIC-based WAN FPs.

There are four emission queues on Multiservice Switch 15000 and Multiservice Switch 20000 FPs, and Multiservice Switch 7400 MSA32 FPs. There are two emission queues on non-channelized SBIC-based WAN FPs. See the figure "Emission priority mapping on a single DLCI" (page 192).

Figure 38
Emission priority mapping on a single DLCI



IP-optimized DLCI class-based forwarding

IP-optimized DLCIs support CoS to QoS mapping over a single DLCI. Nortel Networks Multiservice Switch system IP CoS allows flows of IP packets that are transmitted over a single IP-optimized DLCI to be differentiated by QoS characteristics. This differentiation is done by setting the emission priority and drop precedence of packets in certain IP flows based on their assigned CoS.

The drop precedence assigned to an IP packet arriving on a frame relay connection is determined by the discard eligibility (DE) bit to drop precedence mapping configured on the FRUNI. This can be overridden by the CoS to drop precedence mapping configured on the CoS policy group used by the ingress protocol port.

Multiservice Switch nodes use the policy group configured on the egress protocol port for emission priority mapping and assigns an emission priority to the packet based on the packet's CoS value. IP-optimized DLCIs support four emission priorities.

ATM MPE class-based forwarding

ATM MPE supports CoS to QoS mapping over multiple virtual channel connections (VCCs) or over a single VCC.

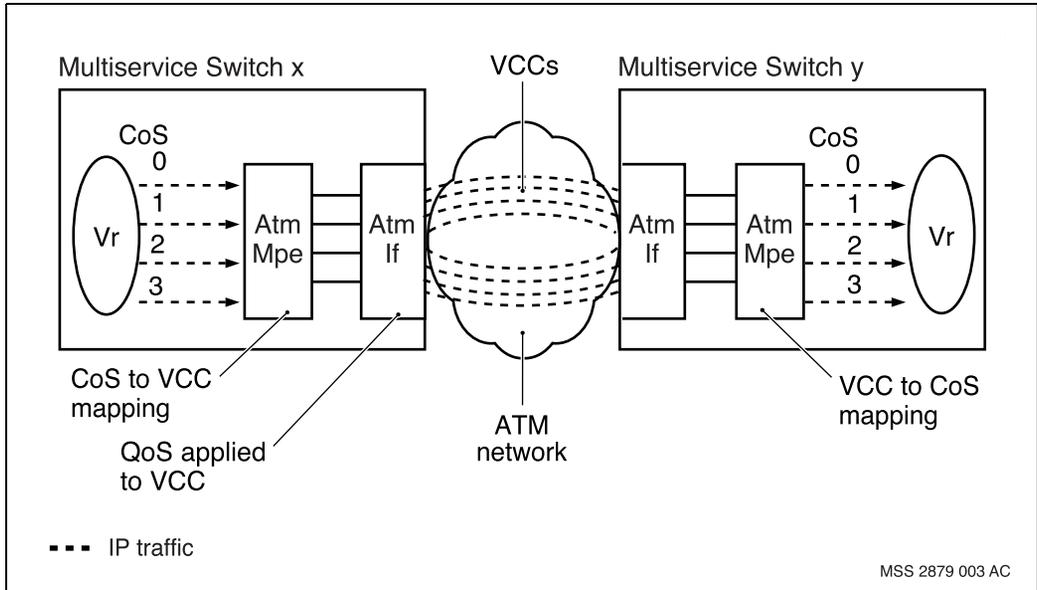
For CoS to QoS mapping over multiple VCCs, IP CoS lets you create up to four VCCs to each next hop router, with each VCC having separate QoS parameters and a different CoS index. Set each associated *AtmMpe* *AtmConnection* *ipcos* attribute.

When a packet goes out on an ATM MPE interface, the node forwards the packet over the VCC whose associated *ipcos* attribute matches the packet's CoS value. See the figure "CoS to QoS mapping on multiple VCCs" (page 194).

The drop precedence assigned to an IP packet arriving on an ATM connection is determined by the cell loss priority (CLP) bit to drop precedence mapping configured on the ATM connection. This can be overridden by the CoS to drop precedence mapping configured on the CoS policy group used by the ingress protocol port.

