



7210 SERVICE ACCESS SWITCH

**7210 SAS 7210 SAS-K 2F2T1C and 7210
SAS-K 2F4T6C OS QoS guide
Release 9.0.R4**

3HE12057AAAATQZZA

Issue: 1

April 2017

Nokia is a registered trademark of Nokia Corporation. Other products and company names mentioned herein may be trademarks or trade names of their respective owners.

The information presented is subject to change without notice. No responsibility is assumed for inaccuracies contained herein.

© 2017 Nokia.

Contains proprietary/trade secret information which is the property of Nokia and must not be made available to, or copied or used by anyone outside Nokia without its written authorization. Not to be used or disclosed except in accordance with applicable agreements.

Table of Contents

Table of Contents	3
List of Figures	11
List of Tables	13
Preface	15
About This Guide	15
Audience	15
List of Technical Publications	15
Getting Started	17
In This Chapter	17
Nokia 7210 SAS-Series Services Configuration Process	17
QoS Policies	19
In This Chapter	19
QoS Overview	20
Overview of QoS Policies on 7210 SAS-K 2F2T1C	21
Overview of QoS Policies on 7210 SAS-K 2F4T6C	23
QoS marking interpretation	24
Summary of major functions of QoS policies	27
Service and Network QoS Policies	30
Network QoS Policies	31
Network QoS Policies on network ports on 7210 SAS-K 2F4T6C	31
Ingress Classification Support for Network ports	33
Egress Marking Support for Network ports	33
Network QoS policies on Access-uplink ports	36
Ingress Classification Support for access-uplink ports	37
Network Queue policies	38
Network Queue policies on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C	38
Queues and Queue Parameters	41
Queue- Queue ID	42
Queue- Committed Information Rate	42
Queue- Peak Information Rate	44
Queue- Adaptation Rule for Queues	44
Queue- Committed Burst Size	45
Queue- Maximum Burst Size	45
Queue Counters	47
Queue – support on various platforms	47
Service Ingress QoS Policies	48
Service Ingress QoS policies	48
Default Service Ingress Policy	49
Service Ingress Classification	50
Service Ingress Classification	50

Table of Contents

Service Ingress Classification– IP and Mac Packet fields.....	51
Service Egress QoS Policies	54
Default Service Egress Policy	55
Buffer Pools.....	55
Buffer pools on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C	56
Ring and Non-Ring Buffer pool	57
Configuration Guidelines for CBS and MBS	60
Slope Policies.....	60
RED Slopes.....	61
Operation and Configuration of RED Slopes	61
Simplified overview of RED for 7210 SAS-K platform.....	61
Slope Policy Parameters	63
Schedulers.....	64
Scheduler on 7210 SAS-K	64
CPU Queues on SAS-K.....	65
.....	65
Egress Port Rate Limiting.....	65
.....	65
Forwarding Classes.....	66
.....	66
Forwarding-Class To Queue ID Mapping.....	67
FC to Queue ID mapping	67
Preclassification on 7210 SAS-K 2F4T6C.....	67
QoS Policy Entities	68
Configuration Notes.....	68
Discard Eligibility Indicator (DEI) based Classification and Marking.....	71
In This Section.....	71
DEI based Classification and Marking	71
DEI based Classification on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.....	71
DEI based marking	72
Port Level Egress Rate-Limiting	73
In This Section.....	73
Overview.....	73
Basic Configurations.....	74
Modifying Port Level Egress-Rate Command	75
Removing Port Level Egress-Rate Command	75
Default Egress-Rate Values.....	75
Port Level Egress-Rate Command Reference	77
Command Hierarchies.....	77
Configuration Commands for 7210 SAS-K 2F2T1C devices	77
Show Commands.....	77
Configuration Descriptions	77
Configuration Commands	77
Show Commands.....	78
Frame Based Accounting	81
In This Section.....	81

Overview.....	81
Frame Based Accounting	81
Enabling and Disabling Frame-Based Accounting	82
Enabling and Disabling Frame-Based Accounting on 7210 SAS-K	83
Configuration Commands for 7210 SAS-K 2F4T6C	84
Frame Based Accounting Command Reference	85
Configuration Descriptions	85
Show Commands	85
DSCP, Dot1p, and, MPLS EXP classification policy.....	89
In This Section.....	89
Overview.....	89
DSCP classification policy	90
Dot1p classification policy	90
MPLS EXP classification policy.....	90
Configuration of DSCP classification policy, Dot1p classification policy, and MPLS EXP classification policy	90
Command reference.....	93
Configuration Commands for 7210 SAS-K 2F2T1C.....	93
Configuration Commands for 7210 SAS-K 2F4T6C.....	93
Show Commands.....	94
Operational Commands	94
Command Description.....	95
DSCP classification policy, Dot1p classification policy, and MPLS EXP classification policy	99
Configuration Descriptions	99
Show Commands.....	99
Network QoS Policies	101
In This Section.....	101
Overview of Network QoS policy on 7210 SAS-K 2F2T1C	101
Resource Allocation for Network QoS policy for 7210 SAS-K 2F2T1C	103
Basic Configurations.....	103
Create a Network QoS Policy on 7210 SAS-K 2F2T1C.....	103
Default Network Policy Values on 7210 SAS-K 2F2T1C	105
Default Network Policy Values	106
7210 SAS-K 2F4T6C Network QoS Policy.....	108
Overview	108
Network QoS policy for access-uplink port on 7210 SAS-K 2F4T6C.....	109
Network QoS policy for network port on 7210 SAS-K 2F4T6C	110
Resource Allocation for Network QoS policy for 7210 SAS-K 2F4T6C.....	111
Basic Configurations	112
Create a Network QoS Policy on 7210 SAS-K 2F4T6C.....	112
Create a Network QoS Policy for access-uplink ports on 7210 SAS-K 2F4T6C.....	112
Create a Network QoS Policy for network ports on 7210 SAS-K 2F4T6C	114
Default Network Policy Values on 7210 SAS-K 2F4T6C	117
DSCP Marking CPU Generated Traffic.....	119
Default DSCP Mapping	120
Service Management Tasks	121
Deleting QoS Policies.....	121

Table of Contents

Remove a Policy from the QoS Configuration	121
Copying and Overwriting Network Policies	122
Editing QoS Policies	122
Network QoS Policy Command Reference	123
Command Hierarchies.....	123
Configuration Commands for 7210 SAS-K 2F2T1C	123
Configuration Commands for 7210 SAS-K 2F4T6C	124
Operational Commands	125
Show Commands	125
Network QoS Policy Command Descriptions	127
Configuration Commands.....	127
Generic Commands	127
Operational Commands	127
Network QoS Policy Commands	131
Network Ingress QoS Policy Commands	133
Network Ingress Queue QoS Policy Commands	137
Show Commands	141
Network QoS Policy Output Fields	144
Network Queue QoS Policies	149
In This Section.....	149
Overview.....	149
Basic Configurations.....	149
Create a Network Queue QoS Policy	150
Applying Network Queue Policies	151
Applying Network Queue Configuration to access-uplink port	151
Applying Network Queue Configuration to network port on 7210 SAS-K 2F4T6C.....	152
Default Network Queue Policy Values.....	152
Default Network Queue policy for 7210 SAS-K 2F2T1C	152
Default Network Queue policy for 7210 SAS-K 2F4T6C	154
Service Management Tasks	156
Deleting Network Queue QoS Policies.....	156
Copying and Overwriting Network Queue QoS Policies	157
Editing Network Queue QoS Policies	158
Network Queue QoS Policy Command Reference.....	159
Command Hierarchies.....	159
Configuration Commands for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.....	159
Operational Commands	159
Show Commands	160
Network Queue Command Descriptions	161
Configuration Commands.....	161
Generic Commands	161
Operational Commands	161
Network Queue QoS Policy Commands	162
Network Queue QoS Policy Queue Commands	163
Show Commands	167
Service Ingress QoS Policies	171
In This Section.....	171

Overview of service ingress policy	171
Configuration Guidelines for SAP ingress Policy	172
Resource Allocation for Service Ingress QoS classification policy.....	173
Default SAP Ingress Policy	174
SAP Ingress Policy Defaults	175
Use of Index file by SAP QoS Ingress policy	176
Basic Configurations	176
Create Service Ingress QoS Policies	176
Service Ingress QoS Policy of 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C	176
Service Ingress QoS Queues for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C	177
SAP Ingress Forwarding Class Configuration for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C	178
Service Ingress Dot1p and IP DSCP Criteria	179
Service Ingress IP Match Criteria	179
Service Ingress MAC Match Criteria	180
Applying Service Ingress Policies	181
Service Management Tasks	182
Deleting QoS Policies.....	182
Remove a QoS Policy from Service SAP(s)	183
Copying and Overwriting QoS Policies	183
Remove a Policy from the QoS Configuration.....	184
Editing QoS Policies.....	184
Service SAP QoS Policy Command Reference	185
Command Hierarchy	185
Service Ingress QoS Policy Command for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C	185
Operational Commands	187
Show Commands	187
Command Descriptions	189
Configuration Commands.....	189
Generic Commands	189
Operational Commands	189
Service Ingress QoS Policy Commands	194
Service Ingress QoS Policy Entry Commands.....	201
Service Ingress MAC QoS Policy Match Commands.....	205
Show Commands	209
.....	209
Service Egress Policies	217
In This Section.....	217
Overview.....	217
Configuration Guidelines for access SAP Egress Policy.....	219
Basic Configurations	220
Create an access SAP Egress Policy	220
Editing QoS Policies.....	221
Service Egress Policy Command Reference.....	223
Command Hierarchies.....	223
Configuration Commands for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.....	223
Copy Commands	223
Show Commands	224
Commands Description	225

Table of Contents

Configuration Commands.....	225
Generic Commands	225
SAP Egress Queue QoS Policy Commands	225
Operational Commands	234
Show Commands	235
Schedulers	237
In This Section.....	237
Overview	237
Slope QoS Policies.....	241
In This Section.....	241
Overview of Buffer pools and Slope policies	241
Basic Configurations.....	242
Create a Slope QoS Policy.....	242
Applying Slope Policies	243
Applying Slope Policies on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C	243
Default Slope Policy Values for 7210 SAS-K devices	244
Deleting QoS Policies.....	246
Remove a Policy from the QoS Configuration.....	246
Copying and Overwriting QoS Policies.....	247
Editing QoS Policies	247
Slope QoS Policy Command Reference	249
Command Hierarchies.....	249
Configuration commands	249
Operational Commands	250
Show Commands.....	250
Command Descriptions	251
Configuration Commands.....	251
Generic Commands	251
Operational Commands	251
.....	251
Slope Policy QoS Commands.....	252
Slope Policy QoS Policy Commands	252
Slope Policy QoS Policy Commands (for SAS-K devices).....	255
RED Slope Commands	255
Show Commands.....	257
Remark Policies for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C	261
In This Section.....	261
Overview	261
Overview of remark policies for 7210 SAS-K 2F2T1C	261
Configuration Guideline for 7210 SAS-K 2F2T1C.....	263
Overview of remark policies for 7210 SAS-K 2F4T6C	264
Configuration Guideline for 7210 SAS-K 2F4T6C.....	267
Basic Configurations on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.....	268
Creating a Remark Policy	268
Editing QoS Policies.....	269
Remark Policy Command Reference	271

- Command Hierarchies271
 - Configuration Commands for 7210 SAS-K 2F2T1C271
 - Configuration Commands for 7210 SAS-K 2F4T6C271
 - Show Commands272
 - Operational Commands272
- Command Descriptions273
 - Configuration Commands273
 - Generic Commands273
 - Operational Commands273
 - 273
 - Remark Policy QoS Commands274
 - Remark Policy Forwarding Class Commands276
 - Show Commands281
 - 281

Standards and Protocol Support1

List of Figures

QoS Policies	19
Figure 1 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C Service and Network Traffic Types and QoS model 31	
Figure 2 Traffic Queuing Model for Forwarding Classes.....	49
Figure 3 RED Slope Characteristics	62
Schedulers	237
Figure 4 Scheduler in 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C	238
Standards and Protocol Support	1

List of Figures

List of Tables

Getting Started	17
Table 1 Configuration Process	17
QoS Policies	19
Table 2 QoS Policy Types and Descriptions for 7210 SAS-K 2F2T1C	27
Table 3 QoS Policy Types and Descriptions for 7210 SAS-K 2F4T6C	29
Table 4 Default Network QoS Policy Egress Marking on 7210 SAS-K 2F4T6C.....	34
Table 5 Default Network QoS policy ID #2 - ingress classification map - Dot1p Mapping to forwarding class 34	
Table 6 Default Network QoS policy mpls-lsp-exp-classification to FC Mapping	35
Table 7 Default Network QoS policy dscp-classification to FC Mapping	35
Table 8 Default Network QoS policy Dot1p to FC Mapping for network egress	37
Table 9 Default Network QoS policy Dot1p to FC Mapping for network ingress	38
Table 10 Default Network Queue Policy Definition on 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C.	39
Table 11 Default Service Ingress Policy ID 1 Definition for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C50	
Table 12 Service Ingress QoS Policy IP Match Criteria for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C51	
Table 13 Service Ingress QoS Policy IPv6 Criteria for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.....	52
Table 14 Service Ingress QoS Policy MAC Match Criteria for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C 53	
Table 15 MAC Match Ethernet Frame Types	53
Table 16 MAC Match Criteria Frame Type Dependencies	54
Table 17 Default Service Egress Policy ID 1 Definition	55
Table 18 Default slope policy definition for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.....	63
Table 19 Forwarding Classes	66
DSCP, Dot1p, and, MPLS EXP classification policy	89
Table 20 Classification policy with platforms support.	89
Network QoS Policies	101
Table 21 Network Policy ID #1 Defaults for access-uplink port	106
Table 22 Default Network QoS Policy ID #1 for Dot1p to FC Mapping for access-uplink port	108
Table 23 DSCP and Dot1p Marking.....	119
Table 24 Show QoS Network Output Fields	142
Network Queue QoS Policies	149
Table 25 Network Queue Labels and Descriptions.....	168
Service Ingress QoS Policies	171
Table 26 SAP Ingress Policy Defaults	175
Slope QoS Policies	241
Table 27 Default Slope Policy for 7210 SAS-K.....	244
Table 28 Show QoS Slope Policy Output Fields	257

List of Tables

Remark Policies for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C	261
Table 29 Summary of remark policy and attachment points for 7210 SAS-K 2F2T1C	262
Table 30 Summary of remark policy and attachment points for 7210 SAS-K 2F4T6C	265
Standards and Protocol Support	1

Preface

About This Guide

This guide describes Quality of Service (QoS) provided by the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C platforms and presents examples to configure and implement various tests and presents examples to configure and implement various protocols and services.

On 7210 SAS devices, not all the CLI commands are supported on all the platforms and in all the modes. In many cases, the CLI commands are mentioned explicitly in this document. In other cases, it is implied and easy to know the CLIs not supported on a particular platform.

This document is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

Audience

This manual is intended for network administrators who are responsible for configuring the 7210 SAS-Series routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Protocols, standards, and services described in this manual include the following:

- CLI concepts
- Quality of Service (QoS) policies and profiles

List of Technical Publications

The 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OS documentation set is composed of the following books:

- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OS Basic System Configuration Guide

About This Guide

- This guide describes basic system configurations and operations.
- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OS System Management Guide
- This guide describes system security and access configurations as well as event logging and accounting logs.
- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OS Interface Configuration Guide
- This guide describes card, Media Dependent Adapter (MDA), link aggregation group (LAG), and port provisioning.
- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OS Router Configuration Guide
- This guide describes logical IP routing interfaces and associated attributes such as an IP address, port, as well as IP and MAC-based filtering.
- 7210 SAS-D and 7210 SAS-E OS Services Guide
7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C OS Services Guide
- This guide describes how to configure service parameters such as customer information, and user services.
- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OS OAM and Diagnostic Guide
- This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
- 7210 SAS-D and 7210 SAS-E OS Quality of Service Guide
7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C OS Quality of Service Guide
- This guide describes how to configure Quality of Service (QoS) policy management.
- 7210 SAS-K 2F4T6C OS MPLS Guide
This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
- 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C-OS Routing Protocols Guide
This guide provides an overview of routing concepts and provides configuration examples for OSPF, IS-IS and route policies.

Getting Started

In This Chapter

This chapter provides process flow information to configure Quality of Service (QoS) policies and provision services.

Nokia 7210 SAS-Series Services Configuration Process

[Table 1](#) lists the tasks necessary to configure and apply QoS policies. This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 1: Configuration Process

Area	Task	Chapter
Policy configuration	Configuring QoS Policies	
	<ul style="list-style-type: none">Egress Rate	Port Level Egress Rate-Limiting
	<ul style="list-style-type: none">Accounting Mode	Frame Based Accounting
	<ul style="list-style-type: none">Network	Network QoS Policies
	<ul style="list-style-type: none">Network queue	Network Queue QoS Policies
	<ul style="list-style-type: none">Service Ingress	Service Ingress QoS Policies
	<ul style="list-style-type: none">Service Egress	Service Egress Policies
	<ul style="list-style-type: none">Schedulers	Schedulers
	<ul style="list-style-type: none">Remark Policies	Remark Policies for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Table 1: Configuration Process (Continued)

Area	Task	Chapter (Continued)
Reference	<ul style="list-style-type: none"><li data-bbox="651 394 984 457">• List of IEEE, IETF, and other proprietary entities	Standards and Protocol Support

In This Chapter

This chapter provides information about Quality of Service (QoS) policy management.

Topics in this chapter include:

- [QoS Overview](#)
 - [Overview of QoS Policies on 7210 SAS-K 2F2T1C](#)
 - [QoS marking interpretation](#)
 - [Summary of major functions of QoS policies](#)
 - [Network QoS Policies](#)
 - [Network QoS policies on Access-uplink ports](#)
 - [Queues and Queue Parameters](#)
 - [Service Ingress QoS Policies](#)
 - [Service Ingress Classification](#)
 - [Service Egress QoS Policies](#)
 - [Buffer Pools](#)
 - [Slope Policies](#)
 - [RED Slopes](#)
 -
- [Schedulers](#)
 - [Scheduler on 7210 SAS-K](#)
- [CPU Queues](#)
 - [Configuration Notes](#)

QoS Overview

The 7210 SAS devices are designed with QoS mechanisms on both ingress and egress to support multiple services per physical port. The 7210 SAS devices are extensive and flexible capabilities to Classify, Policy, Queue, Shape, and mark traffic.



Note: The QoS capabilities supported on different 7210 SAS platforms are different. In other words, not all the platforms support all of the capabilities. Please read through the following chapters to know what is available on different 7210 SAS platforms.

In the Nokia service router's service model, a service is provisioned on the provider-edge (PE) equipment. Service data is encapsulated and then sent in a service tunnel (for example: QinQ tunnel, Dot1q tunnel, IP/MPLS tunnel, etc.) to the far-end Nokia service router where the service data is delivered.

The operational theory of a service tunnel is that the encapsulation of the data between the two Nokia service routers appear like a Layer 2 path to the service data although it is really traversing an QinQ or IP or IP/MPLS core. The tunnel from one edge device to the other edge device is provisioned with an encapsulation and the services are mapped to the tunnel that most appropriately supports the service needs.

The 7210 SAS supports eight forwarding classes internally named: Network-Control, High-1, Expedited, High-2, Low-1, Assured, Low-2 and Best-Effort. The forwarding classes are discussed in more detail in .

7210 SAS devices use QoS policies to control how QoS is handled at distinct points in the service delivery model within the device. There are different types of QoS policies that cater to the different QoS needs at each point in the service delivery model. QoS policies are defined in a global context in the 7210 SAS and only take effect when the policy is applied to a relevant entity.

QoS policies are uniquely identified with a policy ID number or name. Typically, Policy ID 1 or Policy ID "default" (there are a few instances where the default QoS policy uses a different ID) is reserved for the default policy which is used if no policy is explicitly applied.

The QoS policies within the 7210 SAS can be divided into three main types:

- Policies are used for classification, defining metering and queuing attributes and defining marking behavior.
- Slope policies define default buffer allocations and WRED slope definitions.
- Port Scheduler policies, SAP ingress/egress policies and network/network-queue policies determine how queues are scheduled.

Overview of QoS Policies on 7210 SAS-K 2F2T1C

On 7210 SAS-K2F2T1C, QoS policies are applied on service ingress, service egress and access uplink ports (ingress and egress) and define the following:

- Classification rules for how traffic is mapped to forwarding classes
- Forwarding class association with queues.
- Queue parameters for shaping, scheduling and buffer allocation
- QoS marking/interpretation

There are several types of QoS policies:

- Service ingress (for access SAP ingress)
- Service egress (for access SAP egress)
- Network (for access-uplink port ingress and egress)
- Network queue (for access-uplink port egress)
- Slope policies (for all queues)
- Remark policies (for both access SAP egress and access-uplink port egress)
- Dot1p and DSCP classification policies (for access SAP ingress and access-uplink port ingress).

Service ingress QoS policies are applied to the customer-facing Service Access Points (SAPs) on access ports. Traffic that enters through the access SAP is classified to map it to a Forwarding Class (FC) and the user has an option to also assign a profile on SAP ingress. Forwarding class is associated with queues on access SAP ingress and the profile determines the enqueueing priority for the packet, with in-profile packets having a higher chance of getting a buffer, than out-of-profile packets. The mapping of traffic to queues can be based on combinations of customer QoS marking in the packet header (for example: IEEE 802.1p bits, IP DSCP bits, MAC address, etc). The characteristics of the forwarding class queues are defined within the policy as to the number of forwarding class queues to use for unicast traffic and BUM (Broadcast, Unknown-unicast, and Multicast) traffic along with the queue rate and buffer parameters (like CIR, PIR, CBS, MBS). Each of the forwarding classes can be associated with different parameters for unicast traffic and different parameters for multi-point (that is, BUM) traffic. A service ingress QoS policy defines up to 8 queues per policy, with up to 2 queues (that is, Unicast Queue Mapping and Multicast Queue Mapping) per forwarding class. Unicast and multi-point traffic can be mapped to use the same queue or mapped to use different queues per forwarding class with a maximum of up to 2 queues per forwarding class, one each for unicast and for multicast traffic. In the case of VPLS service, four types of forwarding are supported (which is not to be confused with forwarding classes),

unicast, multicast, broadcast, and unknown. Multicast, broadcast, and unknown types are flooded to all destinations within the service while the unicast forwarding type is handled in a point-to-point fashion within the service. All these traffic types use the same multicast-queue mapping defined for forwarding class. In other words, a separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined.

Service egress QoS policies are applied to SAPs and map forwarding classes to service egress queues for a service. The system can allocate a maximum of 8 queues per SAP for the 8 forwarding classes. All traffic types (that is, both unicast and BUM traffic types) share the same queue on service egress. A service egress QoS policy defines the forwarding class queue characteristics and also defines how to remark the forwarding class to priority bits in the packet header (for example: IEEE 802.1p bits in the Ethernet VLAN tag) in the customer traffic.

Network QoS policies are applied to access uplink ports. Access-uplink ports are typically used to connect to the core network and forward customer traffic towards the core network. A network QoS policy defines both ingress and egress behavior. On access-uplink port ingress, traffic that enters through the access-uplink port is classified to map it to a forwarding Class and the user has an option to assign a profile. Forwarding class is associated with ingress queues on access uplink port ingress and the profile determines the enqueueing priority for the packet, with in-profile packets have a higher chance of getting the buffer, than out-of-profile packets. The mapping of traffic to forwarding class queues is based on QoS marking (for example: IEEE 802.1p bits, IP DSCP bits). The characteristics of the forwarding class ingress queues are defined within the policy as to the number of forwarding class queues for unicast traffic type and BUM (Broadcast, Unknown-unicast, and Multicast) traffic type, along with the queue rate and buffer parameters (like CIR, PIR, CBS, MBS, etc.). Each of the forwarding classes can be associated with different ingress and ingress queue parameters for unicast traffic type and for multi-point (that is, BUM) traffic type. A network QoS policy defines up to 8 ingress queues per policy, with up to 2 ingress queues per forwarding class. Unicast and multi-point traffic can be defined to use the same queue or different ingress queues per forwarding class. In the case of VPLS service, four types of forwarding are supported (which is not to be confused with forwarding classes), unicast, multicast, broadcast, and unknown. Multicast, broadcast, and unknown types are sent to multiple destinations within the service while the unicast forwarding type is handled in a point-to-point fashion within the service. All these traffic types use the same queue (in other words, a separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined). On access-uplink port egress, the policy maps forwarding class and profile state to Dot1p and /or IP DSCP values for traffic to be transmitted out of the access-uplink port. All the access-uplink SAPs configured on the same access-uplink port use the same policy and the same set of forwarding class queues. In other words, traffic received and transmitted through all the access-uplink SAPs configured on a given access-uplink port receive similar QoS treatment.

Network queue policies are applied on egress of access uplink ports and map forwarding classes to network egress queues on access uplink ports. The system allocates 8 egress queues per access-uplink port for the 8 forwarding classes. The policy defines the forwarding class queue characteristics (that is, CIR, PIR, CBS, MBS, etc.). All traffic types (that is, both unicast and BUM traffic types) share the same queue on access-uplink port egress. All the access-uplink SAPs configured on the same access-uplink port use the same policy and the same set of forwarding class queues. In other words, traffic transmitted through all the access-uplink SAPs configured on a given access-uplink port receive similar QoS treatment.

Slope policies are applied to service ingress queues, service egress queues, access uplink port ingress, queues, and access uplink port egress queues. Each of these queuing points allocates buffers from the buffer pool and implements WRED for congestion management. During congestion WRED is used to evaluate how buffers from the pool are allocated to different FCs and to in-profile and out-of-profile traffic within a given FC. The slope policies define the WRED parameters to use for in-profile/high-priority packets and for out-of-profile/low-priority packets. The high-slope and low-slope define the parameters for in-profile/high-priority packets and for out-of-profile/low-priority packets respectively. In addition, 7210 SAS-K introduces the concept of ring and non-ring ports with an option for preferential allocation of traffic for ring ports. The system by default treats access-uplink ports as ring ports.

Remark policies are applied to access SAP egress and access-uplink port egress. They are not directly associated with the SAP and access-uplink port egress. Instead they are associated with service egress policy and network qos policy. They define the forwarding class and profile to egress marking values (for example: IEEE 802.1p bits in the Ethernet VLAN tag) to use.

Dot1p classification and DSCP classification allows user to define the map of Dot1p bits and IP DSCP values to forwarding class and assign the profile for the packet on access SAP ingress and access-uplink port ingress.

One service ingress QoS policy and a single service egress QoS policy can be applied to a specific access SAP. One network QoS can be applied to a specific access-uplink port. One network queue policy can be applied to the access uplink port. If no QoS policy is explicitly applied to a SAP or port, a default QoS policy is applied.

Overview of QoS Policies on 7210 SAS-K 2F4T6C

On 7210 SAS-K2F4T6C QoS policies are applied on service ingress, service egress and access uplink ports (ingress and egress) and define the following:

- Classification rules for how traffic is mapped to forwarding classes.
- Forwarding class association with queues.

- Queue parameters for shaping, scheduling and buffer allocation

QoS marking interpretation

There are several types of QoS policies:

- Service ingress (for access SAP ingress)
- Service egress (for access SAP egress)
- Network (for access-uplink port ingress and egress and network port ingress and egress)
- Network queue (for access-uplink port egress and network port egress)
- Slope policies (for all queues)
- Remark policies (for both access SAP egress, access-uplink port egress and network port egress)
- Dot1p, IP DSCP and MPLS EXP classification policies (for access SAP ingress, access-uplink port ingress and network port ingress) Please note, that MPLS EXP classification policies are only for network ports.

Service ingress QoS policies are applied to the customer-facing Service Access Points (SAPs) on access ports. Traffic that enters through the SAP is classified to map it to a Forwarding Class (FC) and the user has an option to also assign a profile on SAP ingress. Forwarding class is associated with queues on SAP ingress and the profile determines the enqueueing priority for the packet, with in-profile packets having a higher chance of getting a buffer, than out-of-profile packets. The mapping of traffic to queues can be based on combinations of customer QoS marking in the packet header (for example: IEEE 802.1p bits, IP DSCP bits, MAC address, etc.). The characteristics of the forwarding class queues are defined within the policy as to the number of forwarding class queues to use for unicast traffic and BUM (Broadcast, Unknown-unicast, and Multicast) traffic along with the queue rate and buffer parameters (like CIR, PIR, CBS, MBS). Each of the forwarding classes can be associated with different parameters for unicast traffic and different parameters for multi-point (that is, BUM) traffic. A service ingress QoS policy defines up to 8 queues per policy, with up to 2 queues (that is, Unicast Queue Mapping and Multicast Queue Mapping) per forwarding class. Unicast and multi-point traffic can be mapped to use the same queue or mapped to use different queues per forwarding class with a maximum of up to 2 queues per forwarding class, one each for unicast and for multicast traffic. In the case of VPLS service, four types of forwarding are supported (which is not to be confused with forwarding classes), unicast, multicast, broadcast, and unknown. Multicast, broadcast, and unknown types are flooded to all destinations within the service while the unicast forwarding type is handled in a point-to-point fashion within the service. All these traffic types use the same multicast-queue mapping defined for forwarding class. In other words, a separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined.

Service egress QoS policies are applied to SAPs and map forwarding classes to service egress queues for a service. The system can allocate a maximum of 8 queues per SAP for the 8 forwarding classes. All traffic types (that is, both unicast and BUM traffic types) share the same queue on service egress. A service egress QoS policy defines the forwarding class queue characteristics and also defines how to remark the forwarding class to priority bits in the packet header (for example: IEEE 802.1p bits in the Ethernet VLAN tag) in the customer traffic.

Network QoS policies are applied to access uplink ports or network ports. Following provides an overview of the network QoS policy applied to access-uplink ports followed by an overview of the network QoS policy applied to network ports.

Access-uplink ports are typically used to connect to the core network using QinQ or Dot1q links and forward customer traffic towards the core network. A network QoS policy defines both ingress and egress behavior on a access-uplink port. On access-uplink port ingress, traffic that enters through the port is classified to map it to a forwarding Class and the user has an option to assign a profile. Forwarding class is associated with ingress queues on access uplink port ingress and the profile determines the en-queuing priority for the packet, with in-profile packets have a higher chance of getting the buffer, than out-of-profile packets. The mapping of traffic to forwarding class ingress queues is based on QoS marking (for example: IEEE 802.1p bits, IP DSCP bits). The characteristics of the forwarding class ingress queues are defined within the policy as to the number of forwarding class queues for unicast traffic type and BUM (Broadcast, Unknown-unicast, and Multicast) traffic type, along with the queue rate and buffer parameters (like CIR, PIR, CBS, MBS, etc.). Each of the forwarding classes can be associated with different ingress queue parameters for unicast traffic type and for multi-point (that is BUM) traffic type. A network QoS policy defines up to 8 ingress queues per policy, with up to 2 ingress queues per forwarding class. Unicast and multi-point traffic can be defined to use the same ingress queue or different ingress queues per forwarding class. In the case of VPLS service, four types of forwarding are supported (which is not to be confused with forwarding classes), unicast, multicast, broadcast, and unknown. Multicast, broadcast, and unknown types are sent to multiple destinations within the service while the unicast forwarding type is handled in a point-to-point fashion within the service. All these traffic types use the same queue (in other words, a separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined). On access-uplink port egress, the policy maps forwarding class and profile state to Dot1p and /or IP DSCP values for traffic to be transmitted out of the access-uplink port. All the access-uplink SAPs configured on the same access-uplink port use the same policy and the same set of forwarding class queues. In other words, traffic received and transmitted through all the access-uplink SAPs configured on a given access-uplink port receive similar QoS treatment.

Network ports are typically used to connect to the core network using IP/MPLS tunnels and forward customer traffic towards the core network. A network QoS policy defines both ingress and egress behavior on a network port. On network port ingress, traffic that enters through the port is classified to map it to a forwarding Class and the user has an option to assign a profile. Forwarding class is associated with ingress queues on network port ingress and the profile determines the en-queuing priority for the packet, with in-profile packets have

a higher chance of getting the buffer, than out-of-profile packets. The mapping of traffic to forwarding class ingress queues is based on QoS marking (for example: MPLS EXP bits, IEEE 802.1p bits, IP DSCP bits). The characteristics of the forwarding class ingress queues are defined within the policy as to the number of forwarding class queues for unicast traffic type and BUM (Broadcast, Unknown-unicast, and Multicast) traffic type, along with the queue rate and buffer parameters (like CIR, PIR, CBS, MBS, etc.). Each of the forwarding classes can be associated with different ingress and ingress queue parameters for unicast traffic type and for multi-point (that is BUM) traffic type. A network QoS policy defines up to 8 ingress queues per policy, with up to 2 ingress queues per forwarding class. Unicast and multi-point traffic can be defined to use the same ingress queue or different ingress queues per forwarding class. In the case of VPLS service, four types of forwarding are supported (which is not to be confused with forwarding classes), unicast, multicast, broadcast, and unknown. Multicast, broadcast, and unknown types are sent to multiple destinations within the service while the unicast forwarding type is handled in a point-to-point fashion within the service. All these traffic types use the same queue (in other words, a separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined). On network port egress, the policy maps forwarding class and profile state to MPLS EXP and/or Dot1p and/or IP DSCP values for traffic to be transmitted out of the network port. All the IP interfaces configured on the same network port use the same policy and the same set of forwarding class queues. In other words, traffic received and transmitted through all the IP interfaces configured on a given network port receive similar QoS treatment.

Network queue policies are applied on egress of access uplink ports and network ports. It maps forwarding classes to network egress queues on access uplink ports and on network ports. The system allocates 8 egress queues per port (per network port and per access-uplink port) for the 8 forwarding classes. The policy defines the forwarding class queue characteristics (that is, CIR, PIR, CBS, MBS, etc.). All traffic types (that is, both unicast and BUM traffic types) share the same egress queue on access-uplink port and network port. All the access-uplink SAPs configured on the same access-uplink port use the same policy and the same set of forwarding class queues. In other words, traffic transmitted through all the access-uplink SAPs configured on a given access-uplink port receive similar QoS treatment. All the IP interfaces configured on the same network port use the same policy and the same set of forwarding class queues. In other words, traffic transmitted through all the IP interfaces configured on a given network port receive similar QoS treatment.

Slope policies are applied to service ingress queues, service egress queues, access uplink port ingress queues, access uplink port egress queues, network port ingress queues, network port egress queues. Each of these queuing points allocates buffers from the buffer pool and implements WRED for congestion management. During congestion WRED is used to evaluate how buffers from the pool are allocated to different FCs and to in-profile and out-of-profile traffic within a given FC. The slope policies define the WRED parameters to use for in-profile/high-priority packets and for out-of-profile/low-priority packets. The high-slope and low-slope define the parameters for in-profile/high-priority packets and for out-of-profile/low-priority packets respectively. In addition, 7210 SAS-K introduces the concept of ring and non-ring ports with an option for preferential allocation of traffic for ring ports. The system by default treats network ports and access-uplink ports as ring ports.

Remark policies are applied to access SAP egress, access-uplink port egress and network port egress. They are associated with service egress policy and network qos policy. They define the forwarding class and profile to egress marking values (for example: IEEE 802.1p bits in the Ethernet VLAN tag) to use.

Dot1p classification policy, IP DSCP classification policy and MPLS EXP classification policy allows user to define the map of Dot1p bits, IP DSCP values and MPLS EXP values respectively to forwarding class and assign the profile for the packet. Dot1p classification policy and IP DSCP classification policy is available on access SAP ingress and access-uplink port ingress. Dot1p classification policy, IP DSCP classification policy, and MPLS EXP classification policy is available on access SAP ingress and access-uplink port ingress, and network port ingress.

One service ingress QoS policy and a single service egress QoS policy can be applied to a specific access SAP. One network QoS can be applied to a specific access-uplink port. One network queue policy can be applied to the network port. If no QoS policy is explicitly applied to a SAP or port, a default QoS policy is applied.

Summary of major functions of QoS policies

A summary of the major functions performed by the QoS policies is listed in [Table 2](#).

Table 2: QoS Policy Types and Descriptions for 7210 SAS-K 2F2T1C

Policy Type	Applied at...	Description	Page
Service Ingress	Access SAP Ingress	<ul style="list-style-type: none"> Defines up to 8 forwarding class queues and queue parameters to define queue characteristics (For example: scheduler priority and weight, rates, etc.) For traffic classification, defines match criteria to map flows to the queues based on Dot1p or IP DSCP criteria or MAC criteria or IP criteria. 	65
Service Egress	Access SAP Egress	<ul style="list-style-type: none"> Allocates up to 8 forwarding class queues and maps forwarding classes to the queues. Defines FC to remarking values, through the use of remark policies. Defines queue parameters to define queue characteristics (For example: scheduler priority and weight, rates, etc.). 	67

Table 2: QoS Policy Types and Descriptions for 7210 SAS-K 2F2T1C (Continued)

Policy Type	Applied at...	Description	Page
Egress Rate	Access port and Access-uplink port	Configures the maximum bandwidth available for traffic sent out of a specified port.	79
Network QoS	Access-uplink port	<ul style="list-style-type: none"> At ingress, defines up to 8 forwarding class queues and queue parameters to define queue characteristics (For example: scheduler priority and weight, rates, etc.). For traffic classification, defines match criteria to map flows to the queues based on Dot1p and DSCP values. At egress, defines FC to remarking values, through the use of remark policies. 	34
Network Queue	Access-uplink port	<ul style="list-style-type: none"> Allocates up to 8 forwarding class queues and maps forwarding classes to the queues. Defines queue parameters to define queue characteristics (For example: scheduler priority and weight, rates, etc.). 	43
Slope policies	SAP queues (both ingress and egress) and Access-uplink port (both ingress and egress queues)	<ul style="list-style-type: none"> Enables or disables the high-slope and low-slope parameters for queue. On 7210 SAS-K 2F2T1C, in addition to high-slope and low-slope, user has an option to use high-slope-ring and low-slope-ring parameters for access-uplink port egress queues. 	
Remark Policies	SAP egress, Access-uplink port egress	<ul style="list-style-type: none"> Defines FC to remarking values; Not directly associated with a SAP or a port. Instead it is associated with SAP egress policy and network qos policy. 	109
Dot1p classification policy and DSCP classification policy	Access SAP ingress and access-uplink port ingress	<ul style="list-style-type: none"> Defines the map of Dot1p bits and IP DSCP values to forwarding class and assign the profile for the packet on access SAP ingress and access-uplink port ingress. 	

Table 3: QoS Policy Types and Descriptions for 7210 SAS-K 2F4T6C

Policy Type	Applied at...	Description	Page
Service Ingress	Access SAP Ingress	<ul style="list-style-type: none"> Defines up to 8 forwarding class queues and queue parameters to define queue characteristics (For example: scheduler priority and weight, rates, etc.) For traffic classification, defines match criteria to map flows to the queues based on Dot1p or IP DSCP criteria or MAC criteria or IP criteria. 	65
Service Egress	Access SAP Egress	<ul style="list-style-type: none"> Allocates up to 8 forwarding class queues and maps forwarding classes to the queues. Defines FC to remarking values, through the use of remark policies. Defines queue parameters to define queue characteristics (For example: scheduler priority and weight, rates, etc.). 	67
Egress Rate	Access port, Access-uplink port, and network port	Configures the maximum bandwidth available for traffic sent out of a specified port.	79
Network QoS	Network port and Access-uplink port	<ul style="list-style-type: none"> At ingress, defines up to 8 forwarding class queues and queue parameters to define queue characteristics (For example: scheduler priority and weight, rates, etc.). For traffic classification, defines match criteria using packet header bits (e.g. MPLS EXP, Dot1p, IP DSCP, etc.) to map flows to the queues. At egress, defines FC to remarking values, through the use of remark policies. 	34
Network Queue	Network port and Access-uplink port	<ul style="list-style-type: none"> Allocates up to 8 forwarding class queues and maps forwarding classes to the queues. Defines queue parameters to define queue characteristics (For example: scheduler priority and weight, rates, etc.). 	43

Table 3: QoS Policy Types and Descriptions for 7210 SAS-K 2F4T6C (Continued)

Policy Type	Applied at...	Description	Page
Slope policies	Access SAP queues (both ingress and egress), network port queues (both ingress and egress) and Access-uplink port (both ingress and egress queues)	<ul style="list-style-type: none"> Enables or disables the high-slope and low-slope parameters for queue. On 7210 SAS-K 2F4T6C, in addition to high-slope and low-slope, user has an option to use high-slope-ring and low-slope-ring parameters for network port and access-uplink port egress queues. 	
Remark Policies	Access SAP egress, network port egress, Access-uplink port egress	<ul style="list-style-type: none"> Defines FC to remarking values 	109
MPLS EXP classification policy, Dot1p classification policy and DSCP Classification policy	Network port ingress, Access SAP ingress and access-uplink port ingress	<ul style="list-style-type: none"> Defines the map of Dot1p bits and IP DSCP values to forwarding class and assign the profile for the packet received on access SAP ingress and access-uplink port. Defines the map of MPLS EXP, Dot1p, and IP DSCP values to forwarding class and assign the profile for packets received on network port. 	

Service and Network QoS Policies

The QoS mechanisms within the 7210 SAS- K are specialized for the type of traffic on the interface.

On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, for customer interfaces, there is service ingress and service egress, and for access uplink port, and network ports, there is network ingress and network egress traffic ([Figure 1](#) below).

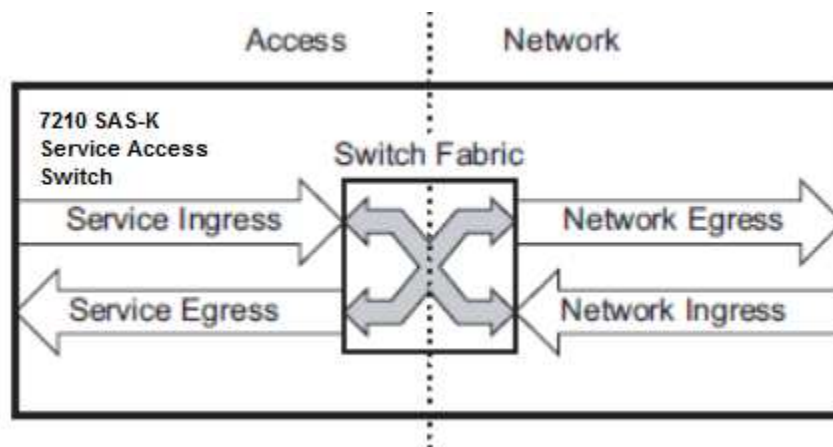


Figure 1: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C Service and Network Traffic Types and QoS model

The 7210 SAS uses QoS policies applied to a SAP for a service or to an access uplink port or to a network port to define the queuing, queue attributes, meter attributes, and QoS marking/interpretation.

The 7210 SAS supports the following major types of service and network QoS policies:

- Service ingress QoS policies
- Service Egress QoS policies
- Network QoS policies
- Network Queue QoS policies

The support of different policies varies across different platforms. More details are available in the following chapters and sections of this chapter.

Network QoS Policies

Network QoS Policies on network ports on 7210 SAS-K 2F4T6C

Platforms supported: 7210 SAS-K 2F4T6C

Network QoS policies define egress QoS marking and ingress QoS interpretation for traffic on received on network ports.

A network QoS policy defines both the ingress and egress handling of QoS on the network ports. The following functions are defined:

- Ingress
 - Option to use MPLS EXP value to map traffic to available forwarding classes and profile state.
 - Option to use Dot1p value to map traffic to available forwarding classes and profile state.
 - Option to use IP DSCP value to map traffic to available forwarding classes and profile state.
 - Option to use all of three above simultaneously along with DEI for forwarding class determination and assigning profile.
 - Defines forwarding class to queue mapping.
- Egress
 - Option to define the forwarding class and profile state to MPLS EXP value markings.
 - Option to define the forwarding class and profile state to Dot1p value markings.
 - Option to define the forwarding class and profile state to IP DSCP value marking
 - Remarking of QoS bits can be enabled or disabled

The required elements to be defined in a network QoS policy are:

- A unique network QoS policy ID.
- Egress forwarding class to priority bits (for example: 802.1p, etc.) used for marking, for each forwarding class.
- A default ingress forwarding class and an optional in-profile/out-of-profile state.
- At least one default unicast forwarding class meter or queue based on the platform. The parameters that can be configured for a meter and queue are discussed below.

Optional network QoS policy elements include:

- Additional queues.
- MPLS EXP value to forwarding class and profile state mappings for all values received
- Dot1p value to forwarding class and profile state mappings for all values received.
- Option to use DEI bit along with Dot1p classification for profile state mapping
- Option to use IP DSCP value to forwarding class and profile state mappings for all DSCP values received.

Ingress Classification Support for Network ports

On ingress of network ports user has an option to use both MPLS EXP and Dot1p bits to map received MPLS packets to forwarding class.

- If both MPLS EXP and Dot1p bits are configured, then the match order for MPLS packets is to match with MPLS EXP entries first and assign a FC if there is match. If no match, match with Dot1p entries, if configured and assign a FC if there is a match with Dot1p entries. If there is no match with both MPLS EXP and Dot1p entries, assign the default FC configured. DEI bit can be used to assign profile state to the MPLS packets on network ingress.
- On ingress of network ports, the user has an option to use both IP DSCP and Dot1p bits to map received IP packets (plain routed packets in the context of network IP interfaces configured on the network port) to forwarding class. If both IP DSCP and Dot1p bits are configured, then the match order for IP packets is to match with IP DSCP entries first and Dot1p entries next. The FC and profile value configured in the entry which matches first is assigned to the packet. If there is no match with either IP DSCP and Dot1p values, then the default FC is assigned to the packet. DEI bit can be used to assign profile state to the IP packets of network ingress.
- On 7210 SAS-K 2F4T6C, MPLS EXP classification entries that map MPLS EXP values to FC, Dot1p classification entries that map Dot1p bits to FC and IP DSCP classification entries that map IP DSCP values to FC is defined using MPLS EXP classification policies, Dot1p-classification policies and DSCP classification policies respectively.

Egress Marking Support for Network ports

On network port egress, option is provided to mark MPLS EXP and Dot1p values for MPLS packets. For IP packets sent out of network port, option is provided to mark IP DSCP and Dot1p values for IP packets. Along with Dot1p user has an option to mark DEI bit for both MPLS and IP DSCP packets.

Network policy ID 2 is reserved as the default network QoS policy applied to network ports. The default policy cannot be deleted or changed. The default network QoS policy is applied to all network ports which do not have another network QoS policy explicitly assigned.

The network QoS policy applied at network egress (that is, on an network port) determines how or whether the profile state is marked in packets transmitted into the service core network. If the profile state is marked in the service packets, out-of-profile packets are preferentially dropped over in-profile packets at congestion points in the network.

For network egress, traffic remarking in the network QoS policy can be enabled or disabled.

Table 4 lists the default mapping of forwarding class to Dot1p values for egress marking.

Table 4: Default Network QoS Policy Egress Marking on 7210 SAS-K 2F4T6C

FC-ID	FC Name	FC Label	DiffServ Name	Egress Dot1p Marking		MPLS EXP values		IP DSCP values	
				In-Profile	Out-Profile	In-Profile	Out-Profile	In-Profile	Out-Profile
7	Network Control	nc	NC2	111 - 7	NC	7	7	NC2	NC2
6	High-1	h1	NC1	110 - 6	H1	6	6	NC1	NC1
5	Expedited	ef	EF	101 - 5	EF	5	5	EF	EF
4	High-2	h2	AF4	100 - 4	H2	4	4	AF4	AF4
3	Low-1	l1	AF2	011 - 3	L1	3	3	AF2	AF2
2	Assured	af	AF1	011-3	AF1	2	2	AF1	AF1
1	Low-2	l2	CS1	001 - 1	L2	1	1	CS1	CS1
0	Best Effort	be	BE	000 - 0	BE	0	0	BE	BE

For network ingress, Table 5 lists the default mapping of Dot1p values to forwarding class and profile state for the default network QoS policy.

Table 5: Default Network QoS policy ID #2 - ingress classification map - Dot1p Mapping to forwarding class

Dot1p Value	FC	Profile
0	be	Out
1	l2	In
2	af	Out
3	af	In
4	h2	In
5	ef	In
6	h1	In
7	nc	In

For network ingress, [Table 6](#) the default mapping of mpls-lsp-exp-classification values to forwarding class and profile state for the default network QoS policy.

Table 6: Default Network QoS policy mpls-lsp-exp-classification to FC Mapping

MPLS EXP value	Ingress FC	Profile
0	be	Out
1	l2	In
2	af	Out
3	af	In
4	h2	In
5	ef	In
6	h1	In
7	nc	In

For network ingress, [Table 7](#) the default mapping of dscp-classification values to forwarding class and profile state for the default network QoS policy.

Table 7: Default Network QoS policy dscp-classification to FC Mapping

IP DSCP value	Ingress FC	Profile
be	be	Out
ef	ef	In
cs1	l2	In
nc1	h1	In
nc2	nc	In
af11	af	In
af12	af	Out
af41	h2	In

Network QoS policies on Access-uplink ports

Platforms supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Network QoS policies define egress QoS marking and ingress QoS interpretation for traffic on received on access-uplink ports.

A network QoS policy defines both the ingress and egress handling of QoS on the access uplink ports. The following functions are defined:

- Ingress
 - Option to use Dot1p value mapping to forwarding classes and profile.
 - Option to use IP DSCP value to map traffic to different forwarding classes and profile.
 - Defines forwarding class to ingress queue mapping
- Egress
 - Option to define the forwarding class and profile to Dot1p value markings.
 - Option to define the forwarding class and profile to IP DSCP value marking
 - Remarking of QoS bits can be enabled or disabled.

The required elements to be defined in a network QoS policy are:

- A unique network QoS policy ID.
- Egress - forwarding class and optional profile state to priority bits (for example: 802.1p, etc.) used for marking, for each forwarding class.
- A default ingress forwarding class and an optional in-profile/out-of-profile state.
- At least one default unicast forwarding class queue. The parameters that can be configured for a ingress queue are discussed below.

Optional network QoS policy elements include:

- Additional queues.
- Dot1p value to forwarding class and profile state mappings for all values received.
- Option to use DEI bit along with Dot1p classification for profile state mapping
- Option to use IP DSCP value to forwarding class and profile state mappings for all DSCP values received.

Ingress Classification Support for access-uplink ports

On ingress of access-uplink ports user has an option to use both IP DSCP and Dot1p bits to map received IP packets to forwarding class. If both IP DSCP and Dot1p bits are configured, then the match order for IP packets is to match with IP DSCP entries first and Dot1p entries next. The FC and profile value configured in the entry which matches first is assigned to the packet. If there is no match with either IP DSCP and Dot1p values, then the default FC is assigned to the packet. DEI bit can be used to assign profile state to the packets of network ingress.

On 7210 SAS-K 2F4T6C, Dot1p classification entries that map Dot1p bits to FC and IP DSCP classification entries that map IP DSCP values to FC is defined by using Dot1p-classification policies and DSCP classification policies respectively.

Network policy ID 1 is reserved as the default network QoS policy applied to access-uplink ports. The default policy cannot be deleted or changed. The default network QoS policy is applied to all access uplink ports which do not have another network QoS policy explicitly assigned.

The network QoS policy applied at network egress (that is, on an access uplink port) determines how or whether the profile state is marked in packets transmitted into the service core network. If the profile state is marked in the service packets, out-of-profile packets are preferentially dropped over in-profile packets at congestion points in the network.

For network egress, traffic remarking in the network QoS policy can be enabled or disabled. See [Table 8](#) for more information.

Table 8: Default Network QoS policy Dot1p to FC Mapping for network egress

Dot1p Value	Ingress FC	Profile
0	be	Out
1	l2	In
2	af	Out
2	l1	In
4	h2	In
5	ef	In
6	h1	In
7	nc	In

For network ingress, [Table 9](#) the default mapping of Dot1p values to forwarding class and profile state for the default network QoS policy.

Table 9: Default Network QoS policy Dot1p to FC Mapping for network ingress

Dot1p Value	Ingress FC	Profile
0	be	Out
1	l2	In
2	af	Out
3	af	In
4	h2	In
5	ef	In
6	h1	In
7	nc	In

Network Queue policies

Network Queue policies on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Network queue policies are applied on egress of access-uplink ports on SAS-K2F2T1C.

On 7210 SAS-K2F4T6C, Network queue policies are applied on egress of access-uplink ports and network ports.

The Network queue policies can be defined with up to a maximum of 8 egress queues. The user has an option to define the policies with less than eight egress queues.

The queue characteristics configured on a per-forwarding class basis are:

- Committed Buffer Size (CBS) in Kilobytes
- Maximum Buffer Size (MBS) in Kilobytes
- Peak Information Rate (PIR) as a percentage of egress port bandwidth
- Committed Information Rate (CIR) as a percentage of egress port bandwidth
- Queue priority and Queue weight

Network queue policies are identified with a unique policy name which conforms to the standard 7210 SAS alphanumeric naming conventions. The system default network queue policy is named default and cannot be edited or deleted. [Table 9](#) describes the default network queue policy definition in access-uplink mode.

Table 10: Default Network Queue Policy Definition on 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C.

Forwarding Class	Queue	Definition
Network-Control (nc)	Queue 8	PIR = 100% CIR = 10% MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F2T1C CBS = 24 Kilobytes -7210 SAS-K 2F4T6C priority = 1 weight = 1
High-1 (h1)	Queue 7	PIR = 100% CIR = 10% MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F2T1C CBS = 24 Kilobytes -7210 SAS-K 2F4T6C priority = 1 weight = 1
Expedited (ef)	Queue 6	PIR = 100% CIR = 100% MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F2T1C CBS = 24 Kilobytes -7210 SAS-K 2F4T6C priority = 1 weight = 1

Table 10: Default Network Queue Policy Definition on 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C.

Forwarding Class	Queue	Definition (Continued)
High-2 (h2)	Queue 5	PIR = 100% CIR = 100% MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F2T1C CBS = 24 Kilobytes -7210 SAS-K 2F4T6C priority = 1 weight = 1
Low-1 (l1)	Queue 4	PIR = 100% CIR = 25% MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F2T1C CBS = 24 Kilobytes -7210 SAS-K 2F4T6C priority = 1 weight = 1
Assured (af)	Queue 3	PIR = 100% CIR = 25% MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F2T1C CBS = 24 Kilobytes -7210 SAS-K 2F4T6C priority = 1 weight = 1

Table 10: Default Network Queue Policy Definition on 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C.

Forwarding Class	Queue	Definition (Continued)
Low-2 (12)	Queue 2	PIR = 100% CIR = 25% MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F2T1C CBS = 24 Kilobytes -7210 SAS-K 2F4T6C priority = 1 weight = 1
Best-Effort (be)	Queue 1	PIR = 100% CIR = 0% CBS = 7% MBS = 200 Kilobytes CBS = 50 Kilobytes -7210 SAS-K 2F2T1C CBS = 24 Kilobytes -7210 SAS-K 2F4T6C priority = 1 weight = 1

Queues and Queue Parameters

This section describes the queue parameters provisioned for queues used in service ingress policy, service egress policy, access egress policy, network qos policy and network queue policy.



Note: Not all 7210 platforms support queues for all the above policies. In addition, the queue parameters support available varies across different platforms. See platform specific QoS overview sections above and the chapter following to know the support available on different platforms.

Queues are available on different platforms as follows:

On 7210 SAS-K 2F2T1C queues are available with the following:

- SAP ingress policies associated with access SAP ingress.

QoS Overview

- SAP egress policies associated with access SAP egress.
- Network Queue policies associated with access-uplink port egress.
- Network QoS policies associated with access-uplink port ingress.

On 7210 SAS-K 2F4T6C queues are available with the following:

- SAP ingress policies associated with access SAP ingress.
- SAP egress policies associated with access SAP egress.
- Network Queue policies associated with access-uplink port egress.
- Network QoS policies associated with access-uplink port ingress.
- Network Queue policies associated with network port egress.
- Network QoS policies associated with network port ingress.

The queue parameters are:

- [Queue- Queue ID](#)
- [Queue- Committed Information Rate](#)
- [Queue- Peak Information Rate](#)
- [Queue- Adaptation Rule for Queues](#)
- [Queue- Committed Burst Size](#)
- [Queue – Ingress Profile Assignment](#)
- [Queue – Weight and Priority](#)
- [Queue Counters](#)

Queue- Queue ID

The queue ID is used to uniquely identify the queue. The queue ID is only unique within the context of the QoS policy within which the queue is defined. It allows user to define multiple queues with different characteristics and identify them while associating it with different forwarding classes.

On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, user has an option to allocate upto 8 queues and assign them queue ID along with option to configure some of the queue parameters which determine the queue characteristics.

Queue- Committed Information Rate

The committed information rate (CIR) for a queue performs two distinct functions:

- **Minimum bandwidth guarantees** — Queue's CIR setting provides the bandwidth which will be given to this queue as compared to other queues on the port competing for a share of the available link bandwidth. The queue CIR does not necessarily guarantee bandwidth in all scenarios and also depends on factors such as CIR over subscription and link port bandwidth capacity. For each packet in a queue, the CIR is checked with the current transmission rate of the queue. If the current rate is at or below the CIR threshold, the queue is considered in-profile. If the current rate is above the threshold, the queue is considered out-of-profile. The in-profile and out-profile state of queue is linked to scheduler prioritizing behavior as discussed below.
- **Scheduler queue priority metric** — The scheduler serving a group of queues prioritizes individual queues based on their current CIR and PIR states. Queues operating below their CIR are always served before those queues operating at or above their CIR. Also see information about schedulers to know the scheduler behavior on different 7210 platforms.

NOTE: On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the in-profile and out-profile state of the ingress queue determines the packets final profile state based on the queue CIR, PIR values. The in-profile and out-profile state of the ingress queue also interacts with the scheduler mechanism and provides the minimum and maximum bandwidth guarantees. This is true only off ingress queues and not of egress queues. In other words, the in-profile and out-profile state of the egress queue does not change the packets final profile state based on the queue CIR, PIR values. The in-profile and out-profile state of the egress queue only interacts with the scheduler mechanism and provides the minimum and maximum bandwidth guarantees. Please see the section on “Queue – Ingress Profile Assignment on 7210 SAS-K” to know more.

When defining the CIR for a queue, the value specified is the administrative CIR for the queue. User has some control over how the administrative CIR is converted to an operational CIR should the hardware not support the exact CIR and PIR combination specified. The interpretation of the administrative CIR is discussed below in [Queue- Adaptation Rule for Queues](#). Although the 7210 SAS is flexible in how the CIR can be configured, there are conventional ranges for the CIR based on the forwarding class of a queue. An egress queue associated with the high-priority class normally has the CIR threshold equal to the PIR rate although the 7210 SAS allows the CIR to be provisioned to any rate below the PIR should this behavior be required.

The CIR of the queue is configurable under the different qos policies that provide the option to configure the queue parameters – example under service ingress and service egress queue policies, access port policies, network queue policies, etc.

Queue- Peak Information Rate

The peak information rate (PIR) defines the maximum rate at which packets are allowed to exit the queue. It does not specify the maximum rate at which packets may enter the queue; this is governed by the queue's ability to absorb bursts. The actual transmission rate of an egress queue depends on more than just its PIR. Each queue is competing for transmission bandwidth with other queues. Each queue's PIR, CIR and the relative priority and/or weight of the scheduler serving the queue, all combine to affect a queue's ability to transmit packets.

When defining the PIR for a queue, the value specified is the administrative PIR for the queue. The user has some control over how the administrative PIR is converted to an operational PIR should the hardware not support the exact CIR and PIR values specified. The interpretation of the administrative PIR is discussed below in [Queue- Adaptation Rule for Queues](#)

The PIR of the queue is configurable under the different qos policies that provide the option to configure the queue parameters – example under service ingress and service egress queue policies, access port policies, network queue policies, etc.

Queue- Adaptation Rule for Queues

The adaptation rule provides the QoS provisioning system with the ability to adapt specific CIR and PIR defined administrative rates to the underlying capabilities of the hardware the queue will be created on to derive the operational rates. The administrative CIR and PIR rates are translated to actual operational rates enforced by the hardware queue. The rule provides a constraint used when the exact rate is not available.

For the CIR and PIR parameters individually, the system will attempt to find the best operational rate depending on the defined constraint. The supported constraints are:

- **Minimum** — Find the hardware supported rate that is equal to or higher than the specified rate.
- **Maximum** — Find the hardware supported rate that is equal to or lesser than the specified rate.
- **Closest** — Find the hardware supported rate that is closest to the specified rate.

Depending on the platform on which the queue is provisioned, the actual operational CIR and PIR settings used by the queue will be dependent on the method the hardware uses to implement and represent the mechanisms that enforce the CIR and PIR rates. The adaptation rule controls the method the system uses to choose the rate step based on the administrative rates defined by the rate command.

NOTE: For 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the 7210 SAS software considers the adaptation rules and the hardware values available to determine the admin PIR/CIR rates.

To illustrate (the example that follows is only for illustration of the use of adaptation rule and the values provided below does not list the actual values supported in hardware), how the adaptation rule constraints minimum, maximum and closest are evaluated in determining the operational CIR or PIR for the 7210 SAS, assume there is a queue where the administrative CIR and PIR values are 90Kbps and 150 Kbps, respectively. If the adaptation rule is minimum, the operational CIR and PIR values will be 128Kbps and 192Kbps respectively, as it is the native hardware rate greater than or equal to the administrative CIR and PIR values. If the adaptation rule is maximum, the operational CIR and PIR values will be 64Kbps and 128Kbps. If the adaptation rule is closest, the operational CIR and PIR values will be 64Kbps and 128Kbps, respectively, as those are the closest matches for the administrative values that are even multiples of the 64 Kbps rate step.

Queue- Committed Burst Size

The committed burst size (CBS) parameters specify the amount of buffers that can be drawn from the reserved buffer portion of the queue's buffer pool. Once the reserved buffers for a given queue have been used, the queue contends with other queues for additional buffer resources up to the maximum burst size.

The CBS of the queue is configurable under the different QoS policies, if supported by the platform, that provide the option to configure the queue parameters – example under service ingress and service egress queue policies, [access port policies](#), network queue policies, etc. The CBS for a queue is specified in Kbytes.

On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the CBS for the queues is user configurable. By default, software assigns a default value. It can be modified by the user as per their needs. The default value are specified in the command description.

Queue- Maximum Burst Size

The maximum burst size (MBS) parameter specifies the maximum queue depth to which a queue can grow. This parameter ensures that a customer that is massively or continuously oversubscribing the PIR of a queue will not consume all the available buffer resources. For high-priority forwarding class service queues, the MBS can be relatively smaller than the other forwarding class queues because the high-priority service packets are scheduled with priority over other service forwarding classes.

The MBS of the queue is configurable under the different QoS policies, if supported by the platform, that provide the option to configure the queue parameters – example under service ingress and service egress queue policies, [access port policies](#), network queue policies, etc. The MBS for a queue is specified in Kbytes.

On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the MBS for the queues is user configurable. By default, software assigns a default value. It can be modified by the user as per their needs. The default values are specified in the command description.

Queue – Ingress Profile Assignment

On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C queues can operate in two modes – profile mode and non-profile mode. On 7210 SAS-K 2F2T1C, SAP Ingress queues and Access uplink port ingress queues operate in either profile mode or non-profile mode.

On 7210 SAS-K 2F4T6C, SAP Ingress queues, Network port ingress queues and Access uplink port ingress queues operate in either profile mode or non-profile mode

In ‘profile mode’, the profile defined in the policy is used to determine the WRED slope to use for ingress queuing, with ‘profile in’ packets using high-slope and ‘profile out’ packets using low-slope. The ingress queue shaper does not change the profile value assigned to a packet that has a user assigned profile value. In other words, if an user assigns a profile value of green and the packet exceeds the CIR rate of the shaper, it is not changed to yellow. Similarly, packets assigned yellow color is not changed by the shaper. The color assigned by the user is also subsequently used at the egress queuing point to determine the slope to use.

In ‘non-profile’ mode, the profile is not specified by the user (and hence the node maps it to ‘undefined’ value. The low WRED slope is used at the ingress queuing point, as all packets received are considered to be the same as ‘profile out’. The packet is then assigned the profile by the ingress queue shaper. It is assigned ‘in’ profile value if it’s within the CIR and assigned ‘out’ profile value if it exceeds the CIR. It is dropped if it exceeds the PIR rate of ingress queue shaper (except for packets which are assigned a profile value of “undefined” on ingress and where the shaper assigns the profile based on CIR/PIR rate). The profile assigned by the ingress queue shaper is also subsequently used at the egress queuing point to determine the slope to use.

The user is provided with different options to assign the profile to the packet (for example: DEI based). It is always assigned on ingress of the packet into the node. Once the profile is assigned at the ingress, it is used at subsequent queuing points in the system. In other words, subsequent modules and queuing points in the system do not change the profile assigned to the packet on ingress. The profile assigned at ingress is also used to subsequently assign different marking/remarking values to in-profile and out-of-profile packets, if the user so desires.

Queue – Weight and Priority

On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, user is provided an option to assign the priority and weight to the queue. The priority determines the service order of the queue when the scheduler schedules multiple queues configured on the same port. The queue weight determines the proportion of the available bandwidth that the scheduler allocates to a queue.

Queue Counters

The router maintains counters for queues within the system for granular billing and accounting.

Each queue maintains the following counters:

- Counters for packets and octets accepted into the queue
- Counters for packets and octets rejected at the queue
- Counters for packets and octets transmitted in-profile
- Counters for packets and octets transmitted out-of-profile

The counters available vary among the 7210 SAS platform. Not all the counters are supported on all the platforms. Additionally there are restrictions on the number of counters that can be used simultaneously with a single queue. Some platforms can only count octets or packets and other can count both packets and octets. Counter (and corresponding accounting record) support available on each of the platform is listed in the 7210 SAS System Management user guide under the “Accounting Records/Logs” section.

Queue – support on various platforms

The following support is available on different platforms:

- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C - Service queue is provisioned on access sap ingress and access sap egress service queues within service ingress QoS policies and service egress QoS policies, respectively.
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C - access-uplink port egress queues are defined within network queue policies.
- On 7210 SAS-K 2F4T6C - network port ingress queues are defined within network QoS policies.
- On 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C - access-uplink port ingress queues are defined within network qos policies.

- On 7210 SAS-K2F4T6C - network port egress queues are defined within network queue policies.

Service Ingress QoS Policies

Service ingress QoS policies define ingress service forwarding class queues or meters and map traffic flows to forwarding class on access SAP ingress.



Note: Not all 7210 platforms support queues and meters on service ingress. The support varies across different platforms. Please read the subsequent chapters/sections for more information.

Service Ingress QoS policies

Service ingress QoS policies define ingress service forwarding class queues and map traffic flows to forwarding class.

On 7210 SAS-K platforms, when a service ingress QoS policy is created, it always has one queue defined that cannot be deleted. The queue is used to queue all the traffic, both the unicast traffic and the multipoint traffic. These queues exist within the definition of the policy. The queues only get instantiated in hardware when the policy is applied to a SAP. Multipoint queues are instantiated only if the SAP ingress policy defines a multipoint queue. By default, software does not allocate any multipoint queues.

In the simplest service ingress QoS policy, all traffic is treated as a single flow and mapped to a single queue, including all flooded traffic.

The required elements to define a service ingress QoS policy are:

A unique service ingress QoS policy ID.

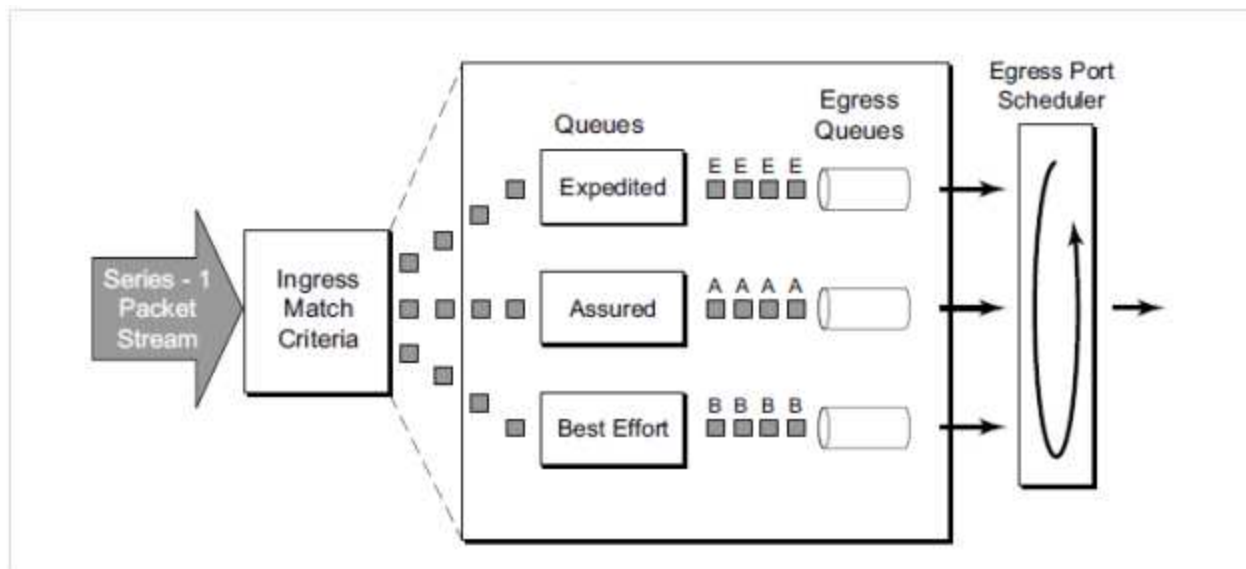
- A QoS policy scope of template or exclusive.
- At least one default forwarding class queue. The parameters that can be configured for a queue are discussed in [Queues and Queue Parameters](#).

Optional service ingress QoS policy elements for 7210 SAS-K platforms include:

- Additional unicast queues or multicast queues up to a total of 8.
- QoS policy match criteria to map packets to a forwarding class.

Each queue can have unique queue parameters to allow individual shaping of the flow mapped to the forwarding class. The [Figure 2](#) below depicts service traffic being classified into three different forwarding classes.

Figure 2: Traffic Queuing Model for Forwarding Classes



Default Service Ingress Policy

Service ingress QoS policy ID 1 is reserved for the default service ingress policy. The default policy cannot be deleted or changed. The default service ingress policy is implicitly applied to all SAPs which do not explicitly have another service ingress policy assigned. In the default policy, all traffic is mapped to the default forwarding class which uses a queue by default. The characteristics of the default policy are listed in [Table 11](#).