You can configure the service class of each VCC under the *AtmIf* component as required, in order to provide different levels of service. Service classes include CBR, rtVBR, nrtVBR, UBR, and ABR. In addition, you can use the virtual path connection (VPC) functionality to apply traffic management functions to multiple VCCs. For more information, see NN10600-705 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Management Fundamentals*.

Figure 39
CoS to QoS mapping on multiple VCCs



Gigabit Ethernet class-based forwarding

Quality of service for gigabit Ethernet is applied using IP CoS when component *Vr Ip Dsd* is not provisioned. When *Vr Ip Dsd* is provisioned, quality of service is applied using differentiated services as described in “Differentiated Services in a Layer 3 network” (page 29).

Gigabit Ethernet supports only layer 2 packet classification. Use attribute *Vr Pp IpPort ipCos* to assign a CoS index. The drop precedence is always set to high (3) and the scheduling class is always set to the lowest priority (0) for IP packets arriving on gigabit Ethernet ports.

Gigabit Ethernet supports CoS interworking with PQC-based ATM media. The interworking media combinations are described in the following figures:

- “Ingress on gigabit Ethernet; egress on PQC-based ATM” (page 196)
- “Ingress on PQC-based ATM; egress on gigabit Ethernet” (page 197)

You may see the following semantic warning coming from the IP port associated with a gigabit Ethernet interface:

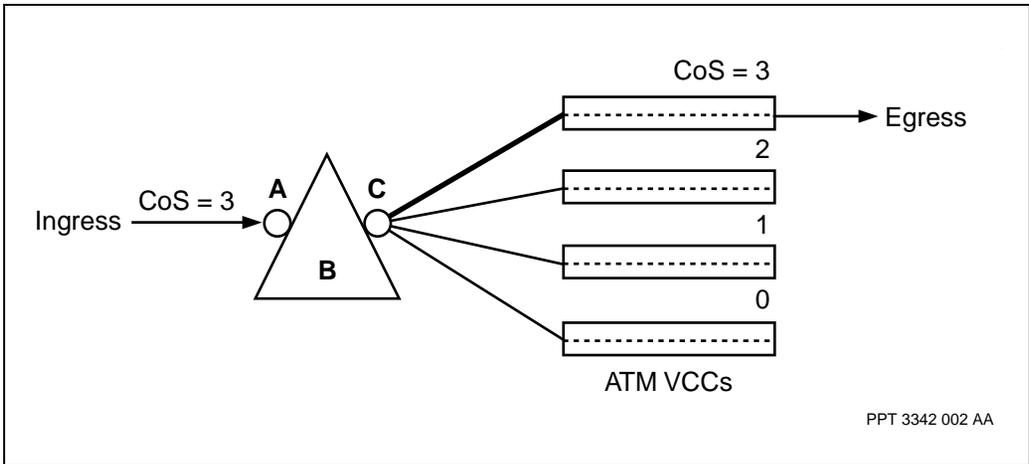
Warning: If a CosPolicyGroup is linked to the Ip or IpPort, then ipCos must be included in the featureList of the LogicalProcessorType linked to the Logical Processor that is used by the media linked to the ProtocolPort.

A Check Prov command produces this warning if

- the Ip or IpPort is linked to a CosPolicyGroup and
- the LogicalProcessorType linked to the LogicalProcessor that is used by the media linked to the ProtocolPort does not include ipCos in the featureList.

To eliminate this warning, you would normally add ipCos to the featureList of the LogicalProcessorType linked to the LogicalProcessor that is used by the media linked to the ProtocolPort, or ensure that the Ip and IpPort are not linked to a CosPolicyGroup; however, since the ipCos feature is not supported on the 4pGe FP, the warning cannot be eliminated.