Table 11: Default Service Ingress Policy ID 1 Definition for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Characteristic	Item	Definition
Queue	Queue 1	1 (one) queue all unicast traffic and multicast traffic: <ul style="list-style-type: none"> • Forward Class: best-effort (be) • CIR = 0 • PIR = max • MBS = 60 Kilobytes • CBS = 10 Kilobytes • Priority= 1 • Weight= 1
Flows	Default Forwarding Class (be)	1 (one) flow defined for all traffic: <ul style="list-style-type: none"> • All traffic mapped to best-effort (be)

Service Ingress Classification

Mapping flows to forwarding classes is controlled by comparing each packet to the match criteria in the Service Ingress QoS policy applied to an access SAP. The ingress packet classification to forwarding class is subject to a classification policy provisioned.

Service Ingress Classification

On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices, on access SAP ingress user has an option to use either Dot1p classification or IPv4 DSCP classification or IPv4 packet header fields or IPv6 packet header fields or MAC packet header fields. The Dot1p or DSCP classification rules to be used are defined in the Dot1p and DSCP classification policy and associated with the SAP ingress policy. The DSCP and Dot1p classification policies can be configured in the same QoS policy. The IPv4 or IPv6 or MAC criteria can be configured in the SAP ingress policy.

When packets are received on an access SAP, the following steps are used to determine the FC to assign to the packet.

- Step 1.** Match IP criteria entries with the IP packet header fields in the packet. Assign the FC corresponding to the first entry which matches with IP packet header field values in the packet. If it is not an IP packet or if there is no match, go to next step.

- Step 2.** Match MAC criteria entries with the MAC packet header fields in the packet. Assign the FC corresponding to the first entry which matches with MAC packet header field values in the packet. If there is no match, go to next step.
- Step 3.** Match the IP DSCP value in the packet with the value configured in each of the IP DSCP entry defined in the DSCP classification policy. Assign the FC corresponding to the first entry which matches with IP DSCP value in the packet. If it is not an IP packet or if there is no match, go to next step.
- Step 4.** Match the Dot1p value in the packet (if available) with the value configured in each of the Dot1p entry defined in the Dot1p classification policy. Assign the FC corresponding to the first entry which matches with Dot1p value in the packet. If there is no match, go to next step.
- Step 5.** Assign the default FC.

Service Ingress Classification– IP and Mac Packet fields

The IP and MAC match criteria can be very basic or quite detailed. IP and MAC match criteria are constructed from policy entries. An entry is identified by a unique, numerical entry ID. A single entry cannot contain more than one match value for each match criteria. Each match entry has an action which specifies: the forwarding class of packets that match the entry. The entries are evaluated in numerical order based on the entry ID from the lowest to highest ID value. The first entry that matches all match criteria has its action performed.

[Table 12](#) lists the packets fields used for match-criteria used for access SAP ingress classification on the different 7210 platforms.

Table 12: Service Ingress QoS Policy IP Match Criteria for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

IP Criteria	Services applicable on 7210 SAS-K 2F2T1C	Services applicable on 7210 SAS-K 2F4T6C
<ul style="list-style-type: none"> IP DSCP 	Access SAPs in Epipe, VPLS, IES, RVPLS services	Access SAPs in Epipe, VPLS, IES, VPRN, RVPLS services

Table 12: Service Ingress QoS Policy IP Match Criteria for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

IP Criteria	Services applicable on 7210 SAS-K 2F2T1C	Services applicable on 7210 SAS-K 2F4T6C
<ul style="list-style-type: none"> • IP source address and mask, IP destination address and mask, IP protocol, TCP/UDP source port and fragment field • TCP/UDP destination port 	Access SAPs in Epipe, VPLS, IES, RVPLS services	Access SAPs in Epipe, VPLS, IES, VPRN, RVPLS services

Table 13: Service Ingress QoS Policy IPv6 Criteria for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

IPv6 Criteria	Services applicable on 7210 SAS-K 2F2T1C	Services applicable on 7210 SAS-K 2F4T6C
<ul style="list-style-type: none"> • DSCP value, Destination IPv6 address and mask match, Destination port TCP/UDP port match, Source IPv6 address and mask match, Source port TCP/UDP port match 	Access SAPs in Epipe, VPLS, RVPLS services	Access SAPs in Epipe, VPLS, RVPLS services.

Table 14: Service Ingress QoS Policy MAC Match Criteria for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

MAC Match Criteria	Services applicable on 7210 SAS-K 2F2T1C	Services applicable on 7210 SAS-K 2F4T6C
<ul style="list-style-type: none"> IEEE 802.1p/Dot1p value/mask (for both inner and outer tag separately), Source MAC address/mask, Destination MAC address/mask, EtherType Value/Mask, outer VLAN, and inner VLAN tag value 	Access SAPs in Epipe, VPLS,RVPLS services	Access SAPs in Epipe, VPLS,RVPLS services.

Note that 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C does not support configuring of the frame-type match criteria and the default frame type configured is Ethernet - II. [Table 15](#)

Table 15: MAC Match Ethernet Frame Types

Note: The default frame type configured is Ethernet - II	
Frame Format	Description
802.3	IEEE 802.3 Ethernet frame. Only the source MAC, destination MAC and IEEE 802.1p value are compared for match criteria.
802dot2-llc	IEEE 802.3 Ethernet frame with an 802.2 LLC header. Only the source MAC and destination MAC address are compared for match criteria.
802dot2-snap	IEEE 802.2 Ethernet frame with 802.2 SNAP header. Only the source MAC and destination MAC address are compared for match criteria.
Ethernet-II	Ethernet type II frame where the 802.3 length field is used as an Ethernet type (Etype) value. Etype values are two byte values greater than 0x5FF (1535 decimal).

[Table 16](#) lists the criteria that can be matched for the various MAC frame types.

Table 16: MAC Match Criteria Frame Type Dependencies

Frame Format	Source MAC	Dest MAC	IEEE 802.1p Value	Etype Value
802.3	Yes	Yes	Yes	No
802dot2-llc	Yes	Yes	Yes	No
802dot2-snap	Yes	Yes	Yes	No
ethernet-II	Yes	Yes	Yes	Yes

Service Egress QoS Policies

Service egress queues are implemented at the transition from the service core network to the service access network on access SAPs. The advantages of per-service queuing before transmission into the access network are:

- Per-service egress sub-rate capabilities.
- More granular, fairer scheduling per-service into the access network.
- Per-service statistics for forwarded and discarded service packets.

The subrate capabilities and per-service scheduling control are required to make multiple services per physical port possible. With egress shaping, it is possible to support more than one service per port. It prevents traffic from single service from bursting to the available port bandwidth and starving other services.

For accounting purposes, per-service statistics can be logged. When statistics from service ingress queues are compared with service egress queues, the ability to conform to per-service QoS requirements within the service core can be measured.

Service egress QoS policies define egress queues and map forwarding class flows to queues. In the simplest service egress QoS policy, all forwarding classes are treated like a single flow and mapped to a single queue. To define a basic egress QoS policy, the following are required:

- A unique service egress QoS policy ID.
- A QoS policy scope of template or exclusive.
- At least one defined default queue.

Optional service egress QoS policy elements include:

- Additional queues up to a total of 8 separate queues. A forwarding class queue is shared by unicast and multipoint (BUM) traffic type mapped to that forwarding class.

- IEEE 802.1p priority value remarking based on forwarding class.
- Option to use IP DSCP values for marking based on FC.

Each queue in a policy is associated with one of the forwarding classes. Each queue can have its individual queue parameters allowing individual rate shaping of the forwarding class(es) mapped to the queue. More complex service queuing models are supported in the router where each forwarding class is associated with a dedicated queue. The forwarding class determination per service egress packet is determined at ingress. If the packet ingressed the service on the same router, the service ingress classification rules determine the forwarding class of the packet. If the packet is received, the forwarding class is marked in the tunnel transport encapsulation (for example: QinQ encapsulated packet).

Default Service Egress Policy

Service egress QoS policy ID 1 is reserved as the default service egress policy. The default policy cannot be deleted or changed. The default access egress policy is applied to all SAPs service egress policy explicitly assigned. The characteristics of the default policy are listed in the following [Table 17](#).

Table 17: Default Service Egress Policy ID 1 Definition

Characteristic	Item	Definition
Queues	Queue 1	1 (one) queue defined for all traffic classes: <ul style="list-style-type: none"> • CIR = 0 • PIR = max (line rate) • MBS = 60 Kilobytes • CBS = 10 Kilobytes • Priority= 1 • Weight= 1
Flows	Default Action	1 (one) flow defined for all traffic classes: <ul style="list-style-type: none"> • All traffic mapped to queue 1 with no marking of IEEE 802.1p values

Buffer Pools

Buffer pools are used to manage the packet buffer memory resources used to store packets and absorb bursts received on a queue.

Buffer pools on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

The total amount of available buffers (~64MB on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C) is divided among the 5 buffer pools listed below. In addition some of the buffers are reserved for system internal use (such as, Multicast queues).

- CBS buffer pool
- Ingress non-ring MBS pool
- Egress non-ring MBS pool
- Ingress ring MBS pool
- Egress ring MBS pool

CBS buffer pool is used to allocate buffers towards committed burst size (CBS) configured for ingress and egress queues on the node and some internal system queues.

The CBS pool allocation on different SAS-K platforms is as given below:

- On 7210 SAS-K 2F2T1C, the CBS pool is used to allocate buffers towards CBS configured for ingress and egress queues on access SAP, and ingress and egress queues on access-uplink port.
- On 7210 SAS-K 2F4T6C, the CBS pool is used to allocate buffers towards CBS configured for ingress and egress queues on access SAP, ingress and egress queues on access-uplink ports and ingress and egress queues on network ports.

MBS pool is divided into four pools as shown above. The Ingress and Egress non-ring MBS buffer pool and the Ingress and Egress ring MBS buffer pool. These MBS buffer pools can be over-subscribed.

The Ingress and Egress ‘non-ring’ MBS buffer pool is used to allocate buffers towards the Maximum Burst Size (MBS) configured for ingress queues and egress queues respectively on ‘non-ring’ ports. On both 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the access ports are considered to be ‘non-ring’ ports by default (user does not have an option to change it). The Ingress non-ring MBS pool is used to allocate buffers towards all ingress queues configured on access SAPs. Similarly the egress non-ring MBS pool is used to allocate buffers towards all egress queues configured on access SAPs.

The Ingress and Egress ‘ring’ MBS buffer pool is used to allocate buffers towards Maximum Burst Size (MBS) configured for ingress queues and egress queues respectively on ‘ring’ ports.

Assignment of ring ports and ring MBS buffer pool on different SAS-K platforms is as given below:

- On 7210 SAS-K 2F2T1C, all access-uplink ports are considered to be `ring' ports and the Ingress and Egress ring MBS buffer pool is used to allocate buffers towards access-uplink port ingress and egress queues respectively.
- On 7210 SAS-K 2F4T6C, all network ports and access-uplink ports are considered to be `ring' ports by default (user does not have an option to change it) and the Ingress and Egress ring MBS buffer pool is used to allocate buffers towards network port and access-uplink port ingress queues and network port and access-uplink port egress queues respectively.

The amount of memory allocated towards these pools is software defined and not user configurable. The show pools <port-id> system can be used to display total amount of buffers per pool and the amount of buffers in use per pool.

```
A:dut-i# show pools 1/1/1 system

=====
Pool Information
=====
Port          : 1/1/1
Application   : System

MMU Total     : 65536 KB

MMU CBS       : 14336 KB           MMU CBS In Use      : 2240 KB

Ingress Ring  : 11776 KB           Ingress Ring In Use: 0 KB
Ingress NonRing : 11776 KB       Ingress NonRing In*: 0 KB

Egress Ring   : 11776 KB           Egress Ring In Use : 0 KB
Egress NonRing : 11776 KB       Egress NonRing In *: 0 KB

----- snipped -----
```

Ring and Non-Ring Buffer pool

When 7210 SAS-K 2F2T1C is deployed in a ring environment, the access-uplink ports are typically used to connect the node to ring, similarly on 7210 SAS-K 2F4T6C, users will typically use the network ports to join the node into a ring. Therefore, these ports are designated as the ring ports. These ring ports carry traffic from the head-end towards the node (that is, 7210 SAS-K), dropping traffic off to user/customer locations. Simultaneously, these ring ports carry traffic from the user/customer to the head-end. In other words, traffic received from the user/customer is added to the ring and sent out towards the service edge, where services are terminated. The traffic in both these directions is typically admitted into the ring after being shaped and groomed. That is in the upstream direction (that is, in the direction of customer to service edge) the SLA is enforced at service ingress points (that is, typically access SAPs) and the traffic is shaped and groomed, similarly in the downstream direction

(that is, in the direction of service edge to customer) it is done by the service edge device or the access aggregation device. In other words, the downstream traffic should typically be allowed to pass through the intermediate nodes of the ring, without contention with and prioritized over the traffic that is received from the customer and being added into the ring by the nodes on the ring.

On 7210 SAS-K 2F2T1C, the access-uplink ports are designated as ‘ring’ ports and access ports are designated as ‘non-ring’ ports. Traffic going from any access-uplink to another access-uplink port is identified as ‘ring’ traffic. Traffic going from an access port to any access-uplink port, or traffic going from any access-uplink port to an access port (in egress), or traffic going from an access port to another access port is identified as ‘non-ring’ traffic.

On 7210 SAS-K 2F4T6C, the network ports and access-uplink ports are designated as ‘ring’ ports and access ports are designated as ‘non-ring’ ports. Traffic going from any network port or access-uplink to another network port or access-uplink port is identified as ‘ring’ traffic. Traffic going from an access port to any network port or access-uplink port, or traffic going from any network port or access-uplink port to an access port (in egress), or traffic going from an access port to another access port is identified as ‘non-ring’ traffic.

To ensure that the traffic received on ring ports is prioritized over traffic received on non-ring access ports, a separate ‘Ring’ MBS buffer pool (one each for ingress and egress) is provided for traffic received and sent out of ring ports. In addition, on network port egress and access-uplink egress (where shaped customer (access) traffic and ‘ring’ traffic share the ‘ring’ pool) two additional ring slopes (for a total of 4 configurable WRED slopes) are provided to prioritize the ‘ring’ traffic. Each egress queue on the network port and access uplink port supports 4 slopes per queue – ring high-slope, ring low-slope, non-ring high-slope and non-ring low-slope (in the CLI command the ring slopes are configured using the high-slope-ring and low-slope-ring and the non-ring slopes are configured using the high-slope and low-slope). Ring high-slope and Ring low-slope is used for in-profile and out-of-profile QoS profile ‘ring’ traffic. Non-ring high-slope and low-slope is used for in-profile and out-of-profile ‘non-ring’ traffic. Slope parameters (start-avg, max-avg, max-Prob) of 4 slopes can be configured such that the ring traffic is prioritized over the ‘non-ring’ traffic (that is, traffic being added onto the ring) in congestion scenarios.

A separate ‘non-Ring’ MBS buffer pool for traffic received and sent out of access ports along with 2 configurable WRED slopes is supported. Each queue on the access ports supports 2 slopes per queue –non-ring high-slope and non-ring low-slope. Non-ring high-slope and low-slope is used for in-profile and out-of-profile ‘non-ring’ traffic. The Non-Ring buffer pool (one each for ingress and egress) is used to allocate buffers for non-ring traffic.

The usage of buffer pools for different traffic flows is as given below:

On 7210 SAS-K 2F2T1C the ring and non-ring buffer pools are used by the following traffic flows:

- Traffic received on access-uplink SAP and sent out of access-uplink SAP, uses the ring MBS buffer pool for MBS buffers on access-uplink port ingress and access-uplink port egress. In this case, ring high-slope is used for in-profile traffic and ring low-slope is used for out-of-profile traffic for both access-uplink ingress and access-uplink egress.
- Traffic received on access SAP and sent out of access-uplink SAP, uses the non-ring MBS buffer pool for MBS buffers on access SAP ingress and uses the ring MBS buffer pool for MBS buffers on access-uplink SAP egress. In this case, non-ring high slope and non-ring low slope is used on both access SAP ingress and access-uplink egress.
- Traffic received on access-uplink SAP and sent out of another access SAP uses ring MBS buffer pool for MBS buffers on access-uplink SAP ingress and the non-ring MBS buffer pool for MBS buffers on access SAP egress. In this case, ring high-slope and ring low-slope is used on access-uplink ingress and non-ring high-slope and non-ring low-slope is used on access egress.
- Traffic received on access SAP and sent out of another access SAP uses the non-ring MBS pool for MBS buffers for both access SAP ingress and access SAP egress. In this case, non-ring high-slope and non-ring low-slope is used on both access ingress and access egress.

On 7210 SAS-K 2F4T6C the ring and non-ring buffer pools are used by the following traffic flows:

- Traffic received on network port and sent out of network port, uses the ring MBS buffer pool for MBS buffers on network port ingress and network port egress. In this case, ring high-slope is used for in-profile traffic and ring low-slope is used for out-of-profile traffic for both network port ingress and port egress.
- Traffic received on access SAP and sent out of port, uses the non-ring MBS buffer pool for MBS buffers on access SAP ingress and uses the ring MBS buffer pool for MBS buffers on network port egress. In this case, non-ring high slope and non-ring low slope is used on both access SAP ingress and network port egress.
- Traffic received on network port and sent out of another access SAP uses ring MBS buffer pool for MBS buffers on network port ingress and the non-ring MBS buffer pool for MBS buffers on access SAP egress. In this case, ring high-slope and ring low-slope is used on network port ingress and non-ring high-slope and non-ring low-slope is used on access egress.
- Traffic received on access-uplink SAP and sent out of access-uplink SAP, uses the ring MBS buffer pool for MBS buffers on access-uplink port ingress and access-uplink port egress. In this case, ring high-slope is used for in-profile traffic and ring low-slope is used for out-of-profile traffic for both access-uplink ingress and access-uplink egress.

- Traffic received on access SAP and sent out of access-uplink SAP, uses the non-ring MBS buffer pool for MBS buffers on access SAP ingress and uses the ring MBS buffer pool for MBS buffers on access-uplink SAP egress. In this case, non-ring high slope and non-ring low slope is used on both access ingress and access-uplink egress.
- Traffic received on access-uplink SAP and sent out of another access SAP uses ring MBS buffer pool for MBS buffers on access-uplink SAP ingress and the non-ring MBS buffer pool for MBS buffers on access SAP egress. In this case, ring high-slope and ring low-slope is used on access-uplink ingress and non-ring high-slope and non-ring low-slope is used on access egress.
- Traffic received on access SAP and sent out of another access SAP uses the non-ring MBS pool for MBS buffers for both access SAP ingress and access SAP egress. In this case, non-ring high-slope and non-ring low-slope is used on both access ingress and access egress.

Configuration Guidelines for CBS and MBS

- For configuring the CBS value, the following must be considered:
 - If Jumbo frames need to be accommodated, then CBS must be set to at least a minimum of 10Kbytes. The default value set for the queue allows for jumbo frames. It is recommended to set the CBS to twice the amount of maximum frame size the queues is expected to carry.
 - CBS pool cannot be oversubscribed.
- For configuring the MBS value, the following must be considered:
 - MBS value determines the maximum delay a packet can experience when using that queue. It should be set to a value such that the delay is acceptable.
 - It is recommended to set the minimum value for MBS to be about 4 to 5 times the maximum size of the frame the queue is expected to carry to ensure better scheduling performance.

Slope Policies

The available buffer space is partitioned into buffer pools as described above. The buffers for a queue are allocated from the buffer pool. Slope policies define the RED slope characteristics.

By default, each queue on the port is associated with slope-policy default which disables the high-slope and low-slope for all the queues.



Note: If WRED is not configured, then taildrop is used.

RED Slopes

Operation and Configuration of RED Slopes

On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C each queue provides user with the following:

- An option to use 2 slopes per queue on non-ring ports - high-priority RED slope and a low-priority RED slope
- An option to use 4 slopes per queue on ring ports – non-ring high-priority RED slope, a non-ring low-priority RED slope, a ring high-priority RED slope, and a ring low-priority RED slope

The high-priority RED slope manages access to the shared portion of the buffer pool for high priority or in-profile packets. The low-priority RED slope manages access to the shared portion of the buffer pool for low-priority or out-of-profile packets. Please read the description in section [Buffer pools on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C](#)

By default, the high-priority and low-priority slopes are disabled.

When a queue depth exceeds the queue's CBS, packets received on that queue must contend with other queues exceeding their CBS for shared buffers. To resolve this contention, RED slopes are used to determine buffer availability on a packet by packet basis. A packet that was either classified as high priority or considered in-profile is handled by the high-priority RED slope. This slope should be configured with RED parameters that prioritize buffer availability over packets associated with the low-priority RED slope. Packets that had been classified as low priority or out-of-profile are handled by this low-priority RED slope.


Simplified overview of RED for 7210 SAS-K platform

The following is a simplified overview of how a RED slope determines shared buffer availability on a packet basis:

- The RED function keeps track of shared buffer utilization and shared buffer average utilization.

- At initialization, the utilization is 0 (zero) and the average utilization is 0 (zero).
- When each packet is received, the current average utilization is plotted on the slope to determine the packet's discard probability.
- A random number is generated associated with the packet and is compared to the discard probability.
- The lower the discard probability, the lower the chances are that the random number is within the discard range.
- If the random number is within the range, the packet is discarded which results in no change to the utilization or average utilization of the shared buffers.
- A packet is discarded if the utilization variable is equal to the shared buffer size or if the utilized CBS (actually in use by queues, not just defined by the CBS) is oversubscribed and has stolen buffers from the shared size, lowering the effective shared buffer size equal to the shared buffer utilization size.
- The new shared buffer average utilization is used as the shared buffer average utilization next time a packet's probability is plotted on the RED slope.
- When a packet is removed from a queue (if the buffers returned to the buffer pool are from the shared buffers), the shared buffer utilization is reduced by the amount of buffers returned. If the buffers are from the CBS portion of the queue, the returned buffers do not result in a change in the shared buffer utilization.

Figure 3: RED Slope Characteristics



A RED slope itself is a graph with an X (horizontal) and Y (vertical) axis. The X-axis plots the percentage of shared buffer average utilization, going from 0 to 100 percent. The Y-axis plots the probability of packet discard marked as 0 to 1. The actual slope can be defined as four sections in (X, Y) points (Figure 3):

- Section A is (0, 0) to (start-avg, 0). This is the part of the slope that the packet discard value is always zero, preventing the RED function from discarding packets when the shared buffer average utilization falls between 0 and start-avg.
- Section B is (start-avg, 0) to (max-avg, max-prob). This part of the slope describes a linear slope where packet discard probability increases from zero to max-prob.
- Section C is (max-avg, max-prob) to (max-avg, 1). This part of the slope describes the instantaneous increase of packet discard probability from max-prob to one. A packet discard probability of 1 results in an automatic discard of the packet.
- Section D is (max-avg, 1) to (100%, 1). On this part of the slope, the shared buffer average utilization value of max-avg to 100% results in a packet discard probability of 1.

Plotting any value of shared buffer average utilization will result in a value for packet discard probability from 0 to 1. Changing the values for start-avg, max-avg and max-prob allows the adaptation of the RED slope to the needs of the different queues (for example: access port queues) using the shared portion of the buffer pool, including disabling the RED slope.

Slope Policy Parameters

The elements required to define a slope policy are:

- A unique policy ID
- The high-slope (for in-profile packets) and low-slope (for out-of-profile packets) per queue. And configurable parameters on each slope are start-avg, max-avg, and max-prob.
- The ring and non-ring high and low slopes for access-uplink port egress.
- On 7210 SAS-K 2F4T6C, the ring and non-ring high and low slopes for network port egress.

Table 18: Default slope policy definition for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Parameter	Description
high-slope	start-avg 70 Percent max-avg 90 Percent max-prob 80 Percent

Schedulers

Table 18: Default slope policy definition for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Parameter	Description
low-slope	start-avg 50 Percent max-avg 75 Percent max-prob 80 Percent
high-slope-ring	start-avg 70 Percent max-avg 90 Percent max-prob 80 Percent
low-slope-ring	start-avg 50 Percent max-avg 75 Percent max-prob 80 Percent

Schedulers

Scheduler on 7210 SAS-K

7210 SAS-K platforms support Strict Priority and WFQ mode of scheduling or a mix of both.

On 7210 SAS-K 2F2T1C, Schedulers are used at SAP ingress, SAP egress, Access Uplink Port ingress and Access-uplink port egress.

On 7210 SAS-K 2F4T6C, Schedulers are used at SAP ingress, SAP egress, Network port ingress, network port egress, Access Uplink Port ingress and Access-uplink port egress.

The scheduler uses 2 loops - the CIR loop and PIR loop, each with 4 priorities. The configured priority of the queue determines the service order of the queue in the CIR loop and the PIR loop. The scheduler first goes through the CIR loop, where it services all the queues which are operating at less than CIR rate according to their priority (that is, higher priority queues get services earlier than lower priority queues). It then goes through the PIR loop, where it services all the queues which are operating above the CIR rate (but less than PIR rate) according to their priority (that is, higher priority queues get services earlier than lower priority queues). If there are multiple queues configured with the same priority, in the CIR

loop the queues are scheduled using WFQ, with the configured weight (that is, pir-weight) of the queue used to determine the proportion of the available bandwidth that is given to the queue. In the PIR loop, the queues are scheduled using WFQ, with the configured weight (that is, pir-weight) of the queue used to determine the proportion of the available bandwidth that is given to the queue (using WFQ).

CPU Queues on SAS-K

The packets that are destined to the CPU are prioritized based on the application. Some of the applications that are prioritized are Layer 2 data packets (a copy of which is sent to CPU for MAC learning), EFM, CFM, STP, LACP, ICMP, etc. The packets destined to the CPU are classified internally and are put into the correct CPU queue. These packets are rate-limited to prevent DoS attacks. The software programs the classification entries to identify these packets and assigns appropriate bandwidth and priority to them. It is not configurable by the user.

Egress Port Rate Limiting

This feature allows the user to limit the bandwidth available on the egress of the port to a value less than the maximum possible link bandwidth. On some platforms, it also allows the user to control the amount of burst sent out.

Forwarding Classes

7210 SAS devices support multiple forwarding classes and class-based queuing, so the concept of forwarding classes is common to all of the QoS policies. Each forwarding class (also called Class of Service (CoS)) is important only in relation to the other forwarding classes. A forwarding class provides network elements a method to weigh the relative importance of one packet over another in a different forwarding class.

Queues are created for a specific forwarding class to determine the manner in which the queue output is scheduled. The forwarding class of the packet, along with the in-profile or out-of-profile state, determines how the packet is queued and handled (the per hop behavior (PHB)) at each hop along its path to a destination egress point. 7210 SAS devices support eight (8) forwarding classes (Table 19).

Table 19: Forwarding Classes

FC-ID	FC Name	FC Designation	DiffServ Name	Notes
7	Network Control	NC	NC2	Intended for network control traffic.
6	High-1	H1	NC1	Intended for a second network control class or delay/jitter sensitive traffic.
5	Expedited	EF	EF	Intended for delay/jitter sensitive traffic.
4	High-2	H2	AF4	Intended for delay/jitter sensitive traffic.
3	Low-1	L1	AF2	Intended for assured traffic. Also is the default priority for network management traffic.
2	Assured	AF	AF1	Intended for assured traffic.
1	Low-2	L2	CS1	Intended for BE traffic.
0	Best Effort	BE	BE	

Note that Table 19 presents the default definitions for the forwarding classes. The forwarding class behavior, in terms of ingress marking interpretation and egress marking, can be changed by QoS Policies.

Forwarding-Class To Queue ID Mapping

There are 8 forwarding classes supported on 7210 SAS devices. Each of these FC is mapped to a specific queue. By mapping FC to different queues the differential treatment is imparted to various classes of traffic.

FC to Queue ID mapping

On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, user has an option to define up to 8 queues with an option to define the FC to queue mapping in service ingress policy, service egress policy, network qos policy and network queue policy.

Preclassification on 7210 SAS-K 2F4T6C

On SAS-K 2F4T6C the front-panel ports oversubscribe the capacity of the forwarding processor. A system defined pre-classification scheme is implemented (it is not user configurable) to prioritize ingress packets for processing by the forwarding processor. It prioritizes packets based on Dot1p, DSCP and identifies some of the untagged L2 control protocols into a high priority and low priority queues maintained on ingress per port. The forwarding processor processes the high priority queue across all the ports before servicing packets from the lower priority queues. In addition, the network ports are favored over access ports by allocating more weight-age to the network ports during scheduling.

Preclassification on 7210 SAS-K 2F4T6C

On 7210 SAS-K 2F4T6C, the front-panel ports oversubscribe the capacity of the forwarding processor. A system defined pre-classification scheme is implemented (it is not user configurable) to prioritize ingress packets for processing by the forwarding processor. It prioritizes packets based on Dot1p, DSCP and identifies some of the untagged L2 control protocols into a high priority and low priority queues maintained on ingress per port. The forwarding processor processes the high priority queue across all the ports before servicing packets from the lower priority queues. In addition, the network ports are favored over access ports by allocating more weight-age to the network ports during scheduling.

QoS Policy Entities

Services are configured with default QoS policies. Additional policies must be explicitly created and associated. There is one default service ingress QoS policy, one default service egress policy, one default access egress QoS policy, one default network QoS policy and one default port scheduler policy. Only one ingress QoS policy and one egress QoS policy can be applied to a SAP or access/access-uplink port or network port.

When you create a new QoS policy, default values are provided for most parameters with the exception of the policy ID, descriptions. Each policy has a scope, default action, a description, and meters for ingress policies and queues for egress policies. By default, all forwarding classes are mapped to Queue 1.

QoS policies can be applied to the following service types:

- Epipe and VPLS
 - On 7210 SAS-K 2F2T1C, SAP ingress policies and SAP egress policies are supported on an Epipe access service access point (SAP), VPLS access SAP, RVPLS SAP, and IES SAP.
 - On 7210 SAS-K 2F4T6C, SAP ingress policies and SAP egress policies are supported on an Epipe service access point (SAP), VPLS access SAP, RVPLS SAP, IES access SAP and VPRN access SAP.

QoS policies can be applied to the following entities:

- Network QoS policy on access uplink port (7210 SAS-K 2F2T1C, 7210 SAS-K 2F4T6C)
- Network queue policy (egress) on access uplink port (7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C);
- Network QoS policy on network port (7210 SAS-K 2F4T6C only)
- Network queue policy (egress) on network port. (7210 SAS-K 2F4T6C only)

Configuration Notes

The following information describes QoS implementation caveats:

- Creating additional QoS policies is optional.
- Default policies are created for service ingress, service egress, access service egress, network, network queue, slope, remark, dot1p and DSCP classification and port scheduler. (the policy types created varies across the platforms)

- Associating a service or access or access uplink with a QoS policy other than the default policy is optional.

Configuration Notes

Discard Eligibility Indicator (DEI) based Classification and Marking

In This Section

This section provides information about the Discard Eligibility Indicator (DEI) feature that describes the requirements for DEI based classification and marking for 7210 platforms.



Note: DEI classification and marking is applicable to 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Topics in this section include:

- [DEI based Classification and Marking](#)

DEI based Classification and Marking

DEI based Classification on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

DEI bit in the received packet can be used to assign the ingress profile for the packet. If in the received packet, DEI = 0, then the packet is considered to be GREEN or in-profile and if DEI = 1, then the packet is considered to be YELLOW or out-of-profile. Use of DEI bit for ingress classification can be enabled per FC. For a given FC, if DEI bit is used for ingress profile assignment, then the profile defined in the ingress classification entry is ignored. For more information, see [“Queue – Ingress Profile Assignment”](#) to understand the behavior when profile is assigned to the packet on ingress.

On 7210 SAS-K 2F2T1C, DEI based classification is supported on access SAP ingress and access-uplink port ingress.

On 7210 SAS-K 2F4T6C, DEI based classification is supported on access SAP ingress, network port ingress and access-uplink port ingress.

DEI based marking

DEI bit can be used to mark the packet to carry the profile, assigned by an operator's trusted node at the ingress to the carrier's network, to the subsequent nodes in the network. It allows high-priority in-profile packet to be allocated appropriate resources by all the network nodes on the path to the final destination. Similarly, it allows out-of-profile packets to be treated with less preference compared to in-profile packets by all the network nodes on the path to the final destination.

The following support is available for DEI based marking:

- On both, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, option to mark DEI bits for access SAP egress on access ports.
- On both, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, option to mark DEI bits for port egress on access uplink ports.
- On 7210 SAS-K 2F4T6C, Option to mark DEI bits for port egress on network ports.
- By default, in-profile packets are marked with DEI bit of 0 and out-of-profile packets are marked with DEI bit of 1. The user has an option to mark all the packets belonging to a FC to the same DEI value irrespective of its profile using the “force-de-mark” option.
- DEI bits can be marked only if the remark policy of remark-type dot1p or dot1p-dscp is used.



Note: For information on the CLI commands for DEI, see [Network QoS Policy Command Reference](#), [Service Egress Policy Command Reference](#) and [Service SAP QoS Policy Command Reference](#).

Port Level Egress Rate-Limiting

In This Section

This section provides information to configure port level egress-rate using the command line interface.

Topics in this section include:

- [Overview](#)
- [Basic Configurations](#)
- [Configuration Descriptions](#)

Overview

Egress port rate limiting allows the device to limit the traffic that egresses through a port to a value less than the available link bandwidth.

This feature is useful when connecting the 7210 SAS to an Ethernet-over-SDH (EoSDH) or microwave network, where the network allocates predetermined bandwidth to the nodes connecting into it, based on the transport bandwidth requirement. When connecting to such a network it is important that the traffic sent into the SDH node does not exceed the configured values, since the SDH network does not have QoS capabilities and buffers required to prioritize the ingress traffic.

Egress rate attributes include:

- Allows for per port configuration of the maximum egress port rate, using the egress-rate CLI command.
- Ethernet ports configured as access and access uplink support this feature.
- The port scheduler distributes the available maximum egress bandwidth based on the CIR/PIR configuration parameters provisioned for the queues.
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C the burst parameter is not user configurable and set to a default by software.

Basic Configurations

- When ports are members of a LAG, all the ports use the same value for the egress-rate and the max-burst parameters.
- If frame overhead accounting (also known as Frame-based accounting) is enabled, then queue scheduler accounts for the Ethernet frame overhead.
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, when an egress-rate sub-rate value is given, the access-uplink port egress queue rates that are specified using percentages will use the egress-rate value instead of the port bandwidth if egress rate is lesser than port bandwidth to configure the appropriate queue rates. Configuration of egress port rate to different values will result in a corresponding dynamic adjustment of rates for the egress queues configured on access-uplink ports.
- On 7210 SAS-K 2F4T6C, when an egress-rate sub-rate value is given, the network port egress queue rates that are specified using percentages will use the egress-rate value instead of the port bandwidth if egress rate is lesser than port bandwidth to configure the appropriate queue rates. Configuration of egress port rate to different values will result in a corresponding dynamic adjustment of rates for the egress queues configured on network ports.
- When the egress-rate sub-rate value is set, CBS/MBS of the associated network queues is not modified automatically. On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C user has an option to change the CBS/MBS values, if need be.

Basic Configurations

To apply port level rate-limiting, perform the following:

- The **egress-rate** command is present in the ***A:Dut-1>config>port>ethernet** context.
- The **egress-rate** configures the maximum rate (in kbps).
- By default there is no explicit egress-rate command set on port and the port operates at the maximum line-rate speed it is operating at.