Figure 40
Ingress on gigabit Ethernet; egress on PQC-based ATM



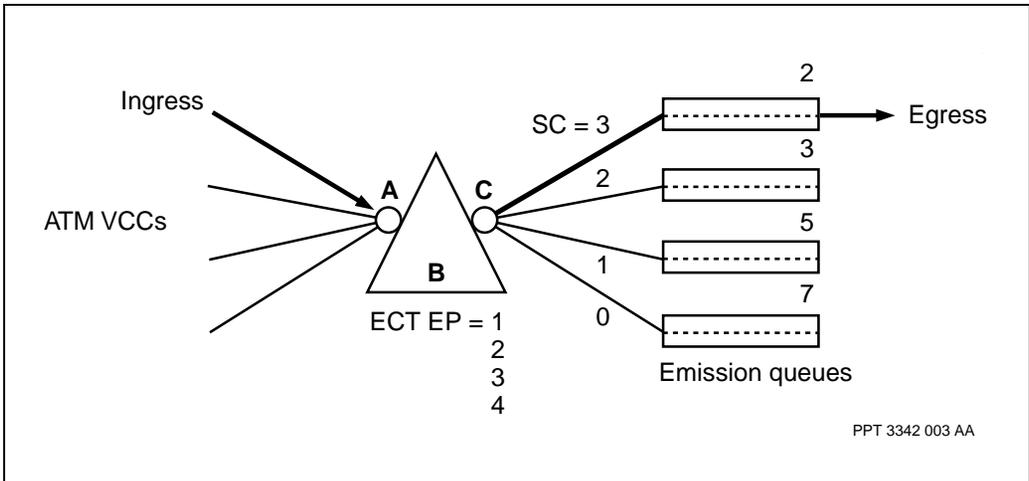
In the above figure, the following sequence occurs:

A The IP packet arrives on the ingress gigabit Ethernet protocol port. The IP packet is assigned a CoS index, using layer 2 classification, in attribute *Vr Pp IpPort ipCos*.

B The CoS index of the IP packet determined in step A is matched against the CoS indexes of the ATM connections belonging to the egress protocol port via attribute *AtmMpe Ac ipCos*. See “ATM MPE class-based forwarding” (page 193) for more information.

C The IP packet is transmitted via the egress ATM connection determined by step B.

Figure 41
Ingress on PQC-based ATM; egress on gigabit Ethernet



In the above figure, the following sequence occurs:

A The IP packet arrives on the ingress protocol port. The IP packet is assigned a CoS index, using layer 2, layer 3, or layer 4 classification.

B The CoS index is matched against a corresponding emission priority (EP), which is defined in attribute *Vr Ip Pg Ect ep*. This EP is mapped to one of eight emission queue of the egress gigabit Ethernet port. For the mapping values, see the table below. If no policy group is assigned to the egress gigabit Ethernet port, then queue 7 is used, regardless of the CoS index assigned to the that corresponds to the ingress PVC.

ECT emission priority	GigE queue number
1	2
2	3
3	5
4-8	7

C The IP packet is transmitted at the egress gigabit Ethernet protocol port on the queue selected by the EP that is now associated with the packet. The traffic management characteristics associated with each EP are defined in component *Lp Ethernet Tm Ep*.

Point-to-point protocol class-based forwarding

Nortel Networks Multiservice Switch system IP CoS allows flows of IP packets that are transmitted over a point-to-point protocol (PPP) interface to be differentiated by QoS characteristics. The differentiation is done by setting the emission priority and drop precedence of packets in certain IP flows based on their assigned CoS.

Multiservice Switch nodes perform layer 2 classification on every IP packet that arrives on a PPP interface based on the incoming protocol port. The node uses the policy group configured on the ingress protocol port for drop precedence mapping.

Multiservice Switch nodes use the policy group configured on the egress protocol port for emission priority mapping. Therefore, when a packet goes out a PPP interface, the assigns an emission priority to the packet based on the packet's CoS value. For CoS to emission priority mapping over PPP, you can specify up to four different emission priorities, one for each CoS value.

Emission priority mapping based on CoS over PPP is not supported on channelized SBIC-based WAN FPs or PQC-based FPs.

IP CoS over virtual media

Virtual media protocol ports support layer 2 classification for multicast packets only. Multicast packets transmitting from one VR to another receive a CoS value at the ingress to the second VR, as configured through the *Vr ProtocolPort IpPort ipCoS* attribute of the associated ingress virtual media protocol port.