The following displays the egress-rate configuration for a port:

```
*A:Dut-1>config>port# info
-----
    ethernet
      egress-rate 120000
    exit
    no shutdown
-----
*A:Dut-1>config>port#
```

Modifying Port Level Egress-Rate Command

To modify egress-rate parameters you can simply apply a egress-rate command with new egress-rate.

Removing Port Level Egress-Rate Command

To remove egress-rate command from a port, use the **no** option with the **egress-rate** command. The rate for the egress-rate option and max-burst should not be used in this case.

CLI Syntax: `config>port>ethernet# no egress-rate`

The following displays the removal of egress-rate configuration from a port:

```
*A:Dut-1>config>port# no ethernet egress-rate
*A:Dut-1>config>port# info
-----
      ethernet
      exit
      no shutdown
-----
*A:Dut-1>config>port#
```

Default Egress-Rate Values

By default no egress-rate is configured for a port. For more information on the CLI and description, see [Port Level Egress-Rate Command Reference](#).

Basic Configurations

Port Level Egress-Rate Command Reference

Command Hierarchies

Configuration Commands for 7210 SAS-K 2F2T1C devices

```

— config
  — port
    — ethernet
      — egress-rate sub-rate
      — no egress-rate
  
```

Show Commands

```

— show
  — port [port-id]
  
```

Configuration Descriptions

Configuration Commands

egress-rate

Syntax egress-rate *sub-rate*
no egress-rate

Context config>port>ethernet

Description Platforms supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

This command configures maximum egress rate for a port. The egress-rate is configured as kbps.

The **no** form of the command removes egress-rate from the port.

Port Level Egress-Rate Command Reference

Default No egress-rate and max-burst is configured for the port.

Parameters *sub-rate* — Specifies an integer value between 1 and 1000000 kbps.

Show Commands

port

Syntax **port** [*port-id*]

Context show

Description This command displays Egress-Rate and Max-Burst value set for port along with other details of the port.

Parameters *port-id* — Displays information about the specific port ID.

Output

Sample Output

```
*A:dut-1>config>qos>network-queue# show port 1/1/1
=====
Ethernet Interface
=====
Description           : 10/100/Gig Ethernet SFP
Interface             : 1/1/1                      Oper Speed       : 1 Gbps
Link-level            : Ethernet                  Config Speed    : 1 Gbps
Admin State           : up                       Oper Duplex      : full
Oper State            : up                       Config Duplex    : full
Physical Link         : Yes                      MTU              : 1514
IfIndex               : 35684352                 Hold time up    : 0 seconds
Last State Change     : 01/17/2011 04:05:37      Hold time down  : 0 seconds
Last Cleared Time     : N/A

Configured Mode       : access                    Encap Type       : null
Dot1Q Ethertype      : 0x8100                   QinQ Ethertype  : 0x8100
Net. Egr. Queue Pol  : default                   Access Egr. Qos *: 1
Egr. Sched. Pol      : default                   Network Qos Pol : n/a
Auto-negotiate       : limited                   MDI/MDX         : MDI
Accounting Policy    : None                      Collect-stats   : Disabled
Egress Rate          : Default                   Max Burst       : Default
Uplink               : No

Down-when-looped     : Disabled                  Keep-alive       : 10
Loop Detected        : False                     Retry            : 120

Configured Address   : 00:78:76:45:54:02
Hardware Address     : 00:78:76:45:54:02
Cfg Alarm            :
Alarm Status         :
```

Port Level Egress Rate-Limiting

Transceiver Data

Transceiver Type : SFP
Model Number : 3HE00027AAAA02 ALA IPUIAELDAB=
TX Laser Wavelength: 850 nm Diag Capable : yes
Connector Code : LC Vendor OUI : 00:0a:1d
Manufacture date : 2008/08/10 Media : Ethernet
Serial Number : OPCPCH08052638
Part Number : TRPAG1SXLAES-TM
Optical Compliance : GIGE-SX
Link Length support: 550m for 50u MMF; 280m for 62.5u MMF;

Traffic Statistics

```
=====
                                     Input           Output
-----
Octets                               0             0
Packets                              0             0
Errors                               0             0
=====
```

* indicates that the corresponding row element may have been truncated.

Port Statistics

```
=====
                                     Input           Output
-----
Unicast Packets                       0             0
Multicast Packets                      0             0
Broadcast Packets                      0             0
Discards                               0             0
Unknown Proto Discards                 0
=====
```

Ethernet-like Medium Statistics

```
=====
Alignment Errors : 0 Sngl Collisions : 0
FCS Errors : 0 Mult Collisions : 0
SQE Test Errors : 0 Late Collisions : 0
CSE : 0 Excess Collisns : 0
Too long Frames : 0 Int MAC Tx Errs : 0
Symbol Errors : 0 Int MAC Rx Errs : 0
=====
```

*A:dut-1>config>qos>network-queue#

Port Level Egress-Rate Command Reference

Frame Based Accounting

In This Section

This section provides information to configure frame-based accounting using the command line interface.

Topics in this section include:

- [Overview](#)
- [Configuration Descriptions](#)

Overview

This feature when enabled let QoS policies to account for the Ethernet frame overhead (for example, it accounts for the IFG (inter-frame gap) and the preamble). Typically, the IFG and preamble constitutes about $12 + 8 = 20$ bytes. The QoS policer/meter and shaper uses this overhead for Ethernet ports when allocating bandwidth.

Frame Based Accounting

On 7210 SAS-K platforms, a configurable CLI command enables accounting of the frame overhead per port. This command affects the behavior of the ingress queue and egress queue rate of all the SAPs configured on the port. When disabled, the ingress and egress queue rates along with port egress-rate do not account for the Ethernet frame overhead. When enabled, the ingress and egress queue rates along with port egress-rate account for the Ethernet frame overhead. By default frame-based accounting is disabled for the port.

Accounting records and statistics account for frame overhead for SAPs configured on the port when FBA is enabled on the port.

On 7210 SAS-K 2F2T1C, frame-based accounting is supported on both access port and access-uplink port.

On 7210 SAS-K 2F4T6C, frame-based accounting is supported on network port, access port and access-uplink port.

Enabling and Disabling Frame-Based Accounting

On 7210 SAS-K platforms, frame-based accounting is supported per port with the capability to enable and disable it per port for both ingress and egress. In other words, it not possible to enable/disable it only for ingress or for egress, both can be enabled together or disabled together.

To enable frame-based-accounting for both ingress and egress on a port, execute the command

— config> port> ethernet> frame-based-accounting

To disable frame-based-accounting for both ingress and egress on a port, execute the command

— config>port> ethernet> frame-based-accounting.

The following output displays the enabling of frame-based-accounting:

```
*A:Dut-1>config>port>ethernet>#info detail
-----
...snipped...
frame-based-accounting;
... snipped ...
-----
*A:Dut-1>config>port>ethernet #
```

The following output displays the disabling of frame-based-accounting:

```
*A:Dut-1>config>port>ethernet># info detail-----
-----
...snipped...
no frame-based-accounting;
... snipped ...
-----
*A:Dut-1>config>port>ethernet #
```

For more information about the command please refer to 7210 SAS-D,7210 SAS-E,7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C Interfaces Guide.

Enabling and Disabling Frame-Based Accounting on 7210 SAS-K

On SAS-K platforms, frame-based accounting is supported per port with the capability to enable and disable it per port for both ingress and egress. In other words, it not possible to enable/disable it only for ingress or for egress, both can be enabled together or disabled together.

To enable frame-based-accounting for both ingress and egress on a port, execute the command **config>port> ethernet> frame-based-accounting**.

To disable frame-based-accounting for both ingress and egress on a port, execute the command **config>port> ethernet> no frame-based-accounting**.

CLI Syntax: `config>port> ethernet> frame-based-accounting`

The following output displays the enabling of frame-based-accounting:

```
*A:Dut-1>config>port>ethernet># info detail
-----
...snipped...
frame-based-accounting;
... snipped ...
-----
*A:Dut-1>config>port>ethernet #
```

The following output displays the disabling of frame-based-accounting:

```
*A:Dut-1>config>port>ethernet># info detail
-----
...snipped...
frame-based-accounting;
... snipped ...
-----
*A:Dut-1>config>port>ethernet #
```

For more information about the command please refer to 7210 SAS-D,7210 SAS-E,7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C Interfaces Guide.

Configuration Commands for 7210 SAS-K 2F4T6C

- config
 - port
 - ethernet
 - **frame-based-accounting**
 - **no frame-based-accounting**

Frame Based Accounting Command Reference

frame-based-accounting

Syntax	frame-based-accounting no frame-based-accounting
Context	config>port>ethernet
Description	<p>Platforms Supported: 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T</p> <p>This command configures per port frame-based accounting. It can be enabled or disabled on each port.</p> <p>When enabled, all the shapers rates and queues statistics on that port also account for the Ethernet Layer 1 overhead (of 20 bytes) in both ingress and egress direction. In other words all ingress queue shaper rates, egress queue shaper rates and aggregate SAP shaper rate account for the ethernet overhead.</p> <p>The no form of the command disables frame-based-accounting.</p>
Default	no frame-based-accounting

Configuration Descriptions

Show Commands

sap-ingress

Syntax	sap-ingress [<i>policy-id</i>] [association] match-criteria [detail]
Context	show>qos
Description	This command displays accounting status of a sap-ingress policy along with other details of the policy. When frame-based-accounting is enabled accounting is shown as frame-based otherwise packet-based.
Parameters	<p><i>policy-id</i> — Displays information about the specific policy ID.</p> <p>associations — Displays the associations of the sap-ingress policy.</p> <p>match-criteria — Displays the match criteria of the sap-ingress policy.</p> <p>detail — Displays the detailed information of the sap-ingress policy.</p>
Output	

Frame Based Accounting Command Reference

Sample Output

```
*A:Dut-1# show qos sap-ingress 1
=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (1)
-----
=====
*A:Dut-1#
```

Sample Output for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

```
*A:dut-a>show>qos# sap-ingress 1
=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (1)
-----
Policy-id           : 1           Scope           : Template
Default FC         : be
Criteria-type      : None
Mac Sub-Criteria   : None           IP Sub-Criteria  : None
IPv6 Enabled       : False
DOT1P Class Policy Id : 0           DSCP Class Policy Id : 0
MPLS Lsp Exp Class Policy*: 0
Name               : default
Description        : Default SAP ingress QoS policy.
=====
* indicates that the corresponding row element may have been truncated.
*A:dut-a>show>qos#
```

network

- Syntax** `network [policy-id] [detail]`
- Context** `show>qos`
- Description** This command displays the accounting status of a network qos policy along with other details of the policy. When frame-based-accounting is enabled accounting is shown as frame-based otherwise packet-based.
- Parameters** *policy-id* — Displays information about the specific policy ID.
detail — Displays the detail policy information.
- Output**

Sample output for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

```
*A:dut-a>show>qos# network 1
```

```

=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id          : 1
Egr Remark        : False                Egr Rem Plcy      : N/A
Forward Class     : be                   Profile           : None
Scope             : Template
DOT1P Class Poli* : 1                    DSCP Class Polic* : 0
MPLS Lsp Exp Cla* : 0
Description       : Default network-port QoS policy.
=====
* indicates that the corresponding row element may have been truncated.
*A:dut-a>show>qos#
*A:dut-a>show>qos# network 1

=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id          : 1
Egr Remark        : False                Egr Rem Plcy      : N/A
Forward Class     : be                   Profile           : None
Scope             : Template
DOT1P Class Poli* : 1                    DSCP Class Polic* : 0
MPLS Lsp Exp Cla* : 0
Description       : Default network-port QoS policy.
=====
* indicates that the corresponding row element may have been truncated.
*A:dut-a>show>qos#

```

network-queue

- Syntax** **network-queue** [*network-queue-policy-name*] [**detail**]
- Context** show>qos
- Description** This command displays accounting status of a network-queue policy along with other details of the policy. When frame-based-accounting is enabled accounting is shown as frame-based otherwise packet-based.
- Parameters** *network-queue-policy-name* — Displays information about the specific Network queue policy.
detail — Displays the detailed policy information.
- Output**

Sample Output

```
*A:Dut-1# show qos network-queue default
```

Frame Based Accounting Command Reference

```
=====
QoS Network Queue Policy
=====
-----
Network Queue Policy (default)
-----
Policy          : default
Accounting      : frame-based
Description     : Default network queue QoS policy.
-----
Associations
-----
Port-id : 1/1/6
Port-id : 1/1/7
Port-id : 1/1/8
Port-id : 1/1/9
Port-id : 1/1/10
Port-id : 1/1/11
Port-id : 1/1/12
Port-id : 1/1/13
Port-id : 1/1/14
Port-id : 1/1/15
Port-id : 1/1/16
Port-id : 1/1/17
Port-id : 1/1/18
Port-id : 1/1/20
Port-id : 1/1/21
Port-id : 1/1/22
Port-id : 1/1/23
Port-id : 1/1/24
=====
*A:Dut-1#
```

DSCP, Dot1p, and, MPLS EXP classification policy

In This Section

This section provides information to configure DSCP classification policy, Dot1p classification policy and MPLS EXP classification policy using the command line interface.

Topics in this section include:

- [Overview](#)
- [Configuration Descriptions](#)

Overview

These policies allow user to define a policy/template that maps the packet priority bits, like Dot1p, IP DSCP and MPLS EXP bits, to forwarding class and profile. The template can then be used in a SAP ingress policy to define the ingress classification of flows to forwarding class.

The support available on different 7210 platforms is provided in the table below:

Table 20: Classification policy with platforms support.

Policy name	Support available
Dot1p-classification	7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C with sap-ingress policy and network-qos policy.
DSCP classification	7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C with sap-ingress policy and network-qos policy.

Table 20: Classification policy with platforms support.

Policy name	Support available
MPLS-EXP classification	Only 7210 SAS-K 2F4T6C with network-qos policies

DSCP classification policy

This policy is used define the map of IP DSCP values in the IP header of the packet to forwarding class and profile (ingress profile).

Dot1p classification policy

This policy is used define the map of Dot1p values in the ethernet frame to forwarding class and profile (ingress profile).

MPLS EXP classification policy

This policy is used define the map of MPLS EXP values in the MPLS header of the packet to forwarding class and profile (ingress profile).

Configuration of DSCP classification policy, Dot1p classification policy, and MPLS EXP classification policy

CLI Syntax:

```
config>qos>dot1p-classification 1 create
config>qos>dscp-classification 1 create
config>qos>mpls-lsp-exp-classification 1 create
```

The following output displays the enabling DSCP classification policy, Dot1p classification policy, and MPLS EXP classification policy:

```
*A:K-SASK12>config>qos># info detail
-----
-----exit
dot1p-classification 1 create
```

```

description "Default Dot1P Classification policy"
dot1p 0 fc "be" profile out
dot1p 1 fc "l2" profile in
dot1p 2 fc "af" profile out
dot1p 3 fc "af" profile in
dot1p 4 fc "h2" profile in
dot1p 5 fc "ef" profile in
dot1p 6 fc "h1" profile in
dot1p 7 fc "nc" profile in
exit
dot1p-classification 10 create
no description
exit
dscp-classification 1 create
description "Default DSCP Classification policy"
dscp be fc "be" profile out
dscp ef fc "ef" profile in
dscp cs1 fc "l2" profile in
dscp ncl fc "h1" profile in
dscp nc2 fc "nc" profile in
dscp af11 fc "af" profile in
dscp af12 fc "af" profile out
dscp af41 fc "h2" profile in
exit
dscp-classification 20 create
no description
exit
mpls-lsp-exp-classification 1 create
description "Default MplsLspExp Classification policy"
lsp-exp 0 fc "be" profile out
lsp-exp 1 fc "l2" profile in
lsp-exp 2 fc "af" profile out
lsp-exp 3 fc "af" profile in
lsp-exp 4 fc "h2" profile in
lsp-exp 5 fc "ef" profile in
lsp-exp 6 fc "h1" profile in
lsp-exp 7 fc "nc" profile in
exit
mpls-lsp-exp-classification 20 create
no description
exit

*A:K-SASK12>config>qos>dot1p-classification#
*A:K-SASK12>config>qos>mpls-lsp-exp-classification# info detail
-----
description "Default MplsLspExp Classification policy"
lsp-exp 0 fc "be" profile out
lsp-exp 1 fc "l2" profile in
lsp-exp 2 fc "af" profile out
lsp-exp 3 fc "af" profile in
lsp-exp 4 fc "h2" profile in
lsp-exp 5 fc "ef" profile in
lsp-exp 6 fc "h1" profile in
lsp-exp 7 fc "nc" profile in
-----

*A:K-SASK12>config>qos>mpls-lsp-exp-classification

```

Overview

Command reference

Configuration Commands for 7210 SAS-K 2F2T1C

- config
 - qos
 - **dscp-classification** *classification id* [**create**]
 - **no dscp-classification** *classification id*
 - **description** *description-string*
 - **no description** *description-string*
 - **dscp fc** *fc-name* **profile** {in|out}
 - **no dscp** *dscp-name*
 - **dot1p-classification** *classification id* [**create**]
 - **no dot1p-classification** *classification id*
 - **description** *description-string*
 - **no description** *description-string*
 - **dot1p fc** *fc-name* **profile** {in|out}
 - **no dot1p** *dot1p-priority*

Configuration Commands for 7210 SAS-K 2F4T6C

- config
 - qos
 - **dscp-classification** *classification id* [**create**]
 - **no dscp-classification** *classification id*
 - **description** *description-string*
 - **no description** *description-string*
 - **dscp fc** *fc-name* **profile** {in|out}
 - **no dscp** *dscp-name*
 - **dot1p-classification** *classification id* [**create**]
 - **no dot1p-classification** *classification id*
 - **description** *description-string*
 - **no description** *description-string*
 - **dot1p fc** *fc-name* **profile** {in|out}
 - **no dot1p** *dot1p-priority*
 - **mpls-lsp-exp-classification** *classification id* [**create**]
 - **no mpls-lsp-exp-classification** *classification id*
 - **description** *description-string*
 - **no description** *description-string*
 - **lsp-exp fc** *fc-name* **profile** {in|out}
 - **no lsp-exp** *lsp-exp-priority*

Show Commands

- show
 - qos
 - **dot1p-classification** *policy-id* association [detail]
 - **dscp-classification** *policy-id* association [detail]
 - **mpls-lsp-exp-classification** *policy-id* association [detail]

Operational Commands

- config
 - qos
 - **copy** dot1p-classification *src-pol* *dst-pol* [overwrite]
 - **copy** dscp-classification *src-pol* *dst-pol* [overwrite]

Command Description

dscp-classification

Syntax	dscp-classification <i>classification id</i> [create] no dscp-classification <i>classification id</i>
Context	config>qos
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. This policy allows user to define the map of IP DSCP value in the IP header of the received packet to forwarding class and ingress profile. The no form of the command removes the definition of the policy. The no form of the policy cannot be executed if the policy is in use (i.e. associated with for example a network qos policy).
Default	default policy 1
Parameters	<i>classification-id</i> — Identifier used to uniquely identify the policy. Values [1..65535] create — The keyword used to create

description

Syntax	[no] description <i>description-string</i>
Context	config>qos>dscp-classification config>qos> dot1p-classification config>qos> mpls-lsp-exp-classification
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the context in the configuration file. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Command Description

dscp

Syntax	dscp fc fc [profile {in out}] no dscp dscp-name
Context	config>qos>dscp-classification
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>This command is used to map the IP DSCP value to a forwarding class and also assign the ingress profile to the packet. The configured value is used to match the value in the received packet and assign the configured forwarding class and profile on a exact match.</p> <p>The no form of the command removes the mapping of the IP DSCP value to forwarding class.</p> <p>NOTE: A Default forwarding class is not assigned on executing the no form. The default forwarding class is assigned in the QoS policy with which this policy is associated with.</p>
Default	Default forwarding class is not assigned. the mapping must be explicitly configured.
Parameters	<p>dscp dscp-name — Specify the IP DSCP value to match by configuring the IP DSCP name corresponding to the DSCP value.</p> <p>Default None, the DSCP name must be specified</p> <p>Values</p> <p>be ef cp1 cp2 cp3 cp4 cp5 cp6 cp7 cp9 cs1 cs2 cs3 cs4 cs5 nc1 nc2 af11 af12 af13 af21 af22 af23 af31 af32 af33 af41 af42 af43 cp11 cp13 cp15 cp17 cp19 cp21 cp23 cp25 cp27 cp29 cp31 cp33 cp35 cp37 cp39 cp41 cp42 cp43 cp44 cp45 cp47 cp49 cp50 cp51 cp52 cp53 cp54 cp55 cp57 cp58 cp59 cp60 cp61 cp62 cp63</p> <p>fc fc-name — Specify the system-defined forwarding class name (it is case-sensitive).</p> <p>Default None, the fc name must be specified</p> <p>Values fc-name</p> <p>profile {in out} — All packets assigned to this forwarding class will be considered in or out of profile based on this command. A value of 'in' defines the packet profile as 'in-profile' and a value of 'out' defines the packet profile to be out-of-profile.</p>

dot1p-classification

Syntax	dot1p-classification classification id [create] no dot1p-classification classification id
Context	config>qos
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

DSCP classification, Dot1p classification and MPLS EXP classification policies

This policy allows user to define the map of Dot1p value in the Ethernet header of the received frame to forwarding class and ingress profile.

The no form of the command removes the definition of the policy. The no form of the policy cannot be executed if the policy is in use (i.e. associated with for example a network qos policy).

Default default policy 1

Parameters *classification-id* — Identifier used to uniquely identify the policy

Values [1..65535]

create — The keyword used to create

dot1p

Syntax **dot1p fc fc [profile {in | out}]**
dot1p dot1p-priority

Context config>qos>dot1p-classification

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command is used to map the Dot1p value to a forwarding class and also assign the ingress profile to the packet. The configured value is used to match the value in the received packet and assign the configured forwarding class and profile on a exact match.

The no form of the command removes the mapping of the dot1p value to forwarding class.

NOTE: A Default forwarding class is not assigned on executing the no form. The default forwarding class is assigned in the QoS policy with which this policy is associated with.

Default Default forwarding class is not assigned. the mapping must be explicitly configured.

Parameters **fc fc-name** — Specify the system-defined forwarding class name (it is case-sensitive).

Default None, the fc name must be specified

Values fc-name

profile {in | out} — All packets assigned to this forwarding class will be considered in or out of profile based on this command. A value of 'in' defines the packet profile as 'in-profile' and a value of 'out' defines the packet profile to be out-of-profile.

dot1p-priority — Specify the Dot1p priority value to match.

Values [0..7]

mpls-lsp-exp-classification

Syntax **mpls-lsp-exp-classification classification id [create]**

Command Description

no mpls-lsp-exp-classification *classification id*

Context	config>qos
Description	<p>Platforms Supported: 7210 SAS-K 2F4T6C.</p> <p>This policy allows user to define the map of MPLS EXP value in the MPLS header of the received packet to forwarding class and ingress profile.</p> <p>The no form of the command removes the definition of the policy. The no form of the policy cannot be executed if the policy is in use (i.e. associated with for example a network qos policy).</p>
Default	default policy 1
Parameters	<p><i>classification-id</i> — Identifier used to uniquely identify the policy.</p> <p>Values [1..65535]</p> <p>create — The keyword used to create</p>

lsp-exp

Syntax	lsp-exp <i>mpls-lsp-exp-priority fcfc-name</i> [profile {in out}] no lsp-exp <i>lsp-exp-priority</i>
Context	config>qos>mpls-lsp-exp-classification
Description	<p>Platforms Supported: 7210 SAS-K 2F4T6C.</p> <p>This command is used to map the MPLS EXP value to a forwarding class and also assign the ingress profile to the packet. The configured value is used to match the value in the received packet and assign the configured forwarding class and profile on a exact match.</p> <p>The no form of the command removes the mapping of the MPLS EXP value to forwarding class. NOTE: A Default forwarding class is not assigned on executing the no form. The default forwarding class is assigned in the QoS policy with which this policy is associated with.</p>
Default	default forwarding class is not assigned. the mapping must be explicitly configured.
Parameters	<p><i>fcfc-name</i> — Specify the system-defined forwarding class name (it is case-sensitive).</p> <p>Default None, the fc name must be specified</p> <p>Values fc-name</p> <p>profile {in out} — All packets assigned to this forwarding class will be considered in or out of profile based on this command. A value of 'in' defines the packet profile as 'in-profile' and a value of 'out' defines the packet profile to be out-of-profile.</p> <p>lsp-exp-priority — Specify the LSP EXP value to match.</p> <p>Values [0..7]</p>

DSCP classification policy, Dot1p classification policy, and MPLS EXP classification policy

Configuration Descriptions

Show Commands

dot1p-classification

Syntax	dot1p-classification [<i>policy-id</i>] association dot1p-classification [<i>policy-id</i>] detail
Context	show>qos
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C
Parameters	<i>policy-id</i> — Displays information about the specific policy ID. associations — Displays the associations of the dot1p-classification policy. detail — Displays the detailed information of the dot1p-classification policy.

Output **Sample Output**

Sample Output for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

```
*A:K-SASK12>show>qos# dot1p-classification
=====
SAS DOT1P Classification Maps
=====
Policy-Id Description
-----
1          Default Dot1P Classification policy
=====
*A:K-SASK12>show>qos#
```

dscp-classification

Syntax	dscp-classification [<i>policy-id</i>] association dscp-classification [<i>policy-id</i>] detail
Context	show>qos

DSCP classification policy, Dot1p classification policy, and MPLS EXP classification policy

- Description** **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C
- Parameters**
- policy-id* — Displays information about the specific policy ID.
 - associations** — Displays the associations of the dscp-classification policy.
 - detail** — Displays the detailed information of the dscp-classification policy.

Output **Sample Output**

Sample Output for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

```
*A:K-SASK12>show>qos# dscp-classification

=====
SAS DSCP Classification Maps
=====
Policy-Id Description
-----
1          Default DSCP Classification policy
=====
```

mpls-lsp-exp-classification

- Syntax** **mpls-lsp-exp-classification** [*policy-id*] **association**
mpls-lsp-exp-classification [*policy-id*] **detail**
- Context** show>qos
- Description** **Platforms Supported:** 7210 SAS-K 2F4T6C
- Parameters**
- policy-id* — Displays information about the specific policy ID.
 - associations** — Displays the associations of the mpls-lsp-exp-classification policy.
 - detail** — Displays the detailed information of the mpls-lsp-exp-classification policy.

Output **Sample Output**

Sample Output for 7210 SAS-K 2F4T6C

```
=====
SAS MPLS-LSP-EXP Classification Maps
=====
Policy-Id Description
-----
1          Default MplsLspExp Classification policy
=====
*A:K-SASK12>show>qos#
```

Network QoS Policies

In This Section

This section provides information to configure network QoS policies using the command line interface.

Topics in this section include:

- [Overview of Network QoS policy on 7210 SAS-K 2F2T1C](#)
- [Basic Configurations](#)
- [Default Network Policy Values](#)
- [Service Management Tasks](#)

Overview of Network QoS policy on 7210 SAS-K 2F2T1C

Network QoS policy has an ingress and egress component, which define the QoS processing behavior to be provided for packets that ingress the access-uplink port and egress the access-uplink port respectively.

The ingress component of the policy defines how the IP DSCP and Dot1p values using the DSCP and dot1p classification policies are mapped to internal forwarding class and profile state for 7210 SAS-K 2F2T1C. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the system. The mapping on each access uplink port defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the access uplink ports. It also defines the bandwidth-limiting parameters for the traffic mapped to each forwarding classes. Traffic mapped to each forwarding class can be limited to configurable bandwidth values using separate queues for unicast traffic and multi-point traffic.

Overview of Network QoS policy on 7210 SAS-K 2F2T1C

7210 SAS platforms provide different mechanisms to limit the bandwidth per forwarding class. On 7210 SAS-K 2F2T1C, users can use the queue with packet buffers and a rate shaper to limit and shape the traffic per forwarding class. Use of queue with shapers typically allows for better TCP traffic behavior in the network.

The egress component of the network QoS policy defines the marking values associated with each forwarding class.

On 7210 SAS-K, user has an option to define the number of queues to use per access-uplink port and map the forwarding class to queues. By default, network QoS policy 1 is used for access-uplink ports, until an explicit policy is associated. The default policy creates 8 queues per access-uplink port. The queues are assigned default values for all the parameters defined with the default policy.

Access uplink port egress marking support:

- For packets sent out of an access-uplink port, the network QoS policy defines the marking values (for example: IEEE 802.1p bits, etc.) to use based on the forwarding class and the profile state.
- The default map of FC to marking values (for example: 802.1p bits) is as shown in default network qos policy, policy id 1.
- All non-default network qos policies inherits the default map and can be modified by the user.
- On 7210 SAS-K 2F2T1C, remarking can be enabled or disabled on access uplink ports.
- On 7210 SAS-K 2F2T1C, option is available to map forwarding class and profile to either IP DSCP and Dot1p bits along with DEI bit.

New (non-default) network policy parameters can be modified. The no form of the command reverts the object to the default values.

Changes made to a policy are applied immediately to all access uplink ports where the policy is applied. For this reason, when a policy requires several changes, it is recommended that you copy the policy to a work area policy-id. The work-in-progress copy can be modified until all the changes are made and then the original policy-id can be overwritten with the config qos copy command.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your devices, refer to CLI Usage chapter in the 7210 SAS D, E, K OS Basic System Configuration Guide.

Resource Allocation for Network QoS policy for 7210 SAS-K 2F2T1C

The queues required for access-uplink port egress is allocated for the egress queue system pool. Queues from the egress queue system pool are allocated for per SAP egress queues and per access-uplink port egress queues.

The Dot1p policy and DSCP policy resources used for network qos ingress FC assignment on access-uplink port, is shared with the per SAP ingress classification criteria from the system pool.

Basic Configurations

A basic network QoS policy must conform to the following:

- Each network QoS policy must have a unique policy ID.
- Specify the default-action.
- Have a QoS policy scope of template or exclusive.
- Based on the 7210 SAS platform being used, have at least one default unicast forwarding class meter/queue.
- Based on the 7210 SAS platform being used, have at least one multipoint forwarding class meter/queue.

Create a Network QoS Policy on 7210 SAS-K 2F2T1C

Configuring and applying QoS policies other than the default policy is optional. A default network policy of the appropriate type is applied to each access uplink port.

To create an network QoS policy, define the following:

- A network policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- Egress Marking/ remarking - Specifies the egress FC to marking value (for example: IEEE 802.1p, etc) map. Otherwise, the default values are applied.
 - Remarkings — If enabled, this command remarks ALL packets that egress on the specified access uplink port. The remarking is based on the forwarding class to marking values mapping defined under the egress node of the network QoS policy. On 7210 SAS-K 2F2T1C, user has an option to enable it or disable it.

Basic Configurations

- Forwarding class criteria — The forwarding class name represents an egress queue. Specify forwarding class criteria to define the marking criteria of packets flowing through it.
- Marking Value— The marking (for example: IEEE 802.1p) value is used for all packets requiring marking that egress on this forwarding class queue that are in or out of profile.
- Ingress criteria — Specifies the criteria to use for forwarding class mapping for all packets.
 - Default action — Defines the default action to be taken for packets that have an undefined Dot1p bits set. The default-action specifies the forwarding class to which such packets are assigned.
 - Dot1p - On 7210 SAS-K, user has an option to specify either Dot1p or IP DSCP to forwarding class mapping for all packets. Ingress traffic that matches the specified criteria are assigned to the corresponding forwarding class.

Use the following CLI syntax to create a network QoS policy for 7210 SAS-K 2F2T1C:

```
CLI Syntax:  config>qos#
                Config>qos>network <network-policy id> create
                description description-string
                scope {exclusive|template}
                egress
                    [no] remarking
                    remark <remark-policy id>
                    no remark
                ingress
                    default-action fc <fc name>
                    dot1p-classification <dot1p-classification id>
                    dscp-classification <dscp-classification id>
                    queue <queue id> create
                fc <fc-name> create
                    queue <queue-id>
                    multicast-queue <queue-id>
                    [no] use-dei
```

The following commands associated a network QoS policy with the access-uplink port.

```
CLI Syntax:  config>port
                ethernet
                    access
                        uplink
                            qos network-policy-id
```

```
CLI Syntax:  config>router
                interface interface-name
                    qos network-policy-id
```

The following output displays the configuration for uplink port 1/1/1 with network policy 600 applied to the interface.

```
A:ALA-7>config# info
#-----
echo "Port Configuration"
#-----
    port 1/1/1
      shutdown
      description "port 1/1/1"
      ethernet
        mode access uplink
        access
          uplink
            qos 600
          exit
        exit
      exit
    exit
  ...
#-----
A:ALA-7>config#
```

Default Network Policy Values on 7210 SAS-K 2F2T1C

The default network policy access uplink ports is identified as policy-id 1. Default policies cannot be modified or deleted. The following displays default network policy parameters:

```
7210SAS>config>qos>network# info
-----
description "Default network-port QoS policy."
ingress
  dot1p-classification 1
  queue 1 create
  exit
  queue 2 create
    rate cir 25
  exit
  queue 3 create
    rate cir 25
  exit
  queue 4 create
    rate cir 25
  exit
  queue 5 create
    rate cir 100
  exit
  queue 6 create
    rate cir 100
  exit
  queue 7 create
    rate cir 10
  exit
  queue 8 create
```

Basic Configurations

```
        rate cir 10
    exit
    fc "af" create
        queue 3
        multicast-queue 3
    exit
    fc "be" create
        queue 1
        multicast-queue 1
    exit
    fc "ef" create
        queue 6
        multicast-queue 6
    exit
    fc "h1" create
        queue 7
        multicast-queue 7
    exit
    fc "h2" create
        queue 5
        multicast-queue 5
    exit
    fc "l1" create
        queue 4
        multicast-queue 4
    exit
    fc "l2" create
        queue 2
        multicast-queue 2
    exit
    fc "nc" create
        queue 8
        multicast-queue 8
    exit
exit
egress
exit
-----
*7210 SAS>config>qos>network#
```

Default Network Policy Values

The default network policy access uplink ports is identified as policy-id **1**. Default policies cannot be modified or deleted. The following displays default network policy parameters:

Table 21: Network Policy ID #1 Defaults for access-uplink port

Field	Default
description	Default network QoS policy.
scope	template

Table 21: Network Policy ID #1 Defaults for access-uplink port (Continued)

Field	Default
ingress	
default-action	fc be profile out
egress	
remarking	No
fc af:	
dot1p-in-profile	3
dot1p-out-profile	2
fc be:	
dot1p-in-profile	0
dot1p-out-profile	0
fc ef:	
dot1p-in-profile	5
dot1p-out-profile	5
fc h1:	
dot1p-in-profile	6
dot1p-out-profile	6
fc h2:	
dot1p-in-profile	4
dot1p-out-profile	4
fc l1:	
dot1p-in-profile	3
dot1p-out-profile	2
fc l2:	
dot1p-in-profile	1
dot1p-out-profile	1
fc nc:	

Table 21: Network Policy ID #1 Defaults for access-uplink port (Continued)

Field	Default
dot1p-in-profile	7
dot1p-out-profile	7

Table 22: Default Network QoS Policy ID #1 for Dot1p to FC Mapping for access-uplink port

Dot1p Value	7210 FC Ingress	Profile
0	be	Out
1	l2	In
2	af	Out
3	af	In
4	h2	In
5	ef	In
6	h1	In
7	nc	In

7210 SAS-K 2F4T6C Network QoS Policy

Overview

Network QoS policy has an ingress and egress component, which define the QoS processing behavior to be provided for packets that ingress the network port and access-uplink port and egress the network port and access-uplink port respectively.