Chapter 14

IP CoS to IP DiffServ migration

Use this section to learn about the conceptual information pertaining to Nortel Networks Multiservice Switch Cos to Diffserv migration.

Navigation

- “Overview of CoS to DiffServ migration” (page 199)
- “IP DiffServ advantages” (page 200)
- “Comparison between IP CoS and IP DiffServ provisioned attributes” (page 208)

Overview of CoS to DiffServ migration

IP CoS and IP DiffServ are the two ways IP differentiated services can be deployed in Nortel Networks Multiservice Switch systems. DiffServ offers several advantages over IP CoS and should be considered as the best choice for differentiated services on Multiservice Switch 15000 and Multiservice Switch 20000 nodes. “IP DiffServ advantages” (page 200) describes the advantages of DiffServ over IP CoS.

The section “Migration stages” (page 205) describes the potential stages to consider in order to convert an IP VPN network from the CoS model to the DiffServ model.

The section “Comparison between IP CoS and IP DiffServ provisioned attributes” (page 208) compares the components used to provision either IP CoS or DiffServ.

IP DiffServ advantages

IP DiffServ is the best solution for deploying differentiated services on Nortel Networks Multiservice Switch 15000 and Multiservice Switch 20000 nodes. IP CoS is limited to four treatments based on four classes associated with the IP packet. DiffServ offers the following improvements.

- more available treatments for each IP packet forwarded by the router (VirtualRouter, Router, or VpnRouteForwarder)
- greater control of DSCP values that can be marked into the DSCP field of each IP packet header
- greater control of packets that are generated locally at the router
- simplified provisioning model, network planning, and deployment
- improved standards compliance
- improved scalability to future deployment and enhancements of IP differentiated services that are not feasible with IP CoS
- the DiffServ architecture allows the user to mark packets with IETF defined DSCP marking rather than bit patterns
- IP DiffServ allows access to all of the IETF defined PHBs
- more PHB support in the carrier's network and the customer access interfaces.
 - depending on the FP hardware, IP DiffServ can support up to 12 PHBs at the ATM backbone interfaces and up to 24 PHBs at the customer access interfaces
 - the ability to provide a QoS mapping between the customer's PHB and the carrier's PHB based on a service level agreement that does not impact customer tandem traffic
- as part of the RFC2764 DiffServ solution, carriers have a mechanism to mark outer IP Tunnel header packets with an SLA Customer-to-Carrier DSCP Mapping function
- access to the IP Policing feature on customer access interfaces.
- the ability to provide QoS capabilities on MS3 architecture-based access and trunk FPs.

Determining migration requirements

Before beginning the actual network migration, the following requirements must be determined:

- the carrier DiffServ domain. Refer to “Determining the carrier DiffServ domain” (page 201).
- the iIP CoS to DiffServ PHB mapping. Refer to “Determining the IP CoS to DiffServ PHB mapping” (page 202).
- the carrier ATM VCC Migration. Refer to “Determining the carrier ATM VCC migration” (page 202).

Determining the carrier DiffServ domain

The first step that needs to be performed before beginning the actual carrier network migration is to determine the set of carrier network PHBs that will be supported in the backbone network.

The set of carrier network PHBs chosen should contain and/or be similar to the QoS treatments currently provided in the existing carrier network.

The PHBs chosen may be DiffServ RFC recommended PHBs and/or PHBs defined in one of several DiffServDomains as default PHBs. Essentially a single provisionable DiffServDomain instance should be chosen that will contain the set the chosen PHBs, and will be applied to all VCGs during the carrier network migration stage.

For example, a carrier may have the following types of traffic in their carrier network:

- interactive traffic (such as VoIP, or Video)
- responsive traffic (such as eCommerce and eBusiness)
- best effort traffic (such as e-mail and ftp)
- network control traffic (such as IP routing traffic)

Table 29, “Example set of carrier network PHBs,” (page 202) illustrates a possible mapping of this traffic to PHBs currently defined under the DiffServDomain instance.