Network QoS policy for access-uplink port on 7210 SAS-K 2F4T6C

The ingress component of the policy defines how the IP DSCP and Dot1p values using the dscp and dot1p classification policies are mapped to internal forwarding class and profile state for 7210 SAS-K 2F4T6C. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the system. The mapping on each access uplink port defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the access uplink ports. It also defines the bandwidth-limiting parameters for the traffic mapped to each forwarding classes. Traffic mapped to each forwarding class can be limited to configurable bandwidth values using separate queues for unicast traffic and multi-point traffic.

NOTE: 7210 SAS platforms provide different mechanisms to limit the bandwidth per forwarding class. On 7210 SAS-K 2F4T6C, users can use the queue with packet buffers and a rate shaper to limit and shape the traffic per forwarding class. Use of queue with shapers typically allows for better TCP traffic behavior in the network.

The egress component of the network QoS policy defines the marking values associated with each forwarding class.

On 7210 SAS-K 2F4T6C, user has an option to define the number of queues to use per access-uplink port and map the forwarding class to queues. By default, network QoS policy 1 is used for access-uplink ports, until an explicit policy is associated. The default policy creates 8 queues per access-uplink port. The queues are assigned default values for all the parameters defined with the default policy.

Access uplink port egress marking support:

- For packets sent out of an access-uplink port, the network QoS policy defines the marking values (for example: IEEE 802.1p bits, etc.) to use based on the forwarding class and the profile state.
- The default map of FC to marking values (for example: 802.1p bits) is as shown in default network qos policy, policy id 1.
- All non-default network qos policies inherits the default map and can be modified by the user.
- On 7210 SAS-K 2F4T6C, remarking can be enabled or disabled on access uplink ports.
- On 7210 SAS-K 2F4T6C, option is available to map forwarding class and profile to either IP DSCP and Dot1p bits along with DEI bit.

New (non-default) network policy parameters can be modified. The no form of the command reverts the object to the default values.

Changes made to a policy are applied immediately to all access uplink ports where the policy is applied. For this reason, when a policy requires several changes, it is recommended that you copy the policy to a work area policy-id. The work-in-progress copy can be modified until all the changes are made and then the original policy-id can be overwritten with the **config qos copy** command.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your devices, refer to **CLI Usage** chapter in the *7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OS Basic System Configuration Guide*.

Network QoS policy for network port on 7210 SAS-K 2F4T6C

The ingress component of the policy defines how the MPLS EXP, IP DSCP and Dot1p values using classification policies are mapped to internal forwarding class and profile state for 7210 SAS-K 2F4T6C. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the system. The mapping on each network port defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the network ports. It also defines the bandwidth-limiting parameters for the traffic mapped to each forwarding classes. Traffic mapped to each forwarding class can be limited to configurable bandwidth values using separate queues for unicast traffic and multi-point traffic.

NOTE: 7210 SAS platforms provide different mechanisms to limit the bandwidth per forwarding class. On 7210 SAS-K 2F4T6C, users can use the queue with packet buffers and a rate shaper to limit and shape the traffic per forwarding class. Use of queue with shapers typically allows for better TCP traffic behavior in the network.

The egress component of the network QoS policy defines the marking values associated with each forwarding class.

On 7210 SAS-K 2F4T6C, user has an option to define the number of queues to use per network port and map the forwarding class to queues. By default, network QoS policy 2 is used for network ports, until an explicit policy is associated. The default policy creates 8 queues per network port. The queues are assigned default values for all the parameters defined with the default policy.

Network port egress marking support:

- For packets sent out of a network port, the network QoS policy defines the marking values (for example: IEEE 802.1p bits, etc.) to use based on the forwarding class and the profile state.

- The default map of FC to marking values (for example: 802.1p bits) is as shown in default network qos policy, policy id 2 for network port.
- All non-default network qos policies inherits the default map and can be modified by the user.
- On 7210 SAS-K 2F4T6C, remarking can be enabled or disabled on access uplink ports.
- On 7210 SAS-K 2F4T6C, option is available to map forwarding class and profile as follows:
 - For MPLS packets - option to mark MPLS EXP values and Dot1p bits along with DEI bit.
 - For IP packets - option to mark IP DSCP and Dot1p bits along with DEI bit.

New (non-default) network policy parameters can be modified. The no form of the command reverts the object to the default values.

Changes made to a policy are applied immediately to all access uplink ports where the policy is applied. For this reason, when a policy requires several changes, it is recommended that you copy the policy to a work area policy-id. The work-in-progress copy can be modified until all the changes are made and then the original policy-id can be overwritten with the **config qos copy** command.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your devices, refer to **CLI Usage** chapter in the *7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C OS Basic System Configuration Guide*.

Resource Allocation for Network QoS policy for 7210 SAS-K 2F4T6C

The queues required for network port and access-uplink port egress is allocated for the egress queue system pool. Queues from the egress queue system pool are shared among SAP egress queues, access-uplink port egress queues and network port egress queues.

The Dot1p policy and DSCP policy resources used for network qos ingress FC assignment on network port and access-uplink port is shared with the per SAP ingress classification criteria from the system pool.

The MPLS EXP policy resources used for network qos ingress FC assignment on network port is allocated from the system pool.

Basic Configurations

A basic network QoS policy must conform to the following:

- Each network QoS policy must have a unique policy ID.
- Specify the default-action.
- Have a QoS policy scope of template or exclusive.
- Based on the 7210 SAS platform being used, have at least one default unicast forwarding class queue.
- Based on the 7210 SAS platform being used, have at least one multi-point forwarding class queue.

Create a Network QoS Policy on 7210 SAS-K 2F4T6C

Create a Network QoS Policy for access-uplink ports on 7210 SAS-K 2F4T6C

For more information, see the [Create a Network QoS Policy for access-uplink ports on 7210 SAS-K 2F4T6C on page 112](#).

Default Network Policy Values for access-uplink port on 7210 SAS-K 2F4T6C

The default network policy for access uplink ports is identified as policy-id **1**. Default policies cannot be modified or deleted. The following displays default network policy parameters:

The following output displays the default network policy ID **1** configuration for access-uplink port on 7210 SAS-K 2F4T6C:

```
*A:K-SASK12>config>qos>network# info
-----
description "Default network-port QoS policy."
ingress
  dot1p-classification 1
  queue 1 create
  exit
  queue 2 create
    rate cir 25
  exit
  queue 3 create
    rate cir 25
  exit
```

```
queue 4 create
    rate cir 25
exit
queue 5 create
    rate cir 100
exit
queue 6 create
    rate cir 100
exit
queue 7 create
    rate cir 10
exit
queue 8 create
    rate cir 10
exit
fc "af" create
    queue 3
    multicast-queue 3
exit
fc "be" create
    queue 1
    multicast-queue 1
exit
fc "ef" create
    queue 6
    multicast-queue 6
exit
fc "h1" create
    queue 7
    multicast-queue 7
exit
fc "h2" create
    queue 5
    multicast-queue 5
exit
fc "l1" create
    queue 4
    multicast-queue 4
exit
fc "l2" create
    queue 2
    multicast-queue 2
exit
fc "nc" create
    queue 8
    multicast-queue 8
exit
exit
egress
exit
-----
*A:K-SASK12>config>qos>network#
```

Create a Network QoS Policy for network ports on 7210 SAS-K 2F4T6C

Configuring and applying QoS policies other than the default policy is optional. A default network policy of the appropriate type is applied to each network port.

To create a network QoS policy, define the following:

- A network policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- Egress Marking/ remarking - Specifies the egress FC to marking value (for example: IEEE 802.1p, etc) map. Otherwise, the default values are applied.
 - Remarking - If enabled, this command remarks ALL packets that egress on the specified access uplink port. The remarking is based on the forwarding class to marking values mapping defined under the egress node of the network QoS policy. On 7210 SAS-K 2F4T6C user has an option to enable it or disable it.
 - Forwarding class criteria - The forwarding class name represents an egress queue. Specify forwarding class criteria to define the marking criteria of packets flowing through it.
 - Marking Value- The marking (for example: IEEE 802.1p) value is used for all packets requiring marking that egress on this forwarding class queue that are in or out of profile.
- Ingress criteria - Specifies the criteria to use for forwarding class mapping for all packets.
 - Default action - Defines the default action to be taken for packets that have an undefined Dot1p bits set. The default-action specifies the forwarding class to which such packets are assigned.
 - On 7210 SAS-K 2F4T6C, user has an option to specify either MPLS EXP and Dot1p and IP DSCP to forwarding class mapping for all packets. Ingress traffic that matches the specified criteria are assigned to the corresponding forwarding class.

Use the following CLI syntax to create a network QoS policy for network port on 7210 SAS-K 2F4T6C.

```
CLI Syntax: >config>qos# network
                network <network-policy-id> [create]
                no network <network-policy-id>
```

The following commands associate a network QoS policy with the network port on 7210 SAS-K 2F4T6C.

```
*A:K-SASK12>config>qos>network# info
```

```

-----
description "Default network QoS policy."
ingress
  dot1p-classification 1
  dscp-classification 1
  mpls-lsp-exp-classification 1
  queue 1 create
  exit
  queue 2 create
    rate cir 25
  exit
  queue 3 create
    rate cir 25
  exit
  queue 4 create
    rate cir 25
  exit
  queue 5 create
    rate cir 100
  exit
  queue 6 create
    rate cir 100
  exit
  queue 7 create
    rate cir 10
  exit
  queue 8 create
    rate cir 10
  exit
  fc "af" create
    queue 3
    multicast-queue 3
  exit
  fc "be" create
    queue 1
    multicast-queue 1
  exit
  fc "ef" create
    queue 6
    multicast-queue 6
  exit
  fc "h1" create
    queue 7
    multicast-queue 7
  exit
  fc "h2" create
    queue 5
    multicast-queue 5
  exit
  fc "ll" create
    queue 4
    multicast-queue 4
  exit
  fc "l2" create
    queue 2
    multicast-queue 2
  exit
  fc "nc" create
    queue 8

```

7210 SAS-K 2F4T6C Network QoS Policy

```
        multicast-queue 8
        exit
    exit
    egress
        remark 2
    exit
-----
*A:K-SASK12>config>qos>network#
```

The following output displays the configuration for network port 1/1/1 with network policy 600 applied to the port.

```
#-----
echo "Port Configuration"
#-----
    port 1/1/1
        shutdown
        ethernet
            network
                qos 2
        exit
    exit
exit
port 1/1/2
    shutdown
    ethernet
        network
            qos 2
    exit
exit
port 1/1/3
    shutdown
    ethernet
        network
            qos 2
    exit
    connection-type copper
exit
port 1/1/4
    shutdown
    ethernet
        network
            qos 2
    exit
exit
port 1/1/5
    ethernet
        mode access
        access
    exit
    mtu 1518
    exit
    no shutdown
exit
#-----
```

Default Network Policy Values on 7210 SAS-K 2F4T6C

The default network policy for network ports is identified as policy-id **2**. Default policies cannot be modified or deleted. The following displays default network policy parameters:

The following output displays the default network policy ID 2 configuration for network port on 7210 SAS-K 2F4T6C:

```
*A:K-SASK12>config>qos>network# info detail
-----
description "Default network QoS policy."
scope template
ingress
  default-action fc be
  dot1p-classification 1
  dscp-classification 1
  mpls-lsp-exp-classification 1
  queue 1 create
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
  exit
  queue 2 create
    rate cir 25 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
  exit
  queue 3 create
    rate cir 25 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
  exit
  queue 4 create
    rate cir 25 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
  exit
  queue 5 create
    rate cir 100 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
```

7210 SAS-K 2F4T6C Network QoS Policy

```
        cbs 24
        weight 1
        priority 1
        slope-policy "default"
    exit
queue 6 create
    rate cir 100 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
exit
queue 7 create
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
exit
queue 8 create
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    weight 1
    priority 1
    slope-policy "default"
exit
fc "af" create
    queue 3
    multicast-queue 3
    no use-dei
exit
fc "be" create
    queue 1
    multicast-queue 1
    no use-dei
exit
fc "ef" create
    queue 6
    multicast-queue 6
    no use-dei
exit
fc "h1" create
    queue 7
    multicast-queue 7
    no use-dei
exit
fc "h2" create
    queue 5
    multicast-queue 5
    no use-dei
exit
fc "l1" create
    queue 4
```

```

        multicast-queue 4
        no use-dei
    exit
    fc "l2" create
        queue 2
        multicast-queue 2
        no use-dei
    exit
    fc "nc" create
        queue 8
        multicast-queue 8
        no use-dei
    exit
exit
egress
    no remarking
    remark 2
exit
-----
*A:K-SASK12>config>qos>network#

```

DSCP Marking CPU Generated Traffic

DSCP marking for CPU generated traffic is not configurable by the user. The default values are given in Table 33:

Note: DSCP and Dot1P values in the table are applicable when remarking is disabled at port level.

Table 23: DSCP and Dot1p Marking

Protocol	IPv4	DSCP Marking	Dot1P Marking	Default FC	DSCP Values	DOT1P Values
SNMP	Yes	Yes	Yes	H2	34	4
NTP	Yes	Yes	Yes	H2	48	7
TELNET	Yes	Yes	Yes	H2	34	4
FTP	Yes	Yes	Yes	H2	34	4
TFTP	Yes	Yes	Yes	H2	34	4
SYSLOG	Yes	Yes	Yes	H2	34	4
TACACS	Yes	Yes	Yes	H2	34	4
RADIUS	Yes	Yes	Yes	H2	34	4
SSH	Yes	Yes	Yes	H2	34	7

Table 23: DSCP and Dot1p Marking

Protocol	IPv4	DSCP Marking	Dot1P Marking	Default FC	DSCP Values	DOT1P Values
ICMP Req	Yes	Yes	Yes	NC	0	7
ICMP Res	Yes	Yes	Yes	NC	0	7
ICMP Unreach	Yes	Yes	Yes	NC	0	7
SCP	NA	NA	Yes	H2	34	4
STP	NA	NA	Yes	NC	-	7
CFM	NA	NA	Yes	NC	-	7
ARP	NA	NA	Yes	NC	-	7
SNMP trap/log	Yes	Yes	Yes	H2	34	4
ICMP ping	Yes	Yes	Yes	NC	0	7
Trace route	Yes	Yes	Yes	NC	0	7
TACPLUS	Yes	Yes	Yes	H2	34	4
IGMP	Yes	Yes	Yes	NC	30	0

Default DSCP Mapping

Default DSCP Mapping Table

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary	Label
Default	0	0x00	0b000000	be
nc1	48	0x30	0b110000	h1
nc2	56	0x38	0b111000	nc
ef	46	0x2e	0b101110	ef
af11	10	0x0a	0b001010	assured
af12	12	0x0c	0b001100	assured
af13	14	0x0e	0b001110	assured
af21	18	0x12	0b010010	11

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary	Label
af22	20	0x14	0b010100	11
af23	22	0x16	0b010110	11
af31	26	0x1a	0b011010	11
af32	28	0x1c	0b011100	11
af33	30	0x1d	0b011110	11
af41	34	0x22	0b0100010	h2
af42	36	0x24	0b100100	h2
af43	38	0x26	0b100110	h2
default*	0			

*The default forwarding class mapping is used for all DSCP names/values for which there is no explicit forwarding class mapping.

Service Management Tasks

Deleting QoS Policies

A network policy is associated by default with access uplink ports.

You can replace the default policy with a non-default policy, but you cannot remove default policies from the configuration. When you remove a non-default policy, the policy association reverts to the appropriate default network policy.

Remove a Policy from the QoS Configuration

To delete a network policy, enter the following commands:

CLI Syntax: `config>qos# no network network-policy-id`

Copying and Overwriting Network Policies

You can copy an existing network policy to a new policy ID value or overwrite an existing policy ID. The overwrite option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos# copy network source-policy-id dest-policy-id [overwrite]`

The following output displays the copied policies:

```
A:ALA-12>config>qos# info detail
-----
...
network 1 create
  description "Default network QoS policy."
  scope template
  ingress
  default-action fc be profile out
...
network 600 create
  description "Default network QoS policy."
  scope template
  ingress
  default-action fc be profile out
...
network 700 create
  description "Default network QoS policy."
  scope template
  ingress
  default-action fc be profile out
...
-----
A:ALA-12>config>qos#
```

Editing QoS Policies

You can change existing policies, except the default policies, and entries in the CLI. The changes are applied immediately to all access uplink ports where the policy is applied. To prevent configuration errors use the copy command to make a duplicate of the original policy to a work area, make the edits, and then overwrite the original policy.

Network QoS Policy Command Reference

- [Configuration Commands for 7210 SAS-K 2F2T1C](#)
- [Configuration Commands for 7210 SAS-K 2F4T6C](#)
- [Operational Commands](#)
- [Show Commands](#)

Command Hierarchies

Configuration Commands for 7210 SAS-K 2F2T1C

```

— config
  — qos
    — [no] network network-policy-id [create]
      — description description-string
      — no description
      — scope {exclusive | template}
      — no scope
      — egress
        — no remark
        — remark policy-id
        — no remarking
      — marking
      — ingress
        — default-action fc fc-name profile {in | out | use-dei}
        — dot1p-classification policy-id
        — no dot1p-classification
        — dscp-classification policy-id
        — no dscp-classification
        — [no] fc fc-name [create]
          — multicast-queue queue-id
          — no multicast-queue
          — no queue
          — queue queue-id
          — no use-dei
          — use-dei
        — queue queue-id
        — no queue
          — [no] adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
          — cbs size-in-kbyte
          — no cbs
          — mbs size-in-kbytes
          — no mbs
          — no priority

```

- **priority** *level*
- **no rate**
- **rate** [**cir** *cir-percent*] [**pir** *pir-percent*]
- **no slope-policy**
- **slope-policy** *name*
- **no weight**
- **weight** *weight*

Configuration Commands for 7210 SAS-K 2F4T6C

- **config**
 - **qos**
 - [**no**] **network** *network-policy-id* [*create*]
 - **egress**
 - **no remark**
 - **remark** *policy-id*
 - **no remarking**
 - **marking**
 - **ingress**
 - **default-action** **fc** *fc-name* **profile** {**in** | **out** | **use-dei**}
 - **dot1p-classification** *policy-id*
 - **no dot1p-classification**
 - **dscp-classification** *policy-id*
 - **no dscp-classification**
 - [**no**] **fc** *fc-name* [*create*]
 - **multicast-queue** *queue-id*
 - **no multicast-queue**
 - **no queue**
 - **queue** *queue-id*
 - **no use-dei**
 - **use-dei**
 - **mpls-lsp-exp-classification** *policy-id*
 - **no mpls-lsp-exp-classification**
 - **queue** *queue-id*
 - [**no**] **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
 - **cbs** *size-in-kbyte*
 - **no cbs**
 - **mbs** *size-in-kbytes*
 - **no mbs**
 - **no priority**
 - **priority** *level*
 - **no rate**
 - **rate** [**cir** *cir-percent*] [**pir** *pir-percent*]
 - **no slope-policy**
 - **slope-policy** *name*
 - **no weight**
 - **weight** *weight*
 - **scope** {**exclusive** | **template**}

- **no scope**

Operational Commands

- **config**
 - **qos**
 - **copy network** *src-pol dst-pol* [**overwrite**]

Show Commands

- **show**
 - **qos**
 - **network** *policy-id* [**detail**]
 - **network** [*network-policy-id*] **association**
 - **network** [*network-policy-id*] [**detail**]

Network QoS Policy Command Descriptions

Configuration Commands

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>network policy-id
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the context in the configuration file. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax	copy network <i>src-pol dst-pol</i> [overwrite]
Context	config>qos
Description	This command copies existing QoS policy entries for a QoS policy-id to another QoS policy-id. The copy command is used to create new policies using existing policies and also allows bulk modifications to an existing policy with the use of the overwrite keyword. Platforms Supported - 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Network QoS Policy Command Descriptions

Parameters **network** *src-pol dst-pol* — Indicates that the source and destination policies are network policy IDs. Specify the source policy that the copy command will copy and specify the destination policy to which the command will duplicate the policy to a new or different policy ID.

Values

Values 1 — 65535

overwrite — Specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, an error will occur if the destination policy ID exists.

```
SR>config>qos# copy network 1 427
```

```
MINOR: CLI Destination "427" exists use {overwrite}.
```

```
SR>config>qos# copy network 1 427 overwrite
```

remark

Syntax **remark** *policy-id*
no remark

Context config>qos>network>egress

Description **Platforms Supported:** 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C

This command specifies the remarking policy ID to use for marking packets on access-uplink port egress or network port egress. The usage is as below for different platforms:

- On 7210 SAS-K2F2T1C this policy is used to configure marking for packets sent out of access -uplink port.
- On 7210 SAS-K2F4T6C this policy is used to configure marking for packets sent out of access-uplink port or network port, depending on the port which the policy is associated with.

For access-uplink port the behavior is as follows:

The remarking policy ID must be associated with the appropriate network QoS policy associated with the access-uplink port and remarking must be enabled in the network QoS policy to enable marking of packets sent out of access-uplink port egress. Only remarking policy of type dot1p, dscp, or dot1p-dscp is allowed to be used when the remark policy is associated with access-uplink port egress. See the *Remark policy* chapter for more information about remark policy types and its usage.

For network port the behavior is as follows:

The remarking policy ID must be associated with the appropriate network QoS policy associated with the network port and remarking must be enabled in the network QoS policy to enable marking of packets sent out of network port egress. On network ports, the dot1p bits are marked by default irrespective of the remarking enabled or disabled. Only remarking policy of type lsp-exp, dot1p, dscp, or dot1p-dscp, dot1p-lsp-exp, dot1p-dscp-lsp-exp is allowed to be used when the remark policy is associated with network port egress. See the Remark policy chapter for more information about remark policy types and its usage.

The `no` form of the command removes the explicit association of the remark policy and associates the default remark policy. In other words, if remarking is enabled and no remark policy is executed, then the default remark policy is used to mark packets sent out. If no remark policy is executed and remarking is disabled, then packets are not remarked at all.

Parameters *policy-id* — The parameter that uniquely identifies the remark policy.

Values

Values 1 — 65535

remarking

Syntax **remarking**
no remarking

Context config>qos>network>egress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

On 7210 SAS-K 2F2T1C this command enables marking for packets sent out of access-uplink ports.

On 7210 SAS-K 2F4T6C this command enables marking for packets sent out of access-uplink or network port, depending on the port which the policy is associated with.

When remarking is enabled, the remark policy configured in the QoS policy context is used to determine the FC to QoS bit mapping. For example, when remarking is enabled in the network QoS policy and remark policy of type `dot1p` is configured in the network QoS policy, then the FC to Dot1p mapping is used to mark packets sent out of the port.

Please refer to the remark policy that can be used to configure FC to priority bit markings in different QoS policies associated with different service entities, in the chapter on remark policies. For more information, see the chapter on remark policies.

The `no` form of the command disables remarking.

Default no remarking - Remarking is disabled by default.

scope

Syntax **scope {exclusive | template}**
no scope

Context config>qos>network policy-id

Description This command configures the network policy scope as exclusive or template.

The `no` form of this command sets the scope of the policy to the default of **template**.

Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Network QoS Policy Command Descriptions

Default	template
Parameters	<p>exclusive — When the scope of a policy is defined as exclusive, the policy can only be applied to one interface. If a policy with an exclusive scope is assigned to a second interface an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface.</p> <p>The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to 1.</p> <p>template — When the scope of a policy is defined as template, the policy can be applied to multiple interfaces on the router.</p> <p>Default QoS policies are configured with template scopes. An error is generated if you try to modify the template scope parameter to exclusive scope on default policies.</p>

Network QoS Policy Commands

network

Syntax	network <i>network-policy-id</i> [create] no network <i>network-policy-id</i>
Context	config>qos
Description	<p>This command creates or edits a QoS network policy. The network policy functionality is as follows:</p> <ul style="list-style-type: none"> • On 7210 SAS-K 2F2T1C, the network policy defines the treatment the packets receive on ingress and egress the access-uplink port. • On 7210 SAS-K 2F4T6C, the network policy defines the treatment the packets receive on ingress and egress the access-uplink port and network port. <p>On 7210 SAS-K 2F2T1C, network qos policy associated with access-uplink port has a ingress and egress component.</p> <ul style="list-style-type: none"> • The ingress component of the policy defines how Dot1p bits or IP DSCP are mapped to internal forwarding class and profile state. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the 7210 SAS. The mapping on each network interface defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the network interface. It also defines the rate-limiting parameters for the traffic mapped to each forwarding classes. On 7210 SAS-K 2F2T1C, traffic mapped to each forwarding class can be rate limited using same or separate queue for unicast and multipoint traffic. • The egress component of the network qos policy defines the forwarding class and profile to packet header priority bits (for example: Dot1p bits). Option is provided to map forwarding class to Dot1p bits and IP DSCP bits. <p>On 7210 SAS-K 2F4T6C, the network qos policy associated with access-uplink port has a ingress and egress component:</p> <ul style="list-style-type: none"> • The ingress component of the policy defines how Dot1p bits or IP DSCP are mapped to internal forwarding class and profile state. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the 7210 SAS. The mapping on each network interface defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the network interface. It also defines the rate-limiting parameters for the traffic mapped to each forwarding classes. On 7210 SAS-K 2F4T6C, traffic mapped to each forwarding class can be rate limited using same or separate queue for unicast and multipoint traffic. • The egress component of the network qos policy defines the forwarding class and profile to packet header priority bits (for example: Dot1p bits). Option is provided to map forwarding class to Dot1p bits and IP DSCP bits. <p>On 7210 SAS-K 2F4T6C the network qos policy associated with network port has a ingress and egress component:</p>

Network QoS Policy Command Descriptions

- The ingress component of the policy defines how MPLS EXP bits or Dot1p bits or IP DSCP are mapped to internal forwarding class and profile state. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the 7210 SAS. The mapping on each network interface defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the network interface. It also defines the rate-limiting parameters for the traffic mapped to each forwarding classes. On 7210 SAS-K 2F4T6C, traffic mapped to each forwarding class can be rate limited using same or separate queue for unicast and multipoint traffic.
- The egress component of the network qos policy defines the forwarding class and profile to packet header priority bits (for example: Dot1p bits). Option is provided to map forwarding class to MPLS EXP bits, Dot1p bits and IP DSCP bits.

The default network policy-id 1 is associated with access-uplink ports that do not have an explicit user configured policy and cannot be modified or deleted. The default network policy-id 2 is associated with network ports that do not have an explicit user configured policy and cannot be modified or deleted. The default network policies define default mapping for packet header bits to forwarding class on ingress and the mapping of FC to queues.

If a new network policy is created (for instance, policy-id 10), only the default action, default queues for unicast and multipoint traffic and egress forwarding class parameters are identical to the default policy. A new network policy does not contain the default QoS bit to FC mapping (e.g. Dot1p-to-FC mapping or EXP to FC mapping) for network QoS policy. The default network policy can be copied (use the copy command) to create a new network policy that includes the default ingress Dot1p to FC mapping (as appropriate). You can modify parameters or use the no modifier to remove an object from the configuration.

Any changes made to an existing policy, using any of the sub-commands, will be applied immediately to all the ports where this policy is applied. For this reason, when many changes are required on a policy, it is highly recommended that the policy be copied to a work area policy-id. That work-in-progress policy can be modified until complete and then written over the original policy-id. Use the config qos copy command to maintain policies in this manner.

The **no** form of this command deletes the network policy. A policy cannot be deleted until it is removed from all entities where it is applied. The default network **policy policy-id 1** cannot be deleted.

Default System Default Network Policy 1

Parameters *network-policy-id* — The policy-id uniquely identifies the policy on the 7210 SAS.

Default none

Values 1— 65535

Network Ingress QoS Policy Commands

fc

Syntax	[no] fc <i>fc-name</i> [<i>create</i>]
Context	config>qos>network>ingress
Description	<p>This command creates a class instance of the forwarding class. Once the <i>fc-name</i> is created, classification actions can be applied and it can be used in match classification criteria.</p> <p>The no form of the command removes all the explicit queue mappings for <i>fc-name</i> forwarding types. The queue mappings revert to the default meters for <i>fc-name</i>.</p> <p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C</p>
Default	Undefined forwarding classes default to the configured parameters in the default policy <i>policy-id</i> 1.
Parameters	<p><i>fc-name</i> — The case-sensitive, system-defined forwarding class name for which policy entries will be created.</p> <p>Values be, l2, af, l1, h2, ef, h1, nc</p> <p>create — The keyword used to create the forwarding class. The create keyword requirement can be enabled/disabled in the environment> create context.</p>

ingress

Syntax	ingress
Context	config>qos>network policy-id
Description	<p>This command is used to enter the CLI node that creates or edits policy entries that specify the QoS bits to forwarding class mapping for all packets.</p> <p>When pre-marked packets ingress on a network port, the QoS treatment through the 7210 SAS-based on the mapping defined under the current node.</p>

default-action

Syntax	default-action fc <i>fc-name</i> [profile {in out} add use-dei]
Context	config>qos>network>ingress
Description	This command defines or edits the default action to be taken for packets do not match any f the configured classification entries. The default-action command specifies the forwarding class to which such packets are assigned.

Network QoS Policy Command Descriptions

Multiple default-action commands will overwrite each previous default-action command.

Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Default default-action fc be profile out

Parameters **fc *fc-name*** — Specify the forwarding class name. All packets with Dot1p or dot1p bits that is not defined will be placed in this forwarding class.

Default None, the fc name must be specified

Values be, l2, af, l1, h2, ef, h1, nc

profile {in | out} — All packets assigned to this forwarding class will be considered in or out of profile based on this command. A value of 'in' defines the packet profile as 'in-profile' and a value of 'out' defines the packet profile to be out-of-profile.

use-dei — use-dei is specified with the FC, the DEI bit received in the ethernet VLAN tag is used assign the profile to the packet. The packet will be considered to be in-profile if the DEI bit value is 0, out-profile otherwise (i.e. DEI bit of '1').

It is mutually exclusive to configuration of explicit profile values using the profile parameter.

use-dei

Syntax [no] use-dei

Context config>qos>network>ingress>fc

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

This command is used to enable DEI based classification. When enabled for packets classified to this FC, the DEI bit is used to determine the ingress profile for the packet. Packets received with DEI bit set to zero are treated as in-profile and packets with DEI bit set to one are treated as out-of-profile packets.

When DEI based classification is enabled under the FC context, it overrules the profile values specified in the classification entry used to assign the FC.

The no form of the command disables use of DEI bit for classification of packets.

Default no use-dei

dot1p-classification

Syntax dot1p-classification *policy*
no dot1p-classification

Context config>qos>network>ingress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

This command allows the user to associate a dot1p classification policy which contain entries to use to map traffic received on network port or access-uplink port to a forwarding class and profile state based on the dot1p bits in the packet.

The no form of the policy disables the use of this policy.

Default	no dot1p-classification-policy
Parameters	<i>policy id</i> — The policy-id uniquely identifies the policy on the 7210 SAS.
Values	[1..65535]

dscp-classification

Syntax	dscp-classification <i>policy</i> no dscp-classification
Context	config>qos>network>ingress
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C This command allows the user to associate a IP DSCP classification policy which contain entries to use to map traffic received on network port or access-uplink port to a forwarding class and profile state based on the IP DSCP bits in the packet. The no form of the policy disables the use of this policy.
Default	no dot1p-classification-policy
Parameters	<i>policy id</i> — The policy-id uniquely identifies the policy on the 7210 SAS.
Values	[1..65535]

mpls-lsp-exp-classification

Syntax	mpls-lsp-exp-classification <i>policy-id</i> mpls-lsp-exp-classification
Context	config>qos>network>ingress
Description	Platforms Supported: 7210 SAS-K 2F4T6C This command allows the user to associate a mpls-lsp-exp classification policy which contain entries to use to map traffic received on network port to a forwarding class and profile state based on the MPLS EXP bits in the MPLS packet.
Default	no mpls-lsp-exp-classification

Network QoS Policy Command Descriptions

Parameters *policy id* — The policy-id uniquely identifies the policy on the 7210 SAS.

Values [1..65535]

egress

Syntax **egress**

Context config>qos>network policy-id

Description This command is used to enter the CLI node that creates or edits egress policy entries that specify the forwarding class to marking values map to be instantiated when this policy is applied to the access-uplink port.

The forwarding class and profile state mapping to marking values (for example: IEEE 802.1p bits, etc.) bits mapping for all packets are defined in this context.

fc

Syntax [**no**] **fc** *fc-name*

Context config>qos>network>egress

Description This command specifies the forwarding class name. The forwarding class name represents an egress queue. The **fc** *fc-name* represents a CLI parent node that contains sub-commands or parameters describing the marking criteria of packets flowing through it. The **fc** command overrides the default parameters for that forwarding class to the values defined in the network default policy.

The no form of this command removes the forwarding class to marking value association. The forwarding class reverts to the mapping defined in the default network policy.

Default Undefined forwarding classes default to the configured parameters in the default network policy policy-id 1.

Parameters *fc-name* — The case-sensitive, system-defined forwarding class name for which policy entries will be created.

Default none

Values be, l2, af, l1, h2, ef, h1, nc

Network Ingress Queue QoS Policy Commands

adaptation-rule

Syntax	adaptation-rule [<i>cir adaptation-rule</i>] [<i>pir adaptation-rule</i>] no adaptation-rule
Context	config>qos>network>ingress>queue
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C</p> <p>This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.</p> <p>The no form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific adaptation-rule is removed, the default constraints for cir and pir apply.</p>
Default	adaptation-rule pir closest cir closest
Parameters	<p><i>adaptation-rule</i> — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.</p> <p>Values</p> <p>pir — Defines the constraints enforced when adapting the PIR rate defined within the queue queue-id rate command. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the rate command is not specified, the default applies.</p> <p>cir — Defines the constraints enforced when adapting the CIR rate defined within the queue queue-id rate command. The cir parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the cir parameter is not specified, the default constraint applies.</p> <p>max — The max (maximum) option is mutually exclusive with the min and closest options. The hardware step size varies with the configured rate.</p> <p>min — The min (minimum) option is mutually exclusive with the max and closest options. The hardware step size varies with the configured rate.</p> <p>closest — The closest parameter is mutually exclusive with the min and max parameter. The hardware step size varies with the configured rate.</p>

Network QoS Policy Command Descriptions

cbs

Syntax	[no] cbs <i>size-in-kbytes</i>
Context	config>qos>network>ingress>queue
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C This command specifies the CBS value (Minimum depth of the queue in kilo bytes). The no form of the command sets this to the default value.
Default	32kbytes
Parameters	<i>size-in-kbytes</i> — Specifies the minimum depth of the queue in kilo bytes. Values [0..10240]

mbs

Syntax	[no] mbs <i>size-in-kbytes</i>
Context	config>qos>network>ingress>queue
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C This command specifies the MBS value (Maximum depth of the queue in kilo bytes). The no form of the command sets this to the default value.
Default	512kbytes
Parameters	<i>size-in-kbytes</i> — Specifies the minimum depth of the queue in kilo bytes. Values [0..12800]

priority

Syntax	[no] priority <i>level</i>
Context	config>qos>network>ingress>queue
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C Defines the queue priority. The queue priority is used by the scheduler to determine the order of service in both the within-cir loop and within-pir loop. Higher priority queues are serviced before lower priority queues. The no form of the command sets the default priority values.
Default	1

Parameters *level* — Specifies the priority of the queue.