Table 29
Example set of carrier network PHBs

Traffic Characteristics	Typical Applications	DSCP/PHB
best effort, delay tolerant	e-mail, ftp	df
responsive	eCommerce, eBusiness	af11
interactive, delay and loss intolerant	VoIP, video	ef
network control traffic	IP routing (OSPF, BGP)	cs6

Determining the IP CoS to DiffServ PHB mapping

Now that the set of carrier network DiffServ PHBs have been determined for different types of traffic in the carrier network, the indices for the migration of the original IP VPN network, which currently provides QoS based on a CoS index, need to be mapped into one of the PHBs.

Determine which CoS values (0 to 3) will be mapped to each DSCP/PHB that is defined in your new DiffServ Domain. Note that control traffic in the CoS model is always transmitted with a CoS of 3. Table 30, “Example of CoS index mapping to a DiffServ PHB,” (page 202) provides an example of CoS index mapping to a DiffServ PHB

Table 30
Example of CoS index mapping to a DiffServ PHB

CoS	DSCP/PHB
0	df
1	af11
2	ef
3	cs6

Determining the carrier ATM VCC migration

In a CoS-based network, locally generated IP routing traffic is transmitted through the virtual circuit with an *ipCos* value of 3.

In a DiffServ-based network, locally generated IP routing traffic is transmitted through the virtual circuit with the *ipCos* value equivalent to the *sc4q* value defined in the traffic class of PHB assigned to locally generated IP routing traffic.

The default definitions of all DiffServDomain instances contain:

- locally generated routing traffic is assigned the PHB “cs6”
- the PHB “cs6” is assigned the traffic class of “network”
- the *sc4q* attribute value of the “network” traffic class is assigned 2

Therefore, in most cases in DiffServ, by default, locally generated IP routing traffic is transmitted through the virtual circuit with an *ipCos* value of 2. Also in most cases, carriers assign different ATM Service Categories to the VCCs corresponding to *ipCos* values 2 and 3, to provide separate QoS forwarding behaviour for different traffic characteristics on each of the VCCs.

During the carrier network migration, you may be required to perform the following provisioning modification in order to preserve the QoS forwarding behaviour of IP packets:

- Change the *ipCos* attribute value on the applicable *atmConnection* attributes. Note that any *atmConnection* attribute value changes (including the *ipCos* attribute) requires the *AtmMpe* component to be manually locked and unlocked for the changes to be recognized. Also note that locking and unlocking the *AtmMpe* component causes a temporary traffic outage on the attached protocol port, and the corresponding ARP entries of the interface will be flashed.

There are typically two IP VPN configurations that customers use for their carrier traffic. The type of configuration used usually depends on whether interactive data traffic or control traffic is more important to a customer within their carrier network.

Control traffic may be more desirable to customers who are looking to ensure network stability. However, since the introduction of the IP CoS feature and RFC2764 IP VPN solution on the Nortel Networks Multiservice Switch, interactive data traffic has become of the highest importance, in order to ensure low loss and low delay in this type of traffic. The IP DiffServ feature

has been developed with this need in mind, and provides default provisioning values for DiffServ PHBs to treat interactive data traffic with a higher importance than control traffic. This is reflected in the various DiffServDomain component instances that are available to be provisioned.

This means that if a customer wishes to switch from a control traffic preference to an interactive data preference, after migrating to DiffServ, no specific changes on the carrier ATM VCCs are required. After the carrier network migration to DiffServ, the traffic to ATM service category mapping is shown in Table 31, “Traffic mapped to ATM SC, with a change to interactive data traffic preference,” (page 204).

Table 31
Traffic mapped to ATM SC, with a change to interactive data traffic preference

Traffic Characteristics	AtmMpe AtmConnection ipCos	Atm VCC Service Category
interactive	3	cbr
network control	2 (as determined by the default local routing source attribute set to DSCP “cs6”, and PHB “cs6” set by default to a sc4q value of 2)	rtvbr
responsive	1	nrtvbr
best effort	0	ubr

If it is still required, after the carrier migration to DiffServ, that control traffic be preferred over interactive data traffic and assigned to the ATM VCC with an ATM “cbr” service category, then one of the three provisioning options described earlier should be performed during the carrier network migration.