Values [1..4]

slope-policy

Syntax **[no] slope-policy** *name*

Context config>qos>network>ingress>queue

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

The slope-policy command is used to override the default slope-policy configuration for the queue. The specified slope-policy-name must exist as a current slope policy name. If the slope policy does not exist, the slope-policy command will fail. If a slope policy is currently associated with a queue, the slope policy cannot be removed from the system.

The slope policy contains the ring and non-ring High and Low WRED slope definitions that will be used by the queue. The non-ring slopes are used by the traffic received on access SAP ingress and sent out of access SAP egress queues. The ring slopes are used by the traffic received on access-uplink port ingress and sent out of access-uplink port egress queues.

For more information about ring and non-ring buffer pools and slope usage see the above section on buffer pools.

If the slope-policy command is not executed or the no slope policy command is executed, the default slope policy will be associated with the queue.

The no form of the command restores the default slope policy to the queue.

Parameters *name* — The name parameter is required and must specify an existing slope policy name. If slope-policy-name does not exist, the slope-policy command will fail.

Values 32 chars max

queue

Syntax **queue** *queue-id* **create**
no queue

Context config>qos>network>ingress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

This command creates the context to modify Queue parameters associated with a particular queue.

The queue ID to FC map is user defined. In other words, user can map FC to queues identified by queue-IDs as per their needs.

The no form of this command deletes the queue.

Network QoS Policy Command Descriptions

Default none

Parameters *queue-id* — Specifies the id of the queue.

Values 1 — 8

create — Mandatory keyword to create a network queue policy.

weight

Syntax **[no] weight** *weight*

Context config>qos>network>ingress>queue

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Allows users to specify the weight of the queue.

The configured weight determines the proportion of available bandwidth that is given to this queue in comparison to other queues contending for bandwidth at the same priority level.

The no form of the command sets the weight to default of one.

Default 1

Parameters *weight* — Specifies the value for the weight parameter. It is an integer value which specifies the proportion of available bandwidth to be allocated to this queue relative to other queues.

Values 1 — 100

rate

Syntax **rate** [*cir* *cir-percent*] [*pir* *pir-percent*]
no rate

Context config>qos>network>ingress>queue

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the port. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by over subscription factors or available egress bandwidth. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.

The rate command can be executed at anytime, altering the PIR and CIR rates for all queues created on the access ports.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (100, 0).

- Parameters** *cir percent* — Defines the percentage of the guaranteed rate allowed for the queue. When the **rate** command is executed, a valid CIR setting must be explicitly defined. When the **rate** command has not been executed, the default **CIR of 0** is assumed. Fractional values are not allowed and must be given as a positive integer.
- The actual CIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.
- Values** 0 — 100
- Default** 0
- pir percent* — Defines the percentage of the maximum rate allowed for the queue. When the **rate** command is executed, the PIR setting is optional. When the **rate** command has not been executed, or the PIR parameter is not explicitly specified, the default PIR of 100 is assumed. Fractional values are not allowed and must be given as a positive integer.
- Values** 1— 100 percent
- Default** 100

Show Commands

network

- Syntax** **network** [*policy-id*] [*detail*]
- Context** show>qos
- Description** **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C
- This command displays network policy information.
- Parameters** *policy-id* — Displays information for the specific policy ID.
- Default** all network policies
- Values** 1 — 65535
- detail** — Includes information about ingress and egress Dot Ip bit mappings and network policy interface associations.
- Output** **Network QoS Policy Output Fields**
- The following table describes network QoS Policy output fields.

Table 24: Show QoS Network Output Fields

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Description	A text string that helps identify the policy’s context in the configuration file.
Forward Class/ FC Name	Specifies the forwarding class name.
Profile	Out Specifies the EXP marking for the packets which are out-of-profile, egressing on this queue.Specifies the Dot1p marking for the packets which are out-of-profile, egressing on this queue.
	In Specifies the EXP marking for the packets which are in-of-profile, egressing on this queue.Specifies the Dot1p markings for in-profile packets egressing this queue.
Port-Id	Specifies the physical port identifier that associates the interface.

Sample output for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C:

```
*A:7210SAS>show>qos# network 1 detail

=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id      : 1
Egr Remark    : False
Forward Class  : be                               Profile      : None
Scope         : Template
DOT1P Class Po*: 1                               DSCP Class P*: 0
Description    : Default network-port QoS policy.

-----
FC          Queue      MCast Queue Use Dei
-----
be         1          1          false
l2         2          2          false
af         3          3          false
l1         4          4          false
h2         5          5          false
ef         6          6          false
h1         7          7          false
nc         8          8          false

-----
Queue Rates and Rules
-----
```

QueueId	CIR	CIR Adpt Rule	PIR	PIR Adpt Rule
Queue1	0	closest	100	closest
Queue2	25	closest	100	closest
Queue3	25	closest	100	closest
Queue4	25	closest	100	closest
Queue5	100	closest	100	closest
Queue6	100	closest	100	closest
Queue7	10	closest	100	closest
Queue8	10	closest	100	closest

Queue Priority and Weight Details

QueueId	Priority	Weight
Queue1	1	1
Queue2	1	1
Queue3	1	1
Queue4	1	1
Queue5	1	1
Queue6	1	1
Queue7	1	1
Queue8	1	1

High Slope Non Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	80
Queue2	Down	70	90	80
Queue3	Down	70	90	80
Queue4	Down	70	90	80
Queue5	Down	70	90	80
Queue6	Down	70	90	80
Queue7	Down	70	90	80
Queue8	Down	70	90	80

Low Slope Non Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	80
Queue2	Down	50	75	80
Queue3	Down	50	75	80
Queue4	Down	50	75	80
Queue5	Down	50	75	80
Queue6	Down	50	75	80
Queue7	Down	50	75	80
Queue8	Down	50	75	80

High Slope Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
---------	-------	--------------	------------	-------------

Network QoS Policy Command Descriptions

```

Queue1      Down      70      90      80
Queue2      Down      70      90      80
Queue3      Down      70      90      80
Queue4      Down      70      90      80
Queue5      Down      70      90      80
Queue6      Down      70      90      80
Queue7      Down      70      90      80
Queue8      Down      70      90      80
-----
Low Slope Ring
-----
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1      Down      50      75      80
Queue2      Down      50      75      80
Queue3      Down      50      75      80
Queue4      Down      50      75      80
Queue5      Down      50      75      80
Queue6      Down      50      75      80
Queue7      Down      50      75      80
Queue8      Down      50      75      80
-----
Slope Policies
-----
-----
QueueId      CBS(KBytes)  MBS(KBytes)  Slope-Policy
-----
Queue1      50          200         default
Queue2      50          200         default
Queue3      50          200         default
Queue4      50          200         default
Queue5      50          200         default
Queue6      50          200         default
Queue7      50          200         default
Queue8      50          200         default
-----
Port Attachments
-----
No Matching Entries

=====
* indicates that the corresponding row element may have been truncated.
*A:7210SAS>show>qos#

```

Network QoS Policy Output Fields

network

Syntax **network** [*network-policy-id*] **association**
network [*network-policy-id*] [**detail**]

Context show>qos
Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

This command displays network policy information.

Parameters *network-policy-id* — Displays network information for the specific policy ID.

Default all network policies

Values 1 — 65535

detail — Includes information about ingress and egress Dot1pEXP bit mappings and network policy interface associations.

association —

Output Network QoS Policy Output Fields

The following table describes network QoS Policy output fields.

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Description	A text string that helps identify the policy’s context in the configuration file.
Forward Class/ FC Name	Specifies the forwarding class name.
Profile	Out Indicates that packets are classified as out-profile.
	In Indicates packets are classified as in-profile.
	None Indicates packets profile is undefined.
DOTIP Class	Specifies the dot1p-classification policy ID that is being used for mapping the packets to different FC under the FCs based on the dot1p bits.
DSCP classification	Specifies the dscp-classification policy ID that is being used for mapping the packets to different FC under the FCs based on the dscp bits.
High Slope Non Ring	Specifies the non-ring high-slope policy values
Low Slope Non Ring	Specifies the non-ring low-slope values.
High Slope Ring	Specifies the ring high-slope values.
Slope Policies	Displays the slope policies applied to the queues.

Sample for 7210 SAS-K 2F2T1C

```
*A:Dut-A>show>qos# network 2 detail
```

```
=====
```

```
=====
```

```
QoS Network Policy
```

Network QoS Policy Command Descriptions

```

=====
-----
Network Policy (2)
-----
Policy-id          : 2
Egr Remark        : False          Egr Rem Plcy      : N/A
Forward Class     : be             Profile           : None
Scope             : Template
DOT1P Class Poli* : 1              DSCP Class Polic* : 1
MPLS Lsp Exp Cla* : 1
Description       : Default network QoS policy.
-----

FC          Queue      MCast Queue Use Dei
-----
be          1          1          false
l2          2          2          false
af          3          3          false
l1          4          4          false
h2          5          5          false
ef          6          6          false
h1          7          7          false
nc          8          8          false
-----

Queue Rates and Rules
-----

QueueId      CIR          CIR Adpt Rule      PIR          PIR Adpt Rule
-----
Queue1       0          closest           100         closest
Queue2       25         closest           100         closest
Queue3       25         closest           100         closest
Queue4       25         closest           100         closest
Queue5       100        closest           100         closest
Queue6       100        closest           100         closest
Queue7       10         closest           100         closest
Queue8       10         closest           100         closest
-----

Queue Priority and Weight Details
-----

QueueId      Priority      Weight
-----
Queue1       1            1
Queue2       1            1
Queue3       1            1
Queue4       1            1
Queue5       1            1
Queue6       1            1
Queue7       1            1
Queue8       1            1
-----

High Slope Non Ring
-----

QueueId      State      Start-Avg(%)      Max-Avg(%)      Max-Prob(%)
-----
Queue1       Down      70                 90                80
Queue2       Down      70                 90                80
-----

```

Queue3	Down	70	90	80
Queue4	Down	70	90	80
Queue5	Down	70	90	80
Queue6	Down	70	90	80
Queue7	Down	70	90	80
Queue8	Down	70	90	80

Low Slope Non Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	80
Queue2	Down	50	75	80
Queue3	Down	50	75	80
Queue4	Down	50	75	80
Queue5	Down	50	75	80
Queue6	Down	50	75	80
Queue7	Down	50	75	80
Queue8	Down	50	75	80

High Slope Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	80
Queue2	Down	70	90	80
Queue3	Down	70	90	80
Queue4	Down	70	90	80
Queue5	Down	70	90	80
Queue6	Down	70	90	80
Queue7	Down	70	90	80
Queue8	Down	70	90	80

Low Slope Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	80
Queue2	Down	50	75	80
Queue3	Down	50	75	80
Queue4	Down	50	75	80
Queue5	Down	50	75	80
Queue6	Down	50	75	80
Queue7	Down	50	75	80
Queue8	Down	50	75	80

Slope Policies

QueueId	CBS(KBytes)	MBS(KBytes)	Slope-Policy
Queue1	24	200	default
Queue2	24	200	default
Queue3	24	200	default
Queue4	24	200	default
Queue5	24	200	default

Network QoS Policy Command Descriptions

Queue6	24	200	default
Queue7	24	200	default
Queue8	24	200	default

----- Port Attachments -----

Port-id : 1/1/1
Port-id : 1/1/2
Port-id : 1/1/3
Port-id : 1/1/4
Port-id : 1/1/5

=====

* indicates that the corresponding row element may have been truncated.

*A:Dut-A>show>qos#

*A:SAH01-051>show>qos# network 1 association

----- QoS Network Policy -----

----- Network Policy (1) -----

Policy-id : 1
Egr Remark : False
Forward Class : be Profile : None
Scope : Template
DOT1P Class Po*: 1 DSCP Class P*: 0
Description : Default network-port QoS policy.

----- Port Attachments -----

No Matching Entries

=====

* indicates that the corresponding row element may have been truncated.

*A:SAH01-051>show>qos#

Network Queue QoS Policies

In This Section

This section provides information to configure network queue QoS policies using the command line interface.

Topics in this section include:

- [Overview](#)
- [Basic Configurations](#)
- [Default Network Queue Policy Values](#)
- [Service Management Tasks](#)

Overview

On 7210 SAS-K 2F2T1C, Network Queue policies define the egress network queuing for the traffic egressing on the access uplink port. Network queue policies are used at the Ethernet port and define the bandwidth distribution for the various FC traffic egressing on the Ethernet port. On 7210 SAS-K 2F2T1C, user can define the number of queues and the mapping of FC to queue per network-queue policy. Each of these queues are shared by unicast and multicast traffic.

On 7210 SAS-K 2F4T6C, Network Queue policies define the egress network queuing for the traffic egressing on the access uplink port and the network port. Network queue policies are used at the Ethernet port and define the bandwidth distribution for the various FC traffic egressing on the Ethernet port. On 7210 SAS-K 2F4T6C, user can define the number of queues and the mapping of FC to queue per network-queue policy. Each of these queues are shared by unicast and multicast traffic.

Basic Configurations

A basic network queue QoS policy must conform to the following:

- Each network queue QoS policy must have a unique policy name.
- Queue parameters can be modified, but cannot be deleted.

Create a Network Queue QoS Policy

Configuring and applying QoS policies other than the default policy is optional. A default network queue policy is applied to all access uplink ports:

To create a network queue policy, define the following:

- Enter a network queue policy name. The system will not dynamically assign a name.
- Include a description. The description provides a brief overview of policy features.
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the FC to queue ID mapping can be defined by the user in the policy. User has an option use to lesser number of queues.

Use the following CLI syntax to create a network queue QoS policy

```
CLI Syntax:  config>qos
                network-queue policy-name
                description description-string
                queue queue-id
                  rate cir cir-percent [pir pir-percent]
                adaptation-rule [cir adaptation-rule] [pir
                  adaptation-rule]
```

```
*A:Dut-B>config>qos>network-queue# info detail
-----
description "Default network queue QoS policy."
queue 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
exit
queue 2
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
exit
queue 3
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
exit
queue 4
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
exit
queue 5
  rate cir 100 pir 100
  adaptation-rule cir closest pir closest
```

```

exit
queue 6
    rate cir 100 pir 100
    adaptation-rule cir closest pir closest
exit
queue 7
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
exit
queue 8
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
exit
-----
*A:Dut-B>config>qos>network-queue#

```

Applying Network Queue Policies

Apply network queue policies to the following entities:

- [Applying Network Queue Configuration to access-uplink port](#)

Applying Network Queue Configuration to access-uplink port

Use the following CLI syntax to apply a network queue policy to an Ethernet port.

CLI Syntax:

```

config>port#
    ethernet
        network
            queue-policy policy-name

```

```

#-----
echo "Port Configuration"
#-----
    port 1/1/1
        ethernet
            access
                uplink
                    queue-policy "nq1-cbs"
            exit
    exit
no shutdown
    exit
#-----

```

Default Network Queue Policy Values

Applying Network Queue Configuration to network port on 7210 SAS-K 2F4T6C

Use the following CLI syntax to apply a network queue policy to an Ethernet network port.

```
CLI Syntax:  config>port#
             ethernet
             network
             queue-policy policy-name
```

```
#-----
echo "Port Configuration"
#-----
    port 1/1/1
      ethernet
      network
      queue-policy "nql-cbs"
    exit
exit
no shutdown
  exit
#-----
```

Default Network Queue Policy Values

The default network queue policies are identified as policy-id **default**. The default policies cannot be modified or deleted.

Default Network Queue policy for 7210 SAS-K 2F2T1C

The following displays default policy parameters for 7210 SAS-K 2F2T1C:

```
*A:dut-i>config>qos>network-queue# info detail
-----
description "Default network queue QoS policy."
queue 1 create
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
```

```
queue 2 create
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 3 create
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 4 create
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 5 create
  rate cir 100 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 6 create
  rate cir 100 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 7 create
  rate cir 10 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
  slope-policy "default"
  priority 1
  weight 1
exit
queue 8 create
  rate cir 10 pir 100
  adaptation-rule cir closest pir closest
  mbs 200
  cbs 50
```

Default Network Queue policy for 7210 SAS-K 2F4T6C

```
        slope-policy "default"
        priority 1
        weight 1
    exit
    fc af create
        queue 3
    exit
    fc be create
        queue 1
    exit
    fc ef create
        queue 6
    exit
    fc h1 create
        queue 7
    exit
    fc h2 create
        queue 5
    exit
    fc l1 create
        queue 4
    exit
    fc l2 create
        queue 2
    exit
    fc nc create
        queue 8
    exit
-----
*A:dut-i>config>qos>network-queue#
```

Default Network Queue policy for 7210 SAS-K 2F4T6C

The following displays default policy parameters for 7210 SAS-K 2F4T6C.

```
*A:K-SASK12>config>qos>network-queue# info detail
-----
description "Default network queue QoS policy."
queue 1 create
    rate cir 0 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    slope-policy "default"
    priority 1
    weight 1
exit
queue 2 create
    rate cir 25 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
```

```
        slope-policy "default"
        priority 1
        weight 1
    exit
queue 3 create
    rate cir 25 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    slope-policy "default"
    priority 1
    weight 1
exit
queue 4 create
    rate cir 25 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    slope-policy "default"
    priority 1
    weight 1
exit
queue 5 create
    rate cir 100 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    slope-policy "default"
    priority 1
    weight 1
exit
queue 6 create
    rate cir 100 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    slope-policy "default"
    priority 1
    weight 1
exit
queue 7 create
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    slope-policy "default"
    priority 1
    weight 1
exit
queue 8 create
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    mbs 200
    cbs 24
    slope-policy "default"
    priority 1
    weight 1
exit
fc af create
```

Service Management Tasks

```
        queue 3
    exit
fc be create
    queue 1
    exit
fc ef create
    queue 6
    exit
fc h1 create
    queue 7
    exit
fc h2 create
    queue 5
    exit
fc l1 create
    queue 4
    exit
fc l2 create
    queue 2
    exit
fc nc create
    queue 8
    exit
```

*A:K-SASK12>config>qos>network-queue#

Service Management Tasks

This section discusses the following service management tasks:

- [Deleting Network Queue QoS Policies](#)
- [Copying and Overwriting Network Queue QoS Policies](#)
- [Editing Network Queue QoS Policies](#)

Deleting Network Queue QoS Policies

A network queue policy is associated by default with all access uplink ports. You can replace the default policy with a customer-configured policy, but you cannot entirely remove a QoS policy. When you remove a QoS policy, the policy association reverts to the default network-queue policy **default**.

To delete a user-created network queue policy, enter the following commands:

CLI Syntax: config>qos# no network-queue policy-name

Example: config>qos# no network-queue nq1

Copying and Overwriting Network Queue QoS Policies

You can copy an existing network queue policy, rename it with a new policy ID name, or overwrite an existing network queue policy. The overwrite option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos# copy network-queue source-policy-id dest-policy-id [overwrite]`

Example: `config>qos# copy network-queue nq1-cbs nq2-cbs`

The following output displays the copied policies

```
*A:card-1>config>qos# info
#-----
echo "QoS Slope and Queue Policies Configuration"
#-----
.....
network-queue "nq1-cbs" create
  queue 1
    rate cir 0 pir 32
    adaptation-rule cir max
  exit
  queue 2
  exit
  queue 3
  exit
  queue 4
  exit
  queue 5
  exit
  queue 6
    rate cir 0 pir 4
  exit
  queue 7
    rate cir 3 pir 93
  exit
  queue 8
    rate cir 0 pir 3
  exit
exit
network-queue "nq2-cbs" create
  queue 1
    rate cir 0 pir 32
    adaptation-rule cir max
  exit
  queue 2
  exit
  queue 3
  exit
  queue 4
  exit
  queue 5
  exit
  queue 6
```

Service Management Tasks

```
        rate cir 0 pir 4
    exit
    queue 7
        rate cir 3 pir 93
    exit
    queue 8
        rate cir 0 pir 3
    exit
    exit
-----
*A:card-1>config>qos# info
```

Editing Network Queue QoS Policies

You can change existing policies, except the default policies, and entries in the CLI. The changes are applied immediately to all ports where the policy is applied. To prevent configuration errors use the copy command to make a duplicate of the original policy to a work area, make the edits, and then overwrite the original policy.

Network Queue QoS Policy Command Reference

Command Hierarchies

- [Configuration Commands for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C](#)
- [Operational Commands](#)
- [Show Commands](#)

Configuration Commands for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

```

— config
  — qos
    — network-queue policy-name [create]
      — description description-string
      — no description
      — [no] fc fc-name [create]
        — no queue
        — queue queue-id
      — queue queue-id
      — no queue
        — adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
        — no adaptation-rule
        — cbs size-in-kbyte
        — no cbs
        — mbs size in kbytes
        — no mbs
        — no priority
        — priority level
        — no rate
        — rate cir cir-percent [pir pir-percent]
        — no slope-policy
        — slope-policy name
        — no weight
        — weight weight

```

Operational Commands

```

— config
  — qos
    — copy network-queue src-name dst-name [overwrite]

```

Show Commands

- show
 - qos
 - **network-queue** [*network-queue-policy-name*] [**detail**]

Network Queue Command Descriptions

Configuration Commands

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>network-queue
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the context in the configuration file. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax	copy network-queue <i>src-name dst-name</i> [overwrite]
Context	config>qos
Description	This command copies or overwrites existing network queue QoS policies to another network queue policy ID. The copy command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the overwrite keyword.

Network Queue Command Descriptions

Parameters **network-queue** *src-name dst-name* — Indicates that the source policy ID and the destination policy ID are network-queue policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

overwrite — specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, a message is generated saying that the destination policy ID exists.

```
SR>config>qos# copy network-queue nq1 nq2  
MINOR: CLI Destination "nq2" exists - use {overwrite}.  
SR>config>qos# copy network-queue nq1 nq2 overwrite
```

Network Queue QoS Policy Commands

network-queue

Syntax **[no] network-queue** *policy-name [create]*

Context config>qos

Description **Platforms supported** - 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

This command creates a context to configure a network queue policy. Network queue policies on the Ethernet port define network egress queuing.

On 7210 SAS-K2F2T1C, network queue policy can be associated with access-uplink ports to define queues to be used on egress.

On 7210 SAS-K2F4T6C, network queue policy can be associated with access-uplink ports and network ports to define queues to be used on egress.

Default default

Parameters *policy-name* — The name of the network queue policy.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

create — Mandatory keyword to create a network queue policy.

Network Queue QoS Policy Queue Commands

queue

Syntax	queue <i>queue-id</i> no queue
Context	config>qos>network-queue
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C</p> <p>This command enables the context to configure a QoS network-queue policy queue.</p> <p>The user has an option to define the FC to queue map. Either one or multiple FCs can be mapped to the same queue.</p> <p>The no form of this command deletes the queue.</p>
Parameters	<p><i>queue-id</i> — The <i>queue-id</i> for the queue, expressed as an integer. The <i>queue-id</i> uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.</p> <p>Values 1 — 8</p>

adaptation-rule

Syntax	adaptation-rule [cir <i>adaptation-rule</i>] [pir <i>adaptation-rule</i>] no adaptation-rule
Context	config>qos>network-queue>queue
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.</p> <p>The no form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific adaptation-rule is removed, the default constraints for pir and cir apply.</p>
Default	adaptation-rule cir closest pir closest

Network Queue Command Descriptions

Parameters *adaptation-rule* — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.

Values **pir** — Defines the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **pir** command is not specified, the default applies.

cir — Defines the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

max — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

cbs

Syntax **[no]** **cbs** *size-in-kbytes*

Context config>qos>network-queue>queue

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command specifies the CBS value (Minimum depth of the queue in kilo bytes).

The no form of the command sets it to the default value.

Default 32kbytes

Parameters *size-in-kbytes* — Specifies the minimum depth of the queue in kilo bytes.

Values [0..10240]

mbs

Syntax **[no]** **mbs** *size-in-kbytes*

Context	config>qos>network-queue>queue
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>This command specifies the MBS value (Maximum depth of the queue in kilo bytes).</p> <p>The no form of the command sets it to the default value.</p>
Default	512kbytes
Parameters	<p><i>size-in-kbytes</i> — Specifies the minimum depth of the queue in kilo bytes.</p> <p>Values [0..12800]</p>

priority

Syntax	[no] priority <i>level</i>
Context	config>qos>network-queue>queue
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>Defines the queue priority. The queue priority is used by the scheduler to determine the order of service in both the within-cir loop and within-pir loop. Higher priority queues are serviced before lower priority queues.</p> <p>The no form of the command sets it to the default value.</p>
Default	1
Parameters	<p><i>level</i> — Specifies the priority of the queue.</p> <p>Values [1..4]</p>

slope-policy

Syntax	[no] slope-policy <i>name</i>
Context	config>qos>network-queue>queue
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C</p> <p>The slope-policy command is used to override the default slope-policy configuration for the queue. The specified slope-policy-name must exist as a current slope policy name. If the slope policy does not exist, the slope-policy command will fail. If a slope policy is currently associated with a queue, the slope policy cannot be removed from the system.</p>

Network Queue Command Descriptions

The slope policy contains the ring and non-ring High and Low WRED slope definitions that will be used by the queue. The non-ring slopes are used by the traffic received on access SAP ingress and sent out of access SAP egress queues. The ring slopes are used by the traffic received on access-uplink port ingress and sent out of access-uplink port egress queues.

For more information about ring and non-ring buffer pools and slope usage see the above section.

If the slope-policy command is not executed or the no slope policy command is executed, the default slope policy will be associated with the queue.

The no form of the command restores the default slope policy to the queue.

Parameters *name* — The name parameter is required and must specify an existing slope policy name. If slope-policy-name does not exist, the slope-policy command will fail.

Values 32 chars max

rate

Syntax **rate** [*cir cir-percent*] [*pir pir-percent*]
no rate

Context config>qos>network-queue>queue

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the port. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by over subscription factors or available egress bandwidth. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.

The rate command can be executed at anytime, altering the PIR and CIR rates for all queues created on the access ports.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (100, 0).

Parameters *cir percent* — Defines the percentage of the guaranteed rate allowed for the queue. When the **rate** command is executed, a valid CIR setting must be explicitly defined. When the **rate** command has not been executed, the default **CIR of 0** is assumed. Fractional values are not allowed and must be given as a positive integer.

The actual CIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Default 0

Values 0 — 100

pir percent — Defines the percentage of the maximum rate allowed for the queue. When the **rate** command is executed, the PIR setting is optional. When the **rate** command has not been executed, or the PIR parameter is not explicitly specified, the default PIR of 100 is assumed. Fractional values are not allowed and must be given as a positive integer.

Default 100

Values 1— 100 percent

weight

Syntax **weight** *weight*
no weight

Context config>qos>network-queue>queue

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Allows users to specify the weight of the queue.

The configured weight determines the proportion of available bandwidth that is given to this queue in comparison to other queues contending for bandwidth at the same priority level.

The no form of the command sets the weight to default of one.

Parameters *weight* — Specifies the weight of the queue.

Values [1...100]

Show Commands

network-queue

Syntax **network-queue** [*network-queue-policy-name*] [**detail**]

Context show>qos

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

This command displays network queue policy information.

Parameters *network-queue-policy-name* — The name of the network queue policy.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Network Queue Command Descriptions

detail — Includes each queue's rates and adaptation-rule and & cbs details. It also shows FC to queue mapping details.

Table 25: Network Queue Labels and Descriptions

Label	Description
Policy	The policy name that uniquely identifies the policy.
Accounting	Displays whether the accounting mode is packet-based or frame-based.
Description	A text string that helps identify the policy's context in the configuration file.
Port-Id	Displays the physical port identifier where the network queue policy is applied.
Queue	Displays the queue ID.
CIR	Displays the committed information rate.
PIR	Displays the peak information rate.
CBS	Displays the committed burst size.
FC	Displays FC to queue mapping.

Output

Sample for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C:

```
*A:SAH01-051>show>qos# network-queue "default" detail
```

```
=====
QoS Network Queue Policy
=====
```

```
-----
Network Queue Policy (default)
-----
```

```
Policy          : default
Description     : Default network queue QoS policy.
Wrr Policy      :
Pkt.Byte Offset: 0
-----
```

```
-----
FC To Queue Mappings
-----
```

```
-----
FC      Queue
-----
```

```
be      1
l2      2
af      3
l1      4
h2      5
-----
```

```
ef      6
hl      7
nc      8
```

 Queue Rates and Rules

QueueId	CIR(%)	CIR Adpt Rule	PIR(%)	PIR Adpt Rule
Queue1	0	closest	100	closest
Queue2	25	closest	100	closest
Queue3	25	closest	100	closest
Queue4	25	closest	100	closest
Queue5	100	closest	100	closest
Queue6	100	closest	100	closest
Queue7	10	closest	100	closest
Queue8	10	closest	100	closest

 Queue Priority and Weight Details

QueueId	Priority	Weight
Queue1	1	1
Queue2	1	1
Queue3	1	1
Queue4	1	1
Queue5	1	1
Queue6	1	1
Queue7	1	1
Queue8	1	1

 High Slope Non Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	80
Queue2	Down	70	90	80
Queue3	Down	70	90	80
Queue4	Down	70	90	80
Queue5	Down	70	90	80
Queue6	Down	70	90	80
Queue7	Down	70	90	80
Queue8	Down	70	90	80

 Low Slope Non Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	80
Queue2	Down	50	75	80
Queue3	Down	50	75	80
Queue4	Down	50	75	80
Queue5	Down	50	75	80
Queue6	Down	50	75	80
Queue7	Down	50	75	80
Queue8	Down	50	75	80

Network Queue Command Descriptions

High Slope Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	80
Queue2	Down	70	90	80
Queue3	Down	70	90	80
Queue4	Down	70	90	80
Queue5	Down	70	90	80
Queue6	Down	70	90	80
Queue7	Down	70	90	80
Queue8	Down	70	90	80

Low Slope Ring

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	80
Queue2	Down	50	75	80
Queue3	Down	50	75	80
Queue4	Down	50	75	80
Queue5	Down	50	75	80
Queue6	Down	50	75	80
Queue7	Down	50	75	80
Queue8	Down	50	75	80

Slope Policies

QueueId	CBS(KBytes)	MBS(KBytes)	Slope-Policy
Queue1	50	200	default
Queue2	50	200	default
Queue3	50	200	default
Queue4	50	200	default
Queue5	50	200	default
Queue6	50	200	default
Queue7	50	200	default
Queue8	50	200	default

Network-Port Associations

No Matching Entries

=====
*A:SAH01-051>show>qos#

Service Ingress QoS Policies

In This Section

This section provides information to configure SAP ingress QoS policies using the command line interface.

Topics in this section include:

- [Overview of service ingress policy](#)
- [Basic Configurations](#)
- [Service Management Tasks](#)

Overview of service ingress policy

There is one default service ingress policy. The default policy allocates a single queue and maps all traffic to the 'be' (best-effort) forwarding class. The default policies can be copied and modified but they cannot be deleted. The default policies are identified as policy ID 1. The default policies are applied to the appropriate interface, by default. For example, the default SAP ingress policy is applied to access ingress SAPs. You must explicitly associate other QoS policies. For information about the tasks and commands necessary to access the command line interface and to configure and maintain your 7210 SAS devices, refer to the CLI Usage chapter in the 7210 SAS OS Basic System Configuration Guide.

In a service ingress QoS policy, user can define up to 8 queues per policy, with up to 2 queues per forwarding class. In the case of VPLS service, four types of forwarding are supported (which is not to be confused with forwarding classes); unicast, multicast, broadcast, and unknown. Multicast, broadcast, and unknown types are flooded to all destinations within the service while the unicast forwarding type is handled in a point-to-point fashion within the service. All these traffic types use the same queue (in other words, a separate queue for multicast, broadcast, and unknown unicast traffic types cannot be defined). Unicast and multipoint traffic can be defined to use the same queue or different queues per forwarding class. In other words, 8 queues are shared by unicast and multicast traffic types and if user allocates a dedicated multicast queue for BUM traffic, the number of queues available for unicast traffic reduces.

Configuration Guidelines for SAP ingress Policy

- User has an option to configure lesser number of ingress queues per SAP. In hardware queues are allocated in groups of 2 and reserved for use by the SAP even though user specifies an odd value.
- FC to queue map can be defined, this allows user to assign the packets classified into a particular FC to any one of the queues configured for the SAP.
- Option is available to use up to 2 queues per FC, with one queue for Unicast traffic and one queue for BUM traffic. This option can be used with multipoint service, example – VPLS service. BUM traffic shares a single queue per FC. In other words, it is not possible to use individual queue for each of broadcast, unknown-unicast and multicast traffic. It is possible to define the same queue for unicast and BUM traffic, if the user so desires. For example, user can assign 2 queues per FC, such that unicast traffic uses one of the queues and the BUM traffic uses the other queue. This allows users to have 4 FCs per SAP with 2 queues per FC; OR user can have 7 FCs per SAP with 1 queue per FC and the 8th queue being shared by BUM traffic of all the FCs, OR a mix and match is allowed. If a multicast queue is not assigned to a FC explicitly it uses queue 1 (the default queue of the policy).
- The queue parameters such as queue shaper rate (CIR/PIR), cbs and mbs, queue priority and weight can be defined. The assigned priority and weight is used to determine the priority and weight of the queue in both the CIR and PIR scheduling loop.
- Allow configuration of WRED slopes (per queue) – high-slope and low-slope. Depending on the queue mode and the profile assigned to the packet on SAP ingress classification, one of the configured WRED slopes is used to evaluate if a buffer can be allocated to the packet. In-profile packets use the high-slope and out-of-profile packets use the low-slope.
- Supports Strict priority (SP) scheduling and Weighted-Fair Queuing (WFQ) scheduling for SAP ingress queues. For more information about service ingress scheduling, see the [Schedulers](#).
- If both IP criteria and MAC criteria are configured in SAP ingress QoS classification then they need resources from two different slices. In other words, the MAC and IP criteria entries cannot be located in the same slice/chunk. If resources are not found in two different slices, then the association of the policy fails.

Resource Allocation for Service Ingress QoS classification policy

The available global pool of ingress internal CAM hardware resources can be allocated as per user needs for use with different features such as SAP ingress QoS policy, ingress ACLs, etc. SAP ingress QoS can be allocated classification for use from this pool. Further on, resources can be allocated for SAP ingress QoS policy classification IPv4, IPv6 and MAC match criteria, based on the operator needs. Users can modify the resource allocated to scale the number of entries available per match criteria or scale the number of SAPs. The resources from the global ingress internal CAM pool are allocated in chunks with fixed number of entries.

The number of chunks to be allotted for SAP ingress QoS policy is specified using the CLI command

```
configure> system> resource-profile> ingress-internal-tcam> qos-sap-ingress-resource
```

User can specify a limit for the amount of resources required for SAP ingress QoS policies and also has an option to limit the amount of resources used per match criteria supported for SAP ingress QoS policies. A given chunk can be used for either MAC criteria or IP criteria or IPv6 criteria.

Before associating an IPv4 criteria or IPv6 criteria or MAC criteria SAP ingress policy to a SAP, resources must be allocated. Until resources are allocated for use, software fails all attempts to associate a policy with a SAP. When the user allocates resources for use by SAP ingress QoS policies using the CLI command `configure> system> resource-profile> ingress-internal-tcam> qos-sap-ingress-resource`, the system allocates resources in chunks of 512 entries. A single entry is used by each entry configured with IP criteria and MAC criteria.

The above resources set the maximum limit on the resources available for use by all the SAP ingress policies in use simultaneously on the system. The software manages the resource chunks allocated to SAP ingress QoS policy pool and allocates the entries in the chunks when a SAP ingress QoS policy is associated with a SAP. In other words, the software allocates the resources required by a SAP from the chunks depending on whether the SAP ingress policy uses ip-criteria or mac-criteria or ipv6-criteria and the number of entries configured in the SAP ingress policy.