Migration stages

There are various stages of converting a VPN network from the CoS model to the DiffServ model. The customer must migrate the following:

- the carrier core network (flash-cut or gradual migration). Refer to “Carrier network migration” (page 205).
- the individual customer VPN at all sites (basic, customer-aware, or customer-edge VPN migration). Refer to “Individual customer VPN migration” (page 206).
- the individual VPN site (per node). Refer to “Individual customer VPN migration” (page 206).
- the individual customer virtual routers (cVRs). Refer to “Single customer router (cVR) migration” (page 208).

Carrier network migration

All the VCGs could be attempted to be converted over to DiffServ all at once (flash-cut migration), or one at a time (gradual migration) since VCGs reside on different nodes, depending on the physical outlay of the carrier network.

Flash-cut migration

For a flash-cut migration, perform the following to all VCGs:

- Provision a *DiffServDomain* (DSD) under the virtual router
- Provision a *VpnCosMap* under the VCG’s DSD with a populated mapping for VPN CoS To DiffServ migration purposes. This maps a customer’s CoS to the carrier’s PHB.
- If needed, migrate the ATM VCCs so that the different traffic types are given the same ATM Service Category in the carrier network as prior to the migration.

Single VCG (Gradual) migration

For gradual migration, determine the first VCG to migrate to DiffServ. Before migrating the VCG, it is recommended that you first provision CoS-based DSCP marking on all the remote protocol ports of the connections used by the VCG you are to migrate. The marking performed on these protocol ports should be the same as the CoS to DSCP mapping that has been determined

for the core network. This marking is required such that the targeted VCG that is being converted into DiffServ will be able to classify IP packets correctly based on the marked DSCPs.

Once the marking has been enabled on the remote VCG protocol ports, migrate your VCG from CoS to DiffServ:

- Remove any possible policy groups (including those performing marking in order to interwork with other VCGs in the carrier network already migrated to DiffServ) that may be assigned to VCG or its protocol ports
- Provision a *DiffServDomain* (DSD) under the VR component of the VCG.
- Provision a *VpnCosMap* under the VCG's DSD with a populated mapping for VPN CoS To DiffServ migration purposes. This maps a customer's CoS to the carrier's PHB.
- If needed, migrate the ATM VCCs such that the different traffic types are given the same ATM Service Category in the carrier network as prior to the migration.

Once the carrier network has been migrated, then migration on the individual customer VPNs and customer virtual routers can be performed.

Individual customer VPN migration

There are a number of ways that customer traffic is treated going in and out of the carrier's network. The treatment typically depends on the SLAs available from the carrier and the specific SLA between the customer and the carrier.

There are typically three different VPN configurations most used by customers:

- “Basic VPN configuration” (page 207)
- “Customer Aware VPN configuration” (page 207)
- “Customer Edge VPN configuration” (page 207)

Basic VPN configuration

In a Basic VPN configuration, typically no SLA exists between the customer and the carrier, and the following occur on the applicable tunnel end-points:

- On Tunnel Entry, the customer's traffic is mapped 1:1 from a CoS perspective into the carrier's network without manipulation.
- On Tunnel Exit, the QoS treatment provided at the access protocol port, into the customer's network, is based on the link classification occurring on the VCG, with no classification at the cVR's tunnel protocol port to look at the customer's header.

Customer Aware VPN configuration

In a Customer Aware VPN configuration, An SLA exists between the customer and the carrier, and the customer's DSCP is marked and understood within the customer's networks.

- On Tunnel Entry, the customer's traffic is mapped into the carrier's network based on IP CoS classification occurring on the cVR's access protocol port.
- On Tunnel Exit, QoS treatment provided at the access protocol port, into the customer's network, is based on IP CoS classification at the tunnelling protocol port of the cVR. Most configurations perform DSCP classification at the tunnelling protocol port to classify the customer's DSCP that was marked within the originating customer network

Customer Edge VPN configuration

An SLA exists between the customer and the carrier, and the customer's DSCP may be marked by the customer network or the cVR. This is a case where some customer sites may understand and act upon the customer DSCP and some may not.

- On Tunnel Entry, the customer's traffic is mapped into the carrier's network based on IP CoS classification occurring on the cVR's access protocol port, and the customer's DSCP is marked by the cVR's tunnelling protocol port according to the CoS value.
- On Tunnel Exit, QoS treatment provided at the access protocol port (into the customer's network) is based on IP CoS classification at the tunnelling protocol port of the cVR. Most configurations perform DSCP classification at the tunnelling protocol port to classify the customer's DSCP that was marked by the originating cVR.

Determining the Customer VPN DiffServ Domain

In almost all cases, the DiffServ domain chosen for the carrier network is the same one provisioned for the individual VPNs. A different DiffServ domain may be chosen for the VPN if the PHBs used by the VPN customer in their network are different and/or are defined differently in regards to QoS treatment. Typically a VPN in a Customer Aware VPN Configuration may require a different DiffServ domain from the carrier network, but in the other VPN configurations (Basic VPN and Customer Edge), the same DiffServ Domain is typically used.

Single customer router (cVR) migration

The migration of a single cVR involves the generic IP CoS to DiffServ migration procedure for a virtual router. The provisioning steps for a Basic VPN configuration, a Customer Aware VPN configuration, and a Customer Edge VPN configuration are found in NN10600-591 *Nortel Networks Multiservice Switch 7400/15000/20000 Layer 3 Traffic Management Configuration*.

Comparison between IP CoS and IP DiffServ provisioned attributes

There is no direct mapping in component provisioning between IP CoS and IP DiffServ. IP CoS is provisioned using *Vr Ip CosPolicyGroup* components, while DiffServ is configured using the *Vr DifferentiatedServicesDomain* and *Vr Ip DifferentiatedServices* profile components.

The following tables show the configuration relationship between DiffServ and IP CoS.

- “IP CoS and IP DiffServ software activation attributes” (page 209)
- “IP CoS and IP DiffServ layer 2 configuration attributes” (page 209)
- “IP CoS and IP DiffServ layer 3 configuration attributes” (page 210)
- “IP CoS and IP DiffServ Phb definition attributes” (page 210)
- “IP CoS and IP DiffServ drop precedence selection attributes” (page 210)
- “IP CoS and IP DiffServ local packet marking attributes” (page 210)
- “IP CoS and IP DiffServ Dscp/Tos classification attributes” (page 211)

Table 32
IP CoS and IP DiffServ software activation attributes

IP CoS configuration	IP DiffServ configuration
<i>Sw Lpt fl ipCos</i>	<i>Sw Lpt fl ipDiffServ</i>
Feature ip DiffServ is required for the DiffServ profile for the interface (<i>Vr Ip DiffServProfile</i>) but not for the DiffServ domain for the virtual router (<i>Vr Dsd</i>)	

Table 33
IP CoS and IP DiffServ provisioning components

IP CoS configuration	IP DiffServ configuration
<i>Vr Ip CosPolicyGroup</i>	<i>Vr Ip DiffServProfile</i> (for classification) <i>Vr Dsd</i> (for treatments)

Table 34
IP CoS and IP DiffServ layer 2 configuration attributes

IP CoS configuration	IP DiffServ configuration
<i>Vr Pp ipPort ipCos</i>	<i>Vr Pp ipPort ipCos</i> <i>Vr Ip DiffServProfile LinkMap</i>
<i>AtmMpe Ac ipCos</i>	<i>AtmMpe Ac ipCos</i> <i>Vr Ip DiffservProfile LinkMap</i>
<i>FrDte StDlci ipCos</i>	<i>FrDte StDlci ipCos</i> <i>Vr Ip DiffservProfile LinkMap</i>
<i>IpDlciGrp FrConnection ipCos</i>	<i>IpDlciGrp FrConnection ipCos</i> <i>Vr Ip DiffservProfile LinkMap</i>