When user allocates chunks of resources using the CLI command, `configure> system> resource-profile> ingress-internal-tcam> qos-sap-ingress-resource`, the resources are used only towards classification entries configured under IPv4 criteria or MAC criteria. A single resource is required for each entry configured using any of these two criteria.

Overview of service ingress policy

If user needs to use IPv6 criteria, then resources need to be allocated using the CLI command `configure> system> resource-profile> ingress-internal-tcam> qos-sap-ingress-resource> mac-ipv4-ipv6-128-match-enable`. The resources allocated using this command is used by software to allocate entries for classification entries configured under `ipv6-criteria` and `ipv4-criteria` `mac-criteria`. Each IPv6 or IPv4 or MAC classification entry consumes 2 resources from this pool, reducing the number of classification entries that can be accommodated in a single chunk to 256. The user can choose to allocate all the chunks allocated for SAP ingress QoS classification for `ipv6` criteria or allocate only a portion of it.

The “`tools> dump> system-resources`” command displays the current usage and availability of the resources. One or more entries per chunk are reserved for system use.

Default SAP Ingress Policy

The default policy 1, maps all traffic to default forwarding class 'be' and maps FC 'be' to queue 1. Queue 1 is configured with cir 0 and pir max.

```
A:SAH01-071>config>qos>sap-ingress# info detail
-----
policy-name "default"
description "Default SAP ingress QoS policy."
no ip-mac-match
scope template
queue 1 create
  adaptation-rule cir closest pir closest
  rate cir 0 pir max
  slope-policy "default"
  mbs 60
  cbs 10
  priority 1
  weight 1
exit
fc "af" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "be" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "ef" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "h1" create
  queue 1
  multicast-queue 1
  no use-dei
exit
```

```

fc "h2" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "l1" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "l2" create
  queue 1
  multicast-queue 1
  no use-dei
exit
fc "nc" create
  queue 1
  multicast-queue 1
  no use-dei
exit
default-fc "be"
no dot1p-classification
no dscp-classification

```

SAP Ingress Policy Defaults

Table 26: SAP Ingress Policy Defaults

Field	Default	
description	"Default SAP ingress QoS policy."	
scope	template	
queue	1	
adaptation-rule	cir closest pir closest	
rate	pir = max, cir= 0	
cbs	10KBytes - default	
mbs	60KBytes - default	
priority	1	
weight	1	
default-fc	be	

Use of Index file by SAP QoS Ingress policy

7210 uses an index file to store the map which indicates the QoS resource allocation to SAPs. This file is used on reboot to ensure that all the SAPs that were created successfully before reboot can be created again on a reboot. Without an index file the system cannot ensure this (that is, without an index file it is possible that all the SAPs that were configured successfully, may fail on a reboot after saving the configuration file). The file is stored in the flash. On reboot if the file is found, the system allocates resources as per the stored map. If the file is not found the system implements a best-fit algorithm and tries to allocate resources for all the SAPs on a first-come-first-served basis (Note : There is no guarantee that resources will be allocated to all SAPs). Hence, when the file is not present it is possible that configuration saved, does not execute successfully after the reboot.



Note: The index file used for QoS map is different from the one used for storing Interface indexes.

Basic Configurations

A basic service ingress QoS policy must conform to the following:

- Have a unique service ingress QoS policy ID.
- Have a QoS policy scope of template or exclusive.
- Have at least one default unicast forwarding class queue.
- Use of multipoint forwarding class queue is optional.

Create Service Ingress QoS Policies

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to a SAP, a default QoS policy is applied.

- [Service Ingress QoS Policy](#)

Service Ingress QoS Policy of 7210 SAS-K 2F2T1C and

7210 SAS-K 2F4T6C

To create a service ingress policy, define the following:

- A policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- Specify a default forwarding class for the policy. All packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class.
- Define forwarding class parameters.
 - Modify the unicast/queue default value to override the default unicast forwarding type queue mapping for **fc** *fc-name*.
 - Modify the **multicast/queue** default value to override the default multicast forwarding type queue mapping for **fc** *fc-name*.
- Specify the appropriate classification criteria - IPv4/IPv6 or MAC criteria or both IP and MAC criteria. You can define IPv4/IPv6, MAC-based and MAC and IP based SAP ingress policies to select the appropriate ingress meter and corresponding forwarding class for matched traffic.
- A SAP ingress policy is created with a template scope. The scope can be modified to exclusive for a special one-time use policy. Otherwise, the **template** scope enables the policy to be applied to multiple SAPs.

The following displays an service ingress policy configuration:

```
A:ALA-7>config>qos>sap-ingress# info
-----
...
      sap-ingress 100 create
          description "Used on VPN sap"
...
-----
A:ALA-7>config>qos>sap-ingress#
```

Service Ingress QoS Queues for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

To create service ingress queue parameters, define the following:

- **A new Queue ID value** — The system will not dynamically assign a value.
- **Queue parameters** — Rate, Slope-policy, CBS, MBS, Priority, and Weight.

```
A:ALA-7>config>qos# info
```

Basic Configurations

```
echo "QoS Policy Configuration"
#-----
      sap-ingress 1 create
      policy-name "default"
      description "Default SAP ingress QoS policy."
      no ip-mac-match
      scope template
      queue 1 create
      adaptation-rule cir closest pir closest
      rate cir 0 pir max
      slope-policy "default"
      mbs 60
      cbs 10
      priority 1
      weight 1
      exit
.....
#-----
A:ALA-7>config>qos#
```

SAP Ingress Forwarding Class Configuration for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

The following displays a SAP ingress forwarding class configurations with unicast queues and multicast queues for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

```
*A:dut-i>config>qos>sap-ingress$ info
-----
.....
      queue 1 create
      exit
      queue 2 create
      rate cir 1000 pir 2000
      exit
      fc "af" create
      queue 1
      multicast-queue 2
      exit
      fc "nc" create
      queue 2
      multicast-queue 1
      exit
.....
-----
config>qos>sap-ingress$ info
```

Service Ingress Dot1p and IP DSCP Criteria

The following displays the configuration of Dot1p classification policy and IP DSCP classification policy used for ingress classification and its association with SAP ingress policy on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

```
A:ALA-7>config>qos>dot1p-classification# info
#-----
.....
dot1p 0 fc "be" profile out
dot1p 1 fc "l2" profile in
dot1p 2 fc "af" profile out
dot1p 3 fc "af" profile in
dot1p 4 fc "h2" profile in
dot1p 5 fc "ef" profile in
dot1p 6 fc "h1" profile in
dot1p 7 fc "nc" profile in
.....
#-----

A:ALA-7>config>qos>sap-ingress# info
#-----
.....

dot1p-classification 1

.....
#-----
A:ALA-7>
```

Service Ingress IP Match Criteria

When specifying SAP ingress match criteria, only one match criteria type can be configured in the SAP ingress QoS policy.

The following displays an ingress IP criteria configuration:

```
A:ALA-7>config>qos# info
...
#-----
echo "QoS Policy Configuration"
#-----
...
          sap-ingress 100 create
...
          ip-criteria
          entry 10 create
          description "Entry 10-FC-AF"
          match dscp af12
          exit
```

Basic Configurations

```
        action fc af
        exit
        entry 20 create
        description "Entry 20-FC-BE"
        match dscp be
        exit
        no action
        exit
        exit
exit
..
#-----
A:ALA-7>config>qos#
```

Service Ingress MAC Match Criteria

To configure service ingress policy MAC criteria, define the following:

- A new entry ID value. Entries must be explicitly created. The system will not dynamically assign entries or a value.
- The action to associate the forwarding class with a specific MAC criteria entry ID.
- A description. The description provides a brief overview of policy features.

The following displays an ingress MAC criteria configuration:

```
A:ALA-7>config>qos# info
...
#-----
echo "QoS Policy Configuration"
#-----
...
        sap-ingress 101 create
...
        mac-criteria
        entry 10 create
        description "Entry10"
        match
            dst-mac 04-67-ff-00-00-01 ff-ff-ff-ff-ff-ff
            dot1p 7 7
        exit
        action fc be
        exit
        exit
        exit
#-----
A:ALA-7>config>qos#
```

Applying Service Ingress Policies

Note: SAP ingress QoS policies is supported only on for access SAPs only on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Apply SAP ingress policies to the following service SAPs:

- [Epipe Service](#)
- [VPLS](#)
- [IES](#)

Epipe Service

The following output displays an Epipe service configuration with SAP ingress policy 100 applied to the SAP.

```
A:ALA-7>config>service# info
-----
      epipe 6 customer 6 vpn 6 create
      description "Epipe service to west coast"
      sap 1/1/10:10 create
      exit
      ingress
      qos 100
      exit
    exit
  exit
A:ALA-7>config>service#
```

VPLS

The following output displays a VPLS service configuration with SAP ingress policy 100.

```
A:ALA-7>config>service# info
-----
      vpls 700 customer 7 vpn 700 create
      description "test"
      stp
      shutdown
      exit
      sap 1/1/9:10 create
      ingress
      qos 100
      exit
    exit
  exit
A:ALA-7>config>service#
```

Service Management Tasks

```
A:ALA-7>config>service#
```

IES

The following output displays a IES service configuration.

```
A:ALA-7>config>service# info
-----
...
ies 1 customer 1 create
interface "to-cl" create
address 11.1.0.1/24
sap 1/1/10:100 create
  ingress
    qos 100
  exit
exit
no shutdown
exit
...
-----
A:ALA-7>config>service#
```

Service Management Tasks

This section discusses the following service management tasks:

- [Deleting QoS Policies](#)
- [Copying and Overwriting QoS Policies](#)
- [Remove a Policy from the QoS Configuration](#)
- [Editing QoS Policies](#)

Deleting QoS Policies

Every service SAP is associated, by default, with the appropriate ingress policy (policy-id **1**). You can replace the default policy with a customer-configured policy, but you cannot entirely remove the policy from the SAP configuration. When you remove a non-default service ingress policy, the association reverts to the default policy-id **1**.

A QoS policy cannot be deleted until it is removed from all SAPs where they are applied.

```
A:ALA-7>config>qos# no sap-ingress 100
```

```
MINOR: CLI SAP ingress policy "100" cannot be removed because it is in use.
A:ALA-7>config>qos#
```

Remove a QoS Policy from Service SAP(s)

The following Epipe service output examples show that the SAP service ingress reverted to policy-id “1” when the non-default policies were removed from the configuration.

```
A:ALA-104>config>service>epipe# info detail
-----
description "Distributed Epipe service to west coast"
    no tod-suite
    dotlag
    exit
    ingress
        qos 1
        no filter
    exit
    egress
        no filter
    exit
    no collect-stats
    no accounting-policy
    no shutdown
-----
A:ALA-7>config>service>epipe#
```

Copying and Overwriting QoS Policies

You can copy an existing service ingress policy, rename it with a new policy ID value, or overwrite an existing policy ID. The overwrite option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos# copy {sap-ingress} source-policy-id dest-policy-id [overwrite]`

```
*A:ALU-7210>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
    sap-ingress 100 create
        description "Used on VPN sap"
        meter 1 create
        exit
        meter 2 multipoint create
        exit
        meter 10 create
            rate cir 11000
        exit
```

Service Management Tasks

```
meter 11 multipoint create
exit
exit
sap-ingress 101 create
description "Used on VPN sap"
meter 1 create
exit
meter 2 multipoint create
exit
meter 10 create
rate cir 11000
exit
meter 11 multipoint create
exit
exit
sap-ingress 200 create
description "Used on VPN sap"
meter 1 create
exit
meter 2 multipoint create
```

Remove a Policy from the QoS Configuration

CLI Syntax: config>qos# no sap-ingress policy-id

Example: config>qos# no sap-ingress 100

Editing QoS Policies

You can change QoS existing policies and entries. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors copy the policy to a work area, make the edits, and then write over the original policy.

Service SAP QoS Policy Command Reference

- [Command Hierarchy](#)
- [Operational Commands](#)
- [Show Commands](#)

Command Hierarchy

- [Service Ingress QoS Policy Command for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C](#)
- [Operational Commands](#)

Service Ingress QoS Policy Command for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

```

— config
  — qos
    — [no] sap-ingress policy-id [create]
      — default-fc fc [profile {in | out | use-dei}]
      — no default-fc
      — description description-string
      — no description
      — dot1p-classification policy-id
      — no dot1p-classification
      — dscp-classification policy-id
      — no dscp-classification
      — [no] fc fc-name [create]
        — multicast-queue queue-id
        — no multicast-queue
        — no queue
        — queue queue-id
        — no use-dei
        — use-dei
      — [no] ip-criteria
        — [no] entry entry-id [create]
          — action [fc fc-name][profile {in | out | use-dei}]
          — no action
          — description description-string
          — no description
          — match [protocol protocol-id]
          — no match
            — dscp dscp-name
            — no dscp

```

Service SAP QoS Policy Command Reference

- **dst-ip** {*ip-address/mask* | *ip-address netmask*}
- **no dst-ip**
- **dst-port** *fc* {**eq**} *dst-port-number*
- **no dst-port**
- **fragment** {**true** | **false**}
- **no fragment**
- **src-ip** {*ip-address/mask* | *ip-address netmask*}
- **no src-ip**
- **src-port** {**eq**} *src-port-number*
- **no src-port**
- **dscp** [*old-entry-id new-entry-id*]
- **[no] ipv6-criteria**
 - **[no] entry** *entry-id* [*create*]
 - **action** [*fc fc-name*] [**profile** {**in** | **out** | **use-dei**}]
 - **no action**
 - **description** *description-string*
 - **no description**
 - **match** [**next-header** *next-header*]
 - **no match**
 - **dscp** *dscp-name*
 - **no dscp**
 - **dst-ip** {*ip-address/mask* | *ip-address netmask*}
 - **no dst-ip**
 - **dst-port** *fc* {**eq**} *dst-port-number*
 - **no dst-port**
 - **fragment** {**true** | **false**}
 - **no fragment**
 - **src-ip** {*ip-address/mask* | *ip-address netmask*}
 - **no src-ip**
 - **src-port** {**eq**} *src-port-number*
 - **no src-port**
 - **dscp** [*old-entry-id new-entry-id*]
- **[no] mac-criteria**
 - **[no] entry** *entry-id* [*create*]
 - **action** [*fc fc-name*] [**profile** {**in** | **out** | **use-dei**}]
 - **no action**
 - **description** *description-string*
 - **no description**
 - **[no] match**
 - **dst-mac** *ieee-address* [*ieee-address-mask*]
 - **no dst-mac**
 - **etype** *0x0600..0xffff*
 - **no etype**
 - **inner-dot1p** *dot1p-value* [*dot1p-mask*]
 - **no inner-dot1p**
 - **inner-tag** *value* [*vid-mask*]
 - **no inner-tag**
 - **no outer-dot1p**
 - **outer-dot1p** *dot1p-value* [*dot1p-mask*]
 - **no outer-tag**
 - **outer-tag** *value* [*vid-mask*]
 - **src-mac** *ieee-address* [*ieee-address-mask*]
 - **no src-mac**
 - **dscp** *old-entry-id new-entry-id*

- **queue** *queue-id*
- **no queue**
 - **[no] adaptation-rule** [*pir adaptation-rule*] [*cir adaptation-rule*]
 - **cbs** *size-in-kbyte*
 - **no cbs**
 - **mbs** *size in kbytes*
 - **no mbs**
 - **no priority**
 - **priority** *level*
 - **no rate**
 - **rate cir** *cir-rate-in-kbps* [**pir** *pir-rate-in-kbps*]
 - **no slope-policy**
 - **slope-policy** *name*
 - **no weight**
 - **weight** *weight*
- **scope** {*exclusive* | *template*}
- **no scope**

Operational Commands

- **config**
 - **qos**
 - **copy** *sap-ingress src-pol dst-pol* [*overwrite*]

Show Commands

- **show**
 - **qos**
 - **sap-ingress** *policy-id* [*detail* | *association* | *match-criteria*]

Command Descriptions

Configuration Commands

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>sap-ingress config>qos>sap-ingress>ip-criteria>entry config>qos>sap-ingress>mac-criteria>entry
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. This command creates a text description stored in the configuration file for a configuration context. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands



Note: The 7210 platform QoS capabilities varies across platforms. In the description the term queue/meter is used and based on the platform capabilities both of them or one of them can be used. The description also mentions the capabilities of the node/platform in certain commands, as applicable.

Command Descriptions

copy

Syntax	copy sap-ingress <i>src-pol dst-pol</i> [overwrite] copy dot1p-classification <i>src-pol dst-pol</i> [overwrite] copy dscp-classification <i>src-pol dst-pol</i> [overwrite]
Context	config>qos
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. This command copies existing QoS policy entries for a QoS policy-id to another QoS policy-id. The copy command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the overwrite keyword.
Parameters	sap-ingress <i>src-pol dst-pol</i> — Indicates that the source policy ID and the destination policy ID are SAP ingress policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy. Values 1 — 65535 overwrite — Specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If overwrite is not specified, an error will occur if the destination policy ID exists. dot1p-classification — Indicates that the source policy ID and the destination policy ID are Dot1p classification policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy. Values 1 — 65535 <i>dscp-classification</i> — Indicates that the source policy ID and the destination policy ID are IP DSCP classification policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy. Values 1 — 65535

dscp

Syntax	dscp <i>dscp-name</i> no dscp
Context	config>qos>network>ingress>ip-criteria>entry>match config>qos>network>ingress>ipv6-criteria>entry>match
Description	This command configures a DiffServ Code Point (DSCP) code point to be used as a network ingress QoS policy match criterion. The no form of this command removes the DSCP match criterion.

Parameters *dscp-name* — Specifies a dscp name that has been previously mapped to a value using the *dscp-name* command. The DiffServ code point can only be specified by its name.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

dst-ip

Syntax **dst-ip** {*ip-address/mask* | *ip-address netmask*}
dst-ip {*ipv6-address/prefix-length* | *ipv6-address ipv6-address-mask*}
no dst-ip

Context config>qos>network>ingress>ip-criteria>entry>match
 config>qos>network>ingress>ipv6-criteria>entry>match

Description This command configures a destination address range to be used as a network ingress QoS policy match criterion.

To match on the destination address, specify the address and its associated mask, e.g., 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used.

The **no** form of this command removes the destination IP address match criterion.

Parameters *ip-address* — The IP address of the destination IP or IPv6 interface. This address must be unique within the subnet and specified in dotted decimal notation.

Values

ip-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
	prefix-length: 1 to 128

dst-port

Syntax **dst-port** {*eq*} *dst-port-number*
no dst-port

Context config>qos>network>ingress>ip-criteria>entry>match
 config>qos>network>ingress>ipv6-criteria>entry>match

Description This command configures a destination TCP or UDP port number or port range for a network ingress QoS policy match criterion.

Command Descriptions

The **no** form of this command removes the destination port match criterion.

Default n/a

Parameters **eq** *dst-port-number* — The TCP or UDP port numbers to match specified as equal to (eq) to the destination port value specified as a decimal integer.

Values 1 to 65535 (decimal)

range *start end* — The range of TCP or UDP port values to match specified as between the start and end destination port values inclusive.

Values 1 to 65535 (decimal)

fragment

Syntax **fragment** {**true** | **false**}
no fragment

Context config>qos>ingress>ip-criteria>entry>match

Description This command configures fragmented or non-fragmented IP packets as a network ingress QoS policy match criterion.

The **no** form of this command removes the match criterion and matches all packets regardless of whether they are fragmented or not.

Parameters **true** — Configures a match on all fragmented IP packets. A match will occur for all packets that have either the MF (more fragment) bit set OR have the Fragment Offset field of the IP header set to a non-zero value.

false — Configures a match on all non-fragmented IP packets. Non-fragmented IP packets are packets that have the MF bit set to zero and have the Fragment Offset field also set to zero.

src-ip

Syntax **src-ip** {*ip-address/mask* | *ip-address ipv4-address-mask*}
src-ip {*ipv6-address/prefix-length* | *ipv6-address ipv6-address-mask*}
no src-ip

Context config>qos>network>ingress>ip-criteria>entry>match
config>qos>network>ingress>ipv6-criteria>entry>match

Description This command configures a source IPv4 or IPv6 address range to be used as a network ingress QoS policy match criterion.

To match on the source IPv4 or IPv6 address, specify the address and its associated mask, for example, 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used for IPv4.

The **no** form of the command removes the source IPv4 or IPv6 address match criterion.

Default	No source IP match criterion.
Parameters	<p>ip-address — Specifies the source IPv4 address specified in dotted decimal notation.</p> <p>Values ip-address: a.b.c.d</p> <p>mask — Specifies the length in bits of the subnet mask.</p> <p>Values 1 to 32</p> <p>ipv4-address-mask — Specifies the subnet mask in dotted decimal notation.</p> <p>Values a.b.c.d (dotted quad equivalent of mask length)</p> <p>ipv6-address — Specifies the IPv6 prefix for the IP match criterion in hex digits.</p> <p>Values ipv6-address: x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D</p> <p>prefix — Specifies the IPv6 prefix length for the ipv6-address expressed as a decimal integer.</p> <p>Values 1 to 128</p> <p>mask — Specifies the eight 16-bit hexadecimal pieces representing bit match criteria.</p> <p>Values x:x:x:x:x:x:x (eight 16-bit pieces)</p>

src-port

Syntax	src-port {eq} <i>src-port-number</i> no src-port
Context	config>qos>network>ingress>ip-criteria>entry>match config>qos>network>ingress>ipv6-criteria>entry>match
Description	<p>This command configures a source TCP or UDP port number or port range for a network ingress QoS policy match criterion.</p> <p>The no form of this command removes the source port match criterion.</p>
Default	No src-port match criterion.
Parameters	<p>eq <i>src-port-number</i> — The TCP or UDP port numbers to match specified as equal to (eq) to the source port value specified as a decimal integer.</p> <p>Values 1 to 65535 (decimal)</p> <p>range <i>start end</i> — The range of TCP or UDP port values to match specified as between the start and end source port values inclusive.</p> <p>Values 1 to 65535 (decimal)</p>

Command Descriptions

renum

Syntax	renum <i>old-entry-id new-entry-id</i>
Context	config>qos>sap-ingress>ip-criteria config>qos>sap-ingress>ipv6-criteria config>qos>sap-ingress>mac-criteria
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. This command renumbers existing QoS policy criteria entries to properly sequence policy entries. This can be required in some cases since the 7210 SAS exits when the first match is found and executes the actions in accordance with the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.
Parameters	<i>old-entry-id</i> — Enter the entry number of an existing entry. Default none Values 1 — 64 <i>new-entry-id</i> — Enter the new entry-number to be assigned to the old entry. Default none Values 1 — 64

Service Ingress QoS Policy Commands

sap-ingress

Syntax	[no] sap-ingress <i>policy-id [create]</i>
Context	config>qos
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. This command is used to create or edit the ingress policy. The ingress policy defines the Service Level Agreement (SLA) enforcement service packets receive as they ingress a SAP. SLA enforcement is accomplished through the definition of meters/queues (depends on the support available on a platform) that have Forwarding Class (FC), Committed Information Rate (CIR), Peak Information Rate (PIR), Committed Burst Size (CBS), and Maximum Burst Size (MBS) characteristics. The simplest policy defines a single queue/meter that all ingress traffic flows through. Complex policies have multiple meters/queues combined with classification entries that indicate which meter/queue a packet will flow through. Policies in effect are templates that can be applied to multiple services as long as the scope of the policy is template. Meters defined in the policy are not instantiated until a policy is applied to a service SAP.

Depending on the support available on different 7210 platforms, SAP ingress policies can be defined with either Dot1p as the match criteria or IP DSCP as the match criteria or IP headers as the match criteria or MAC headers or both as the match criteria. On 7210 SAS user has an option to use Dot1p, IP DSCP, IPv4/IPv6 criteria and MAC criteria.

Only one service ingress policy can be provisioned. The SAP ingress policy with *policy-id* 1 is a system-defined policy applied to services when no other policy is explicitly specified. The system SAP ingress policy can be modified but not deleted. The **no sap-ingress** command restores the factory default settings when used on *policy-id* 1. The default SAP Ingress policy for different platforms are listed above [Default SAP Ingress Policy](#) .

Any changes made to the existing policy, using any of the sub-commands are applied immediately to all services where this policy is applied. For this reason, when many changes are required on a policy, it is recommended that the policy be copied to a work area policy ID. That work-in-progress policy can be modified until complete and then written over the original policy-id. Use the **config qos copy** command to maintain policies in this manner.

Note:



- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, SAP ingress policy with multiple criteria can be associated with the SAP. The system defines a match-order to match against the classification rules specified in the policy. Please see the QoS overview section to know more about the match order.
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, before using IPv4, IPv6 and MAC match criteria, resources must be allocated using the CLI command `config> system> resource-profile>ingress-internal-tcam> qos-sap-ingress-resource`. Please read the Service Ingress Qos Policies Chapter above and the 7210 Basic Systems Guide for more information about this CLI command and resource allocation.
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, service ingress only queues with shapers are supported.

The **no sap-ingress** *policy-id* command deletes the SAP ingress policy. A policy cannot be deleted until it is removed from all services where it is applied. The system default SAP ingress policy is a special case; the **no** command restores the factory defaults to policy-id 1.

Parameters *policy-id* — The *policy-id* uniquely identifies the policy.

Values 1 — 65535

create — Keyword used to create a sap ingress policy.

dot1p-classification

Syntax **dot1p-classification** *policy-id*
no dot1p-classification

Context config>qos>sap-ingress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Command Descriptions

This command allows the user to associate a dot1p classification policy which contain entries to use to map traffic received on network port or access-uplink port to a forwarding class and profile state based on the dot1p bits in the packet.

When it is defined in a network QoS policy and associated with an access-uplink port, a packet received on the access-uplink port is used to match with the Dot1p values defined in this policy. If a match is found, then the corresponding forwarding class and profile is assigned to the packet.

When it is defined in a service ingress QoS policy and associated with an access SAP ingress, a packet received on the SAP is used to match with the Dot1p values defined in this policy. If a match is found, then the corresponding forwarding class and profile is assigned to the packet.

The no form of the command disables use of dot1p-classification policy.

Parameters *policy-id* — The policy-id that uniquely identifies the policy.

Values [1..65535]

dscp-classification

Syntax **dscp-classification**
no dscp-classification

Context config>qos>sap-ingress policy-id

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command allows the user to associate a IP DSCP classification policy which contain entries to use to map traffic received on network port or access-uplink port to a forwarding class and profile state based on the IP DSCP bits in the packet.

When it is defined in a network QoS policy and associated with an access-uplink port, a packet received on the access-uplink port is used to match with the IP DSCP values defined in this policy. If a match is found, then the corresponding forwarding class and profile is assigned to the packet.

When it is defined in a service ingress QoS policy and associated with an access SAP ingress, a packet received on the SAP is used to match with the IP DSCP values defined in this policy. If a match is found, then the corresponding forwarding class and profile is assigned to the packet.

The no form of the command disables use of dot1p-classification policy.

Parameters *policy-id* — The policy-id that uniquely identifies the policy.

Values [1..65535]

scope

Syntax **scope {exclusive | template}**
no scope

Context	config>qos>sap-ingress
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>This command configures the Service Ingress QoS policy scope as exclusive or template.</p> <p>The no form of this command sets the scope of the policy to the default of template.</p>
Default	template
Parameters	<p>exclusive — When the scope of a policy is defined as exclusive, the policy can only be applied to one SAP. If a policy with an exclusive scope is assigned to a second SAP an error message is generated. If the policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP.</p> <p>template — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs on the router.</p> <p>Default QoS policies are configured with template scopes. An error is generated when the template scope parameter to exclusive scope on default policies is modified.</p>

default-fc

Syntax	default-fc <i>fc</i> [profile {in out use-dei}]
Context	config>qos>sap-ingress
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>This command configures the default forwarding class for the policy. In the event that an ingress packet does not match a higher priority (more explicit) classification command, the default forwarding class will be associated with the packet. Unless overridden by an explicit forwarding class classification rule, all packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class.</p> <p>The default forwarding class is best effort (be). The default-fc settings are displayed in the show configuration and save output regardless of inclusion of the detail keyword.</p>
Default	be
Parameters	<p><i>fc</i> — Specify the forwarding class name for the queue/meter. The value given for <i>fc</i> must be one of the predefined forwarding classes in the system.</p> <p>profile {in out} — The profile assignment action is optional. When specified, packets matching the classification entry will be explicitly classifier to either in-profile or out-of-profile. The remove the profile action for an classification entry, the action command must be re-executed without the profile action defined. The profile assigned by the user is used subsequently to determine the slope to use at the ingress and egress queuing points and is used for egress marking (if enabled).</p>

Command Descriptions

in — The in parameter is mutually exclusive to the out parameter following the profile classification action keyword. Either in or out must be specified when the profile keyword is present. When in is specified, any packets matching the classification rule will be treated as in-profile.

out — The out parameter is mutually exclusive to the in parameter following the profile classification action keyword. Either in or out must be specified when the profile keyword is present. When out is specified, any packets matching the classification rule will be treated as out-of-profile.

use-dei

Syntax	[no] use-dei
Context	config>qos>sap-ingress>fc
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>This command is used to enable DEI based classification. When enabled, the packet classified to this FC the DEI bit is used to determine the ingress profile for the packet. Packets received with DEI bit set to zero are treated as in-profile and packets with DEI bit set to one are treated as out-of-profile packets.</p> <p>When DEI based classification is enabled, it overrules the profile values specified in the classification entry used to assign the FC.</p> <p>The no form of the command disables use of DEI bit for classification of packets.</p>
Default	no use-dei

fc

Syntax	[no] fc <i>fc-name</i> [<i>create</i>]
Context	config>qos>sap-ingress
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>The fc command creates a class instance of the forwarding class <i>fc-name</i>. Once the <i>fc-name</i> is created, classification actions can be applied and can be used in match classification criteria.</p> <p>The no form of the command removes all the explicit queue mappings for <i>fc-name</i> forwarding types. The queue mappings revert to the default meters for <i>fc-name</i>.</p>

Parameters	<i>fc-name</i> — Specifies the forwarding class name for the queue. The value given for the <i>fc-name</i> must be one of the predefined forwarding classes for the system.
Default	None (Each class-name must be explicitly defined)
Values	fc: class class: be, l2, af, 11, h2, ef, h1, nc
	create — Mandatory keyword to create a forwarding class.

multicast-queue

Syntax	multicast-queue <i>queue-id</i> no multicast-queue
Context	config>qos>sap-ingress>fc config>qos>network>fc
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>This command overrides the default multicast forwarding type queue mapping for <i>fc fc-name</i>. The specified <i>queue-id</i> must exist within the policy as a multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all broadcast, unknown unicast, and multicast (BUM) traffic on a SAP or a access-uplink port using this policy is forwarded using the <i>queue-id</i>.</p> <p>The <i>no</i> form of the command removes the association of the FC and the queue. Once the <i>no</i> form is executed, the BUM traffic uses the default multicast forwarding type queue.</p> <p>The default multicast forwarding type queue is the default unicast queue (queue 1). In other words, if no multicast queue is assigned to a FC, it uses default queue 1.</p>
Default	1
Parameters	<i>queue-id</i> — specifies the queue-id for the queue.
	Values [1..8]

ip-criteria

Syntax	[no] ip-criteria
Context	config>qos>sap-ingress
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>IP criteria-based SAP ingress policies are used to select the forwarding class for matched traffic.</p> <p>This command is used to enter the context to create or edit policy entries that specify IP criteria used to match supported fields from the IP packet header, including IP DSCP.</p>

Command Descriptions

7210 SAS OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.



Note: Before associating a SAP ingress policy configured to use IPv4 criteria with a SAP, resources must be allocated to it. Please read the 7210 Basic Systems Guide for more information about this CLI command and resource allocation.

The **no** form of this command deletes all the entries specified under this node. Once IP criteria entries are removed from a SAP ingress policy, the IP criteria is removed from all services where that policy is applied.

ipv6-criteria

Syntax **[no] ipv6-criteria**

Context config>qos>sap-ingress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

IPv6 criteria-based SAP ingress policies are used to select the forwarding class for matched traffic.

This command is used to enter the node to create or edit policy entries that specify IPv6 criteria such as IP quintuple lookup or DiffServ code point.

The 7210 OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.



Note: Before associating a SAP ingress policy configured to use IPv6 criteria with a SAP, resources must be allocated to it. Please read the 7210 Basic Systems Guide for more information about this CLI command and resource allocation.

The **no** form of this command deletes all the entries specified under this node. Once ipv6-criteria entries are removed from a SAP ingress policy, the ipv6-criteria is removed from all services where that policy is applied.

mac-criteria

Syntax **[no] mac-criteria**

Context config>qos>sap-ingress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

The **mac-criteria** based SAP ingress policies are used to select the forwarding class for matched traffic.

This command is used to enter the node to create or edit policy entries that specify MAC criteria.

7210 SAS OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.



Note: Before associating a SAP ingress policy configured to use MAC criteria with a SAP, resources must be allocated to it. Please read the 7210 Basic Systems Guide for more information about this CLI command and resource allocation.

The **no** form of this command deletes all the entries specified under this node. Once mac-criteria entries are removed from a SAP ingress policy, the mac-criteria is removed from all services where that policy is applied.

queue

Syntax `[no] queue queue-id`

Context `config>qos>sap-ingress>fc`

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command overrides the default unicast forwarding type queue mapping for fc fc-name. The specified queue-id must exist within the policy as a non-multipoint queue before the mapping can be made. Once the forwarding class mapping is executed, all unicast traffic (this includes all traffic, even broadcast and multicast for services) on a SAP or an access-uplink port using this policy is forwarded using the queue-id.

The no form of this command sets the unicast (point-to-point) queue-id back to the default queue for the forwarding class (queue 1).

Parameters *queue-id* — Specifies the queue-id for the queue.

Service Ingress QoS Policy Entry Commands

action

Syntax `action [fc fc-name] profile {in|out}`
`no action`

Context `config>qos>sap-ingress>ip-criteria>entry`

Command Descriptions

```
config>qos>sap-ingress>mac-criteria>entry
```

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This mandatory command associates the forwarding class with specific IP or MAC criteria entry ID. The action command supports setting the forwarding class parameter. Packets that meet all match criteria within the entry have their forwarding class overridden based on the parameters included in the **action** parameters.

The **action** command must be executed for the match criteria to be added to the active list of entries.

Each time action is executed on a specific entry ID, the previous entered values for **fc fc-name** is overridden with the newly defined parameters.

The **no** form of the command removes the entry from the active entry list. Removing an entry on a policy immediately removes the entry from all SAPs using the policy. All previous parameters for the action is lost.

Default Action specified by the **default-fc**.