Table 35
IP CoS and IP DiffServ layer 3 configuration attributes

IP CoS configuration	IP DiffServ configuration
<i>Vr Ip cosPolicyAssignment</i>	(assignment to VR npt supported)
<i>Vr Pp IpPort cosPolicyAssignment</i>	<i>Vr Pp IpPort Is linktoDiffservProfile</i> (for classification) <i>Vr Pp IpPort Es linktoDiffservProfile</i> (for marking)

Table 36
IP CoS and IP DiffServ Phb definition attributes

IP CoS configuration	IP DiffServ configuration
<i>Vr Ip CosPolicyGroup Ect ep</i>	<i>Vr DiffServDomain Phb trafficClass</i> <i>Vr DiffservDomain trafficClass sc8q/</i> <i>sc4q</i>
<i>Vr Ip CosPolicyGroup lct dp</i>	<i>Vr DiffServDomain Phb dp</i>

Table 37
IP CoS and IP DiffServ drop precedence selection attributes

IP CoS configuration	IP DiffServ configuration
<i>Vr Ip CosPolicyGroup lct dp</i>	<i>Vr Ip DiffServ IpPort Is dpMode</i>

Table 38
IP CoS and IP DiffServ local packet marking attributes

IP CoS configuration	IP DiffServ configuration
Hard coded to dscp equals 0	<i>Vr Dsd phbGeneralSource</i>
Hard coded to dscp equals cs6	<i>Vr Dsd phbRoutingSource</i>

Table 39
IP CoS and IP DiffServ Dscp/Tos classification attributes

IP CoS configuration	IP DiffServ configuration
<i>Vr Ip Pg Policy assignedCos</i> <i>Vr Ip Pg Policy TosMap tos</i>	<i>Vr Ip DiffServProfile DscpMap assigned Phb</i> (where <i>Vr Ip DiffServProfile</i> is assigned to <i>Vr Pp IpPort Ingress Services</i>)

Table 40
IP CoS and IP DiffServ Dscp marking attributes

IP CoS configuration	IP DiffServ configuration
<i>Vr Ip Pg Ect tosMask</i>	(always OxFC, not provisioned)
<i>Vr Ip Pg Ect setTos I p Pg Ect tos</i>	<i>Vr Ip DiffServProfile DscpMap assigned Phb</i> (where <i>Vr Ip DiffServProfile</i> is assigned to <i>Vr Pp IpPort Egress Services</i>)

Table 41
IP CoS and IP DiffServ Flow/Mf classification attributes

IP CoS configuration	IP DiffServ configuration
<i>Vr Ip Pg Policy IpAddrLayer4Flow prefix</i>	<i>Vr Ip DiffServProfile MfMap IpAddrLayer4Flow prefix</i>
<i>Vr Ip Pg Policy IpAddrLayer4Flow prefixLength</i>	<i>Vr Ip DiffServProfile MfMap IpAddrLayer4Flow prefixLength</i>
<i>Vr Ip Pg Policy IpAddrLayer4Flow protocol</i>	<i>Vr Ip DiffServProfile MfMap IpAddrLayer4Flow protocol</i>
(Sheet 1 of 2)	

Table 41 (continued)
IP CoS and IP DiffServ Flow/Mf classification attributes

IP CoS configuration	IP DiffServ configuration
<i>Vr Ip Pg Policy IpAddrLayer4Flow portNumberRange</i>	<i>Vr Ip DiffServProfile MfMap IpAddrLayer4Flow portNumberRange</i>
<i>Vr Ip Pg Policy assignedCos</i>	<i>Vr Ip DiffServProfile MfMap assignedPhb (where Vr Ip DiffServProfile is assigned to Vr Pp IpPort Ingress Services)</i>
(Sheet 2 of 2)	

Nortel Networks Multiservice Switch 7400/15000/20000 Layer 3 Traffic Management Fundamentals

Release 6.1

Copyright © 2004 Nortel Networks.
All Rights Reserved.

NORTEL, NORTEL NETWORKS, the globemark design, the NORTEL NETWORKS corporate logo, DPN and PASSPORT are trademarks of Nortel Networks.

Publication: NN10600-590
Document status: Standard
Document version: 6.1S2
Document date: November 2004
Printed in Canada