Parameters **fc fc-name** — The value given for **fc fc-name** must be one of the predefined forwarding classes in the system. Specifying the **fc fc-name** is required. When a packet matches the rule, the forwarding class is only overridden when the **fc fc-name** parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

Values be|l2|af|l1|h2|ef|h1|nc

profile {in | out} — The profile assignment action is optional. When specified, packets matching the classification entry will be explicitly classifier to either in-profile or out-of-profile. The remove the profile action for an classification entry, the action command must be re-executed without the profile action defined. The profile assigned by the user is used subsequently to determine the slope to use at the ingress and egress queuing points and is used for egress marking (if enabled).

in — The in parameter is mutually exclusive to the out parameter following the profile classification action keyword. Either in or out must be specified when the profile keyword is present. When in is specified, any packets matching the classification rule will be treated as in-profile.

out — The out parameter is mutually exclusive to the in parameter following the profile classification action keyword. Either in or out must be specified when the profile keyword is present. When out is specified, any packets matching the classification rule will be treated as out-of-profile.

use-dei — Specifies whether DEI must be used to determine the initial profile of the packet.

entry

Syntax **[no] entry entry-id [create]**

Context config>qos>sap-ingress>ip-criteria

```
config>qos>sap-ingress>mac-criteria
```

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command is used to create or edit an IP or MAC criteria entry for the policy. Multiple entries can be created using unique *entry-id* numbers.

The list of flow criteria is evaluated in a top down fashion with the lowest entry ID at the top and the highest entry ID at the bottom. If the defined match criteria for an entry within the list matches the information in the egress packet, the system stops matching the packet against the list and performs the matching entries reclassification actions. If none of the entries match the packet, the IP flow reclassification list has no effect on the packet.

An entry is not populated in the list unless the action command is executed for the entry. An entry that is not populated in the list has no effect on egress packets. If the action command is executed without any explicit reclassification actions specified, the entry is populated in the list allowing packets matching the entry to exit the list, preventing them from matching entries lower in the list. Since this is the only flow reclassification entry that the packet matched and this entry explicitly states that no reclassification action is to be performed, the matching packet will not be reclassified.

The **no** form of this command removes the specified entry from the policy. Entries removed from the policy are immediately removed from all services where that policy is applied.

Default none

Parameters *entry-id* — The *entry-id*, expressed as an integer, uniquely identifies a match criterion and the corresponding action. It is recommended that multiple entries be given *entry-ids* in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

An entry cannot have any match criteria defined (in which case, everything matches) but must have at least the keyword **action fc** *fc-name* for it to be considered complete. Entries without the action keyword will be considered incomplete and hence will be rendered inactive.

Default none

Values 1— 64

create — Required parameter when creating a flow entry when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the flow entry already exists.

match

Syntax [no] match [protocol *protocol-id*]

Context config>qos>sap-ingress>ip-criteria>entry

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Command Descriptions

This command creates a context to configure match criteria for SAP ingress QoS policy match criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured then all criteria must be satisfied (AND function) before the action associated with the match will be executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

The **no** form of the command removes the match criteria for the entry-id.

Parameters **protocol** *protocol-id* — Specifies an IP protocol to be used as a SAP QoS policy match criterion. The protocol type such as TCP / UDP / OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17).
Values) 0 — 255; check the node for values.

match

Syntax **match**
no match

Context config>qos>sap-ingress>mac-criteria>entry

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command creates a context for entering/editing match MAC criteria for ingress SAP QoS policy match criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured then all criteria must be satisfied (AND function) before the action associated with the match will be executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

The **no** form of the command removes the match criteria for the *entry-id*.

match

Syntax **match** [**next-header** *next-header*]
no match

Context config>qos>sap-ingress>ipv6-criteria>entry

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command creates a context to configure match criteria for ingress SAP QoS policy match IPv6 criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured, then all criteria must be satisfied (AND function) before the action associated with the match is executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

The **no** form of this command removes the match criteria for the *entry-id*.

Parameters

next-header next-header — Specifies the next meader to match.

The protocol type such as TCP / UDP / OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17).

Values protocol numbers accepted in DHB: 0 — 42, 45 — 49, 52 — 59, 61 — 255

keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-icmp, ipv6-no-nxt, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp
* — udp/tcp wildcard

Service Ingress MAC QoS Policy Match Commands

dst-mac

Syntax	dst-mac <i>ieee-address</i> [<i>ieee-address-mask</i>] no dst-mac
Context	config>qos>sap-ingress>mac-criteria>entry
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C Configures a destination MAC address or range to be used as a Service Ingress QoS policy match criterion. The no form of this command removes the destination mac address as the match criterion.
Default	none
Parameters	<i>ieee-address</i> — The MAC address to be used as a match criterion. Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit <i>ieee-address-mask</i> — A 48-bit mask to match a range of MAC address values. This 48-bit mask can be configured using the following formats:

Command Descriptions

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHHH	0xFFFFFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

All packets with a source MAC OUI value of 00-03-FA subject to a match condition should be specified as: 0003FA000000 0xFFFFFFFF000000

Default 0xFFFFFFFFFFFFFFF (hex) (exact match)

Values 0x0000000000000000 — 0xFFFFFFFFFFFFFFF (hex)

etype

Syntax **etype** *0x0600..0xffff*
no etype

Context config>qos>sap-ingress>mac-criteria>entry

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Configures an Ethernet type II value to be used as a service ingress QoS policy match criterion.

The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame. For e.g. 0800 is used to identify the IP v4 packets.

The Ethernet type field is used by the Ethernet version-II frames. IEEE 802.3 Ethernet frames do not use the type field. For IEEE 802.3 frames use the dsap, ssap or snap-pid fields as match criteria.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and cannot be part of the same match criteria.

The no form of this command removes the previously entered etype field as the match criteria.

Default None

Parameters *etype-value* — The Ethernet type II frame Ether type value to be used as a match criterion expressed in hexadecimal.

Values 0x0600 — 0xFFFF [1536..65535] - accepts in decimal or hex

inner-dot1p

Syntax **inner-dot1p** *dot1p-value [dot1p-mask]*
no inner-dot1p

Context config>qos>sap-ingress>mac-criteria>entry

Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C</p> <p>Configures the Dot1p value to be used as a service ingress QoS policy match criterion to match against the Dot1p value in the inner tag (the one that follows the outermost tag in the packet) of the received packet.</p> <p>The no form of this command removes the previously entered dot1p value as the match criteria.</p>
Default	no inner-dot1p
Parameters	<p><i>dot1p-value</i> — Specify the Dot1p value to match.</p> <p>Values [0..7]</p> <p><i>dot1p-mask</i> — Specify the mask value to match a range of Dot1p values.</p> <p>Values [0..7] - accepts decimal hex or binary</p>

inner-tag

Syntax	<p>inner-tag <i>value</i> [<i>vid-mask</i>]</p> <p>no inner-tag</p>
Context	config>qos>sap-ingress>mac-criteria>entry
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>Configures the VLAN value to be used as a service ingress QoS policy match criterion to match against the VLAN value in the inner tag (the one that follows the outermost tag in the packet) of the received packet.</p> <p>The no form of this command removes the previously entered VLAN tag value as the match criteria.</p>
Default	no inner-tag
Parameters	<p><i>value</i> — Specify the VLAN value to use for the match</p> <p>Values [0..4095] decimal or [0x0..0xFFFF] hex</p> <p><i>vid-mask</i> — Specify the mask value to match a range of VLAN values.</p> <p>Values [1..4095] decimal or [0x1..0xFFFF] hex</p>

outer-dot1p

Syntax	<p>outer-dot1p <i>dot1p-value</i> [<i>dot1p-mask</i>]</p> <p>no outer-dot1p</p>
Context	config>qos>sap-ingress>mac-criteria>entry
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Command Descriptions

Configures the Dot1p value to be used as a service ingress QoS policy match criterion to match against the Dot1p value in the outermost tag of the received packet.

The no form of this command removes the previously entered dot1p value as the match criteria.

Default no outer-dot1p

Parameters *dot1p-value* — Specify the Dot1p value to match.

Values [0..7]

dot1p-mask — Specify the mask value to match a range of Dot1p values.

Values [0..7] - accepts decimal hex or binary

outer-tag

Syntax **outer-tag** *value* [*vid-mask*]
no outer-tag

Context config>qos>sap-ingress>mac-criteria>entry

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Configures the VLAN value to be used as a service ingress QoS policy match criterion to match against the VLAN value in the outermost tag in the packet of the received packet.

The no form of this command removes the previously entered VLAN tag value as the match criteria.

Default no outer-tag

Parameters *value* — Specify the VLAN value to use for the match.

Values [0..4095] decimal or [0x0..0xFFFF] hex

vid-mask — Specify the mask value to match a range of VLAN values.

Values [1..4095] decimal or [0x1..0xFFFF] hex

src-mac

Syntax **src-mac** *ieee-address* [*ieee-address-mask*]
no src-mac

Context config>qos>sap-ingress>mac-criteria>entry

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command configures a source MAC address or range to be used as a service ingress QoS policy match criterion.

The **no** form of this command removes the source mac as the match criteria.

Default none

Parameters *ieee-address* — Enter the 48-bit IEEE mac address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask — This 48-bit mask can be configured using:

This 48 bit mask can be configured using the following formats

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHHH	0x0FFFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

To configure all packets with a source MAC OUI value of 00-03-FA are subject to a match condition, then the entry should be specified as: 003FA000000 0xFFFFFFFF000000

Default 0xFFFFFFFFFFFF (hex) (exact match)

Values 0x0000000000000000 — 0xFFFFFFFFFFFF (hex)

Show Commands

sap-ingress

Syntax **sap-ingress** [*policy-id*] [**detail** | **association** | **match-criteria**]

Context show>qos

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

This command displays SAP ingress QoS policy information.

Parameters *policy-id* — Displays information about the specific policy ID.

Default all SAP ingress policies

Values 1 — 65535

detail — Displays detailed policy information including policy associations.

associations- — Displays the policy associations of the sap-ingress policy.

match-criterion- — Displays the match-criterion of the sap-ingress policy.

Command Descriptions

Output

Sample Output

Show SAP Ingress Output

The following table describes SAP ingress show command output.

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Scope	Exclusive Implies that this policy can only be applied to a single SAP. Template Implies that this policy can be applied to multiple SAPs on the router.
Description	A text string that helps identify the policy's context in the configuration file.
Default FC	Specifies the default forwarding class for the policy.
Criteria-type	IP Specifies that an IP criteria-based SAP ingress policy is used to select the appropriate ingress meter and corresponding forwarding class for matched traffic. MAC Specifies that a MAC criteria-based SAP is used to select the appropriate ingress meters and corresponding forwarding class for matched traffic.
Meter	Displays the meter ID.
Mode	For 7210 SAS-E: Specifies the configured mode of the meter (trTcm1 or srTcm). For 7210 SAS-D: Specifies the configured mode of the meter (trTcm1,trTcm2 or srTcm).
CIR Admin	Specifies the administrative Committed Information Rate (CIR) parameters for the meters.
CIR Rule	min The operational CIR for the meters will be equal to or greater than the administrative rate specified using the rate command. max The operational CIR for the meter will be equal to or less than the administrative rate specified using the rate command. closest The operational PIR for the meters will be the rate closest to the rate specified using the rate command without exceeding the operational PIR.
PIR Admin	Specifies the administrative Peak Information Rate (PIR) parameters for the meters.

Label	Description (Continued)
PIR Rule	<p><code>min</code></p> <p>The operational PIR for the meter will be equal to or greater than the administrative rate specified using the rate command.</p> <p><code>max</code></p> <p>The operational PIR for the meters will be equal to or less than the administrative rate specified using the rate command.</p> <p><code>closest</code></p> <p>The operational PIR for the meters will be the rate closest to the rate specified using the rate command.</p>
CBS	<p><code>def</code></p> <p>Specifies the default CBS value for the meters.</p> <p><code>value</code></p> <p>Specifies the value to override the default reserved buffers for the meters.</p>
MBS	<p><code>def</code></p> <p>Specifies the default MBS value.</p> <p><code>value</code></p> <p>Specifies the value to override the default MBS for the meter.</p>
UCastM	Specifies the default unicast forwarding type meters mapping.
MCastM	Specifies the overrides for the default multicast forwarding type meter mapping.
BCastM	Specifies the default broadcast forwarding type meters mapping.
UnknownM	Specifies the default unknown unicast forwarding type meters mapping.
Match Criteria Entry	Specifies an IP or MAC criteria entry for the policy.
DSCP	Specifies a DiffServ Code Point (DSCP) name used for an ingress SAP QoS policy match.
FC	Specifies the entry's forwarding class.
Src MAC	Specifies a source MAC address or range to be used as a Service Ingress QoS policy match.
Dst MAC	Specifies a destination MAC address or range to be used as a Service Ingress QoS policy match.
Dot1p	Specifies a IEEE 802.1p value to be used as the match.
Ethernet-type	Specifies an Ethernet type II Ether type value to be used as a Service Ingress QoS policy match.
FC	Specifies the entry's forwarding class.
Service Association	
Service-Id	The unique service ID number which identifies the service in the service domain.
Customer-Id	Specifies the customer ID which identifies the customer to the service.

Command Descriptions

Label	Description (Continued)
SAP	Specifies the a Service Access Point (SAP) within the service where the SAP ingress policy is applied.
Classifiers required	Indicates the number of classifiers for a VPLS or Epipe service.
Meters required	Indicates the number of meters for a VPLS or Epipe service.

```
A:Dut-A>show>qos# sap-ingress 1 detail
```

```
=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (1)
-----
Policy-id           : 1                Scope           : Template
Default FC         : be
Criteria-type      : None
Mac Sub-Criteria   : None                IP Sub-Criteria  : None
IPv6 Enabled       : False
DOT1P Class Policy Id : 0                DSCP Class Policy Id : 0
MPLS Lsp Exp Class Policy*: 0
Name               : default
Description        : Default SAP ingress QoS policy.
-----
FC      Queue      MCast Queue Use Dei
-----
be      1           1           false
l2      1           1           false
af      1           1           false
l1      1           1           false
h2      1           1           false
ef      1           1           false
h1      1           1           false
nc      1           1           false
-----
Queue Rates and Rules
-----
QueueId      CIR      CIR Adpt Rule      PIR      PIR Adpt Rule
-----
Queue1      0           closest           max      closest
-----
Queue Priority and Weight Details
-----
QueueId      Priority      Weight
-----
Queue1      1           1
-----
High Slope Non Ring
-----
QueueId      State      Start-Avg(%)      Max-Avg(%)      Max-Prob(%)
-----
Queue1      Down      70                90                80
```

```

-----
Low Slope Non Ring
-----
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down       50             75           80
-----
High Slope Ring
-----
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down       70             90           80
-----
Low Slope Ring
-----
-----
QueueId      State      Start-Avg(%)  Max-Avg(%)  Max-Prob(%)
-----
Queue1       Down       50             75           80
-----
Slope Policies
-----
-----
QueueId      CBS(KBytes)  MBS(KBytes)   Slope-Policy
-----
Queue1       10           60            default
-----
Match Criteria
-----
No Matching Criteria.

SAP Associations
-----
No Associations Found.

=====
* indicates that the corresponding row element may have been truncated.
A:Dut-A>show>qos#

```

dscp-classification

- Syntax** `dscp-classification [policy-id] [detail]`
- Context** `show>qos`
- Description** **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C
- This command displays dscp-classification QoS policy information.

Command Descriptions

Parameters *policy-id* — Displays information about the specific policy ID.
detail — Displays detailed policy information.

Output

Sample for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C:

```
*A:SAH01-071>config>qos# show qos dscp-classification 6335 detail
```

```
=====
DSCP Classification Maps
=====
-----
Dscp Class Id      : 6335
Description        : (Not Specified)
-----

Dscp Bit Map                Forwarding Class                Profile
-----
cp7                          nc                               In
cp15                         h1                               In
cp23                         ef                               In
cp31                         h2                               In
cp39                         l1                               In
cp47                         af                               In
cp55                         l2                               In
-----

Network Policy Associations
-----
No Network Policy Associations found.
-----

SAP Ingress Associations
-----
SAP Ingress Id          : 17
=====
```

Label	Description
Dot1P Class Id	The ID that uniquely identifies the policy.
Dot1P Bit Map	The Dot1p value specified in the policy
Forwarding Class	The forwarding class to assigned to the packet if the received packet's Dot1p values match the Dot1p value configured.
Description	A text string that helps identify the policy's context in the configuration file.
Profile	Specifies the profile to be assigned to the packet.

dot1p-classification

Syntax	dot1p-classification [policy-id] [detail]
Context	show>qos
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C This command displays dot1p-classification QoS policy information.
Parameters	<i>policy-id</i> — Displays information about the specific policy ID. <i>detail</i> — Displays detailed policy information.

Output

Label	Description
Dot1P Class Id	The ID that uniquely identifies the policy.
Dot1P Bit Map	The Dot1p value specified in the policy
Forwarding Class	The forwarding class to assigned to the packet if the received packet's Dot1p values match the Dot1p value configured.
Description	A text string that helps identify the policy's context in the configuration file.
Profile	Specifies the profile to be assigned to the packet.

Sample output

```
A:SAH01-050>config>qos>sap-ingress# show qos dot1p-classification 10
=====
DOT1P Classification Maps
=====
-----
Dot1P Class Id      : 10
Description         : (Not Specified)
-----
A:SAH01-050>config>qos>sap-ingress#

A:SAH01-050>config>qos>sap-ingress# show qos dot1p-classification 10 detail
=====
DOT1P Classification Maps
=====
-----
Dot1P Class Id      : 10
Description         : (Not Specified)
-----
-----
Dot1P Bit Map      Forwarding Class      Profile
-----
6                   h1                       None
7                   nc                       None
-----
Network Policy Associations
```

Command Descriptions

```
-----  
No Network Policy Associations found.  
-----
```

```
-----  
SAP Ingress Associations  
-----
```

```
SAP Ingress Id           : 10  
-----  
=====
```

Service Egress Policies

In This Section

This section provides information to configure Service egress QoS policies using the command line interface.

Topics in this section include:

- [Overview](#)

Overview

The Service Egress policy defines the Service Level Agreement (SLA) for service packets as they egress on the access SAP. Service Egress QoS policies allow the definition of queue parameters along with remark policy.

With the default service egress policy, the system allocates 1 (one) queue. All the 8 FCs are mapped to use the same queue. User has an option to define up to 8 queues per policy and define the forwarding class to queue mapping. In addition, the policy allows the user to define the queue parameters. The hardware does not support a linear range of values for the rate parameters (both cir and pir). The user can specify the computation method of rates to match the rates supported by the hardware, through the configuration of adaptation-rules.

The SAP Egress policy for access SAPs supports the following:

- Per SAP egress queuing and shaping, Hierarchical shaping on SAP egress (with 3 levels of shaping) with per FC/queue shaper, per SAP aggregate shaper and per port egress rate shaper.
- SAP Egress Queues, Shaping and Scheduling
 - Provides an option to configure 2 - 8 egress queues per SAP. The hardware queues are allocated in groups of 2 and reserved for use by the SAP even though user specifies an odd value.

Overview

- Provides an option to configure the FC to queue map, allowing the user to assign the packets classified into a particular FC to any one of the queues configured for the SAP
- On SAP egress, only a single queue can be configured per FC and all traffic (unicast and BUM) share the single queue
- Allow configuration of queue shaper rate (CIR/PIR), cbs and mbs, queue priority and weight. The assigned priority and weight is used to determine the priority and weight of the queue in both the CIR and PIR scheduling loop.
- Allow configuration of WRED slopes (per queue) – high-slope and low-slope. One of the configured WRED slopes is used to allocate buffer to the packet. In-profile packets use the high-slope and out-of-profile packets use the low-slope. The profile of the packet is determined at the ingress (access uplink port ingress or Access SAP ingress) and carried through to be used at SAP egress to determine the WRED slope to apply and also to determine the egress marking value to use (if remarking/marketing is enabled).
- Support Strict priority (SP) scheduling and Weighted-Fair Queuing (WFQ) scheduling for SAP egress queues.
- SAP Egress remarking/marketing
 - Dot1p and/or IP DSCP marking must be supported on access SAP egress; Provide support for configuration of per FC Dot1p and/or IP DSCP marking, with the capability to assign different Dot1p and/or IP DSCP values for in-profile and out-of-profile packets. In addition, support for marking DEI value is available.

The following displays the access SAP egress QoS policy configuration for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C:

```
*A:SAH01-071>config>qos>sap-egress# info detail
-----
description "Default SAP egress QoS policy."
scope template
no remarking
remark 1
queue 1 create
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs 60
    cbs 10
    slope-policy "default"
    priority 1
    weight 1
exit
fc af create
    queue 1
exit
fc be create
    queue 1
exit
fc ef create
```

```

        queue 1
    exit
    fc h1 create
        queue 1
    exit
    fc h2 create
        queue 1
    exit
    fc l1 create
        queue 1
    exit
    fc l2 create
        queue 1
    exit
    fc nc create
        queue 1
    exit

```

Configuration Guidelines for access SAP Egress Policy

- User has an option to configure lesser number of egress queues per access SAP. In hardware queues are allocated in groups of 2 and reserved for use by the SAP even though user specifies an odd value.
- FC to queue map can be defined, this allows user to assign the packets classified into a particular FC to any one of the queues configured for the SAP.
- Both unicast traffic and BUM traffic share a single queue per FC. In other words, unlike service ingress policy, it is not possible to assign different queues for BUM traffic and unicast traffic.
- The queue parameters such as queue shaper rate (CIR/PIR), cbs and mbs, queue priority and weight can be defined. The assigned priority and weight is used to determine the priority and weight of the queue in both the CIR and PIR scheduling loop.
- Allow configuration of WRED slopes (per queue) - high-slope and low-slope. Depending on the queue mode and the profile assigned to the packet on SAP ingress classification, one of the configured WRED slopes is used to evaluate if a buffer can be allocated to the packet. In-profile packets use the high-slope and out-of-profile packets use the low-slope.
- Supports Strict priority (SP) scheduling and Weighted-Fair Queuing (WFQ) scheduling for SAP ingress queues. For more information about service ingress scheduling see the .

Basic Configurations

A basic service egress QoS policy must conform to the following:

- Have a unique service egress QoS policy ID.
- Have a QoS policy scope of template or exclusive.
- Have at least one forwarding class queue.

Create an access SAP Egress Policy

To create a new access SAP Egress policy, define the following:

- A SAP Egress policy name.
- Provide a brief description of the policy features.
- Provide the queue parameters for all the queues

Use the following CLI syntax to configure a SAP Egress policy:

```
*A:SAH01-051>config>qos# sap-egress
- no sap-egress <policy-id>
- sap-egress <policy-id> [create]

<policy-id>          : [1..65535]|<name:64 char max>
<create>             : keyword - mandatory while creating an entry.

[no] description     - Description for this sap-egress policy
[no] fc              + Configure forwarding-class mappings
[no] queue           + Configure a queue
[no] remark          - Specify Remarking policy for this policy
[no] remarking       - Enable/disable remarking
[no] scope           - Specify scope of the policy

*A:SAH01-051>config>qos# info detail
sap-egress 1 create
  description "Default SAP egress QoS policy."
  scope template
  no remarking
  remark 1
  queue 1 create
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs 60
    cbs 10
    slope-policy "default"
    priority 1
    weight 1
  exit
  fc af create
    queue 1
  exit
```

```
fc be create
  queue 1
exit
fc ef create
  queue 1
exit
fc h1 create
  queue 1
exit
fc h2 create
  queue 1
exit
fc l1 create
  queue 1
exit
fc l2 create
  queue 1
exit
fc nc create
  queue 1
exit
exit
```

Editing QoS Policies

Existing policies and entries can be edited through the CLI or NMS. The changes are applied immediately to all services where the policy is applicable.

To prevent configuration errors perform the following:

- Copy the policy to a work area
- Edit the policy
- Over write the original policy

Overview

Service Egress Policy Command Reference

Command Hierarchies

- [Configuration Commands for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C](#)
- [Copy Commands](#)
- [Show Commands](#)

Configuration Commands for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

```

— config
  — qos
    — sap-egress policy-id create
      — [no] description description-string
      — [no] fc fc-name [create]
        — no queue
        — queue queue-id
      — queue queue-id
      — no queue
        — [no] adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
        — cbs size-in-kbyte
        — no cbs
        — mbs size in kbytes
        — no mbs
        — no priority
        — priority level
        — no rate
        — rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]
        — no slope-policy
        — slope-policy name
        — no weight
        — weight weight
      — scope {exclusive | template}
      — [no] remark policy-id
      — [no] remarking

```

Copy Commands

```

— config
  — qos
    — copy sap-egress src-pol dst-pol [overwrite]

```

Show Commands

- show
 - qos
 - **sap-egress** [*policy-id*] [**detail** | **association**]

Commands Description

Configuration Commands

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>sap-egress
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The description command associates a text string with a configuration context to help identify the context in the configuration file.</p> <p>The no form of this command removes any description string from the context.</p>
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

SAP Egress Queue QoS Policy Commands

adaptation-rule

Syntax	adaptation-rule [<i>cir adaptation-rule</i>] [<i>pir adaptation-rule</i>] no adaptation-rule
Context	config>qos>sap-egress>queue
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Commands Description

This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **cir** and **pir** apply.

Default adaptation-rule pir closest cir closest

Parameters *adaptation-rule* — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.

Values

- pir** — Defines the constraints enforced when adapting the PIR rate defined within the **queue queue-id rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.
- cir** — Defines the constraints enforced when adapting the CIR rate defined within the **queue queue-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.
- max** — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. The hardware step size varies with the configured rate.
- min** — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. The hardware step size varies with the configured rate.
- closest** — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. The hardware step size varies with the configured rate.

cbs

Syntax **[no] cbs** *size-in-kbytes*

Context config>qos>sap-egress>queue

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command specifies the CBS value (Minimum depth of the queue in kilo bytes).

The no form of the command sets this to the default value.

Default 32Kbits

Parameters *size-in-kbytes* — Specifies the minimum depth of the queue in kilo bytes.
Values [0..10240]

mbs

Syntax **[no] mbs** *size-in-kbytes*

Context config>qos>sap-egress>queue

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.
 This command specifies the MBS value (Maximum depth of the queue in kilo bytes).
 The no form of the command sets this to the default value.

Default 512Kbits

Parameters *size-in-kbytes* — Specifies the minimum depth of the queue in kilo bytes.
Values [0..12800]

priority

Syntax **[no] priority** *level*

Context config>qos>sap-egress>queue

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.
 Defines the queue priority. The queue priority is used by the scheduler to determine the order of service in both the within-cir loop and within-pir loop. Higher priority queues are serviced before lower priority queues.
 The no form of the command sets this to the default value.

Default 1

Parameters *level* — Specifies the priority of the queue.
Values [1..4]

slope-policy

Syntax **[no] slope-policy** *name*

Context config>qos>sap-egress>queue

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Commands Description

The slope-policy command is used to override the default slope-policy configuration for the queue. The specified slope-policy-name must exist as a current slope policy name. If the slope policy does not exist, the slope-policy command will fail. If a slope policy is currently associated with a queue, the slope policy cannot be removed from the system.

The slope policy contains the ring and non-ring High and Low WRED slope definitions that will be used by the queue. The non-ring slopes are used by the traffic received on access SAP ingress and sent out of access SAP egress queues. The ring slopes are used by the traffic received on access-uplink port ingress and sent out of access-uplink port egress queues. For more information about ring and non-ring buffer pools and slope usage see the section above on buffer pools.

If the slope-policy command is not executed or the no slope policy command is executed, the default slope policy will be associated with the queue.

The no form of the command restores the default slope policy to the queue.

Parameters *name* — The name parameter is required and must specify an existing slope policy name. If slope-policy-name does not exist, the slope-policy command will fail.

Values 32 chars max

queue

Syntax **queue** *queue-id* **create**

Context config>qos>sap-egress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command creates the context to modify Queue parameters associated with a particular queue. On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the queue-ID to FC map is user defined.

The **no** form of this command is not supported

Default none

Parameters *queue-id* — Specifies the id of the queue.

Values 1 — 8

create — Mandatory keyword to create a network queue policy.

weight

Syntax [**no**] **weight** *weight*

Context config>qos>sap-egress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Allows users to specify the weight of the queue.

The configured weight determines the proportion of available bandwidth that is given to this queue in comparison to other queues contending for bandwidth at the same priority level.

The no form of the command sets the weight to default of one.

Default 1

Parameters *weight* — Specifies the value for the weight parameter. It is an integer value which specifies the proportion of available bandwidth to be allocated to this queue relative to other queues.

Values 1 — 100

sap-egress

Syntax **sap-egress** *policy-id* [**create**]

Context config>qos

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command enables the context to configure a SAP Egress policy. The SAP egress policy determines the QoS treatment to packets at service egress.

When the policy is created, by default there will be only one queue that is created. User can create up to 8 queues [1..8] and associate it to different FCs on the saps to which this sap-egress policy is attached. SAP egress policy allows the user to define the queue parameters for the eight queues.

Default 1

Parameters *policy-id* — The ID of the SAP Egress policy.

Values 1 — 65535

remark

Syntax **remark** *policy-id*
no remark

Context config>qos>sap-egress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command specifies the remarking policy ID to use for marking packets on service egress (also known as Access SAP egress).

Commands Description

The remarking policy ID must be associated with the appropriate SAP egress policy and remarking must be enabled in the service egress policy to enable marking of packets sent out of service (SAP) egress. Only remarking policy of type dot1p, dscp, or dot1p-dscp is allowed to be used when the remark policy is associated with service egress. See the [Table 29](#) for different remark policies supported on the node and its use.

The no form of the command removes the explicit association of the remark policy and associates the default remark policy. In other words, if remarking is enabled and no remark policy is executed, then the default remark policy is used to mark packets sent out. If no remark policy is executed and remarking is disabled, then packets are not remarked at all.

Parameters *policy-id* — The parameter that uniquely identifies the remark policy.

Values 1 — 65535

remarking

Syntax **[no] remarking**

Context config>qos>sap-egress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command enables the system to remark egress packets sent out of service egress (also known as, access SAP egress).

When remarking is enabled, the remark policy configured in the QoS policy context is used to determine the FC to QoS bit mapping. For example, when remarking is enabled in the sap-egress QoS policy, the remark policy associated with sap-egress QoS policy is used to determine the FC to Dot1p mapping to use for marking packets sent out of access SAPs.

Please refer to the remark policy that can be used to configure FC to priority bit markings in different QoS policies associated with different service entities, in the chapter on remark policies. For more information, see [Table 29](#).

The no form of the command disables remarking.

Default no remarking - Remarking is disabled by default.

scope

Syntax **scope {exclusive | template}**
no scope

Context config>qos>sap-egress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command configures the scope as exclusive or template. The policy's scope cannot be changed if the policy is applied to an interface/multiple ports.

The **no** form of this command sets the scope of the policy to the default of **template**.

Default	template
Parameters	<p>exclusive — When the scope of a policy is defined as exclusive, the policy can only be applied to one interface/port. If a policy with an exclusive scope is assigned to a second interface an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface.</p> <p>The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to default QoS policies are configured with template scope. An error is generated if you try to modify the scope parameter from template to exclusive scope on default policies.</p> <p>template — When the scope of a policy is defined as template, the policy can be applied to multiple interface/ports on the router.</p>

adaptation-rule

Syntax	adaptation-rule [cir adaptation-rule] [pir adaptation-rule] no adaptation-rule
Context	config>qos>sap-egress>queue
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.</p> <p>The no form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific adaptation-rule is removed, the default constraints for cir and pir apply.</p>
Default	adaptation-rule pir closest cir closest
Parameters	<p><i>adaptation-rule</i> — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.</p> <p>Values</p> <p>pir — Defines the constraints enforced when adapting the PIR rate defined within the meter meter-id rate command. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the rate command is not specified, the default applies.</p>

Commands Description

cir — Defines the constraints enforced when adapting the CIR rate defined within the **meter *meter-id* rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

max — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. The hardware step size varies with the configured rate.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. The hardware step size varies with the configured rate.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. The hardware step size varies with the configured rate.

queue

Syntax	queue <i>queue-id</i> create no queue
Context	config>qos>sap-egress
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. This command is used to configure the queue parameters. This command creates the context to modify Queue parameters associated with a particular queue. The queue is identifiable by <i>queue-id</i> and FCs are mapped into the queues. The no form of this command deletes the queue.
Default	none
Parameters	<i>queue-id</i> — Specifies the id of the queue. Specifies the <i>queue-id</i> associated with a FC. Values 1 — 8 <i>create</i> — Mandatory keyword to create a network queue policy.

rate

Syntax	rate cir <i>cir-rate-in-kbps</i> [pir <i>pir-rate-in-kbps</i>] no rate
Context	config>qos>sap-egress>queue
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the port. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.

The `rate` command can be executed at anytime, altering the PIR and CIR rates for all queues created on the access ports.

The `no` form of this command returns all queues created with the queue-id by association with the QoS policy to the default PIR and CIR parameters (max, 0).

Default `rate 0 pir max`— The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the `pir-rate` value.

Parameters `cir cir-rate-in-kbps` — The `cir` parameter overrides the default administrative CIR used by the queue. If the `rate` command is not executed or the `cir` parameter is not explicitly specified, the default CIR value is used.

Default 0

Values 0 — 10000000, max

`pir pir-rate-in-kbps` — Defines the administrative PIR rate, in kilobits, for the queue. When the `rate` command is executed, a PIR setting is optional. If the `rate` command is not executed, the default PIR of maximum value is used.

Default max

Values 1 — 10000000, max

scope

Syntax `scope {exclusive | template}`
`no scope`

Context `config>qos>sap-egress`

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command configures the scope as exclusive or template. The policy's scope cannot be changed if the policy is applied to an interface multiple ports.

The `no` form of this command sets the scope of the policy to the default of **template**.

Default template

Commands Description

- Parameters**
- exclusive** — When the scope of a policy is defined as exclusive, the policy can only be applied to one interfaceport. If a policy with an exclusive scope is assigned to a second interface an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface.
The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to 1default access-egress policy (policy-id 1).
 - template** — When the scope of a policy is defined as template, the policy can be applied to multiple interfaceports on the router.
Default QoS policies are configured with template scope. An error is generated if you try to modify the scope parameter from **template** to exclusive **scope** on default policies.

Operational Commands

copy

- Syntax** **copy sap-egress** *<src-pol>* *<dst-pol>* [**overwrite**]
- Context** config>qos
- Description** **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.
This command copies the existing SAP egress QoS policy entries to another SAP egress QoS policy.
The copy command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the overwrite keyword.
If the destination policy already exists, the key word overwrite must be specified.
- Default** template
- Parameters** *src-pol* — Specifies the source policy.
Values 1—65535
dst-pol — Specifies the destination policy.
Values 1—65535
overwrite — The information in the destination policy is overwritten by the information in the source policy.

Show Commands

sap-egress

Syntax	sap-egress [<i>policy-id</i>] [association detail]
Context	show>qos
Description	Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. This command displays sap egress QoS policy information.
Parameters	<i>policy-id</i> — Displays the policy id of the sap-egress policy. association — Displays associations related to the specified sap-egress policy. detail — Displays detailed policy information including the policy associations.

Output **SAP Egress Output**

The following table describes SAP egress show command output.

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Remark	True — Remarking is enabled for all the Dot1q-tagged packets that egress the ports on which the sap-egress QoS policy is applied and remarking is enabled. False — Remarking is disabled for the policy.
Remark Pol Id	Displays the policy id of the remarking policy.
Accounting	Specifies whether the accounting mode is packet-based or frame-based.
Scope	Exclusive — Implies that this policy can be applied only to a single access egress port. Template — Implies that this policy can be applied to multiple access ports on the router.
Description	Template — Implies that this policy can be applied to multiple access ports on the router. A text string that helps identify the policy's context in the configuration file
Queue Rates and Rules	

Commands Description

In This Section

This section provides information about the scheduler support available in the 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices.

Topics in this section include:

- [Overview](#)

Overview

Schedulers are supported on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, with strict Priority and WFQ mode of scheduling or a mix of both, available for use.

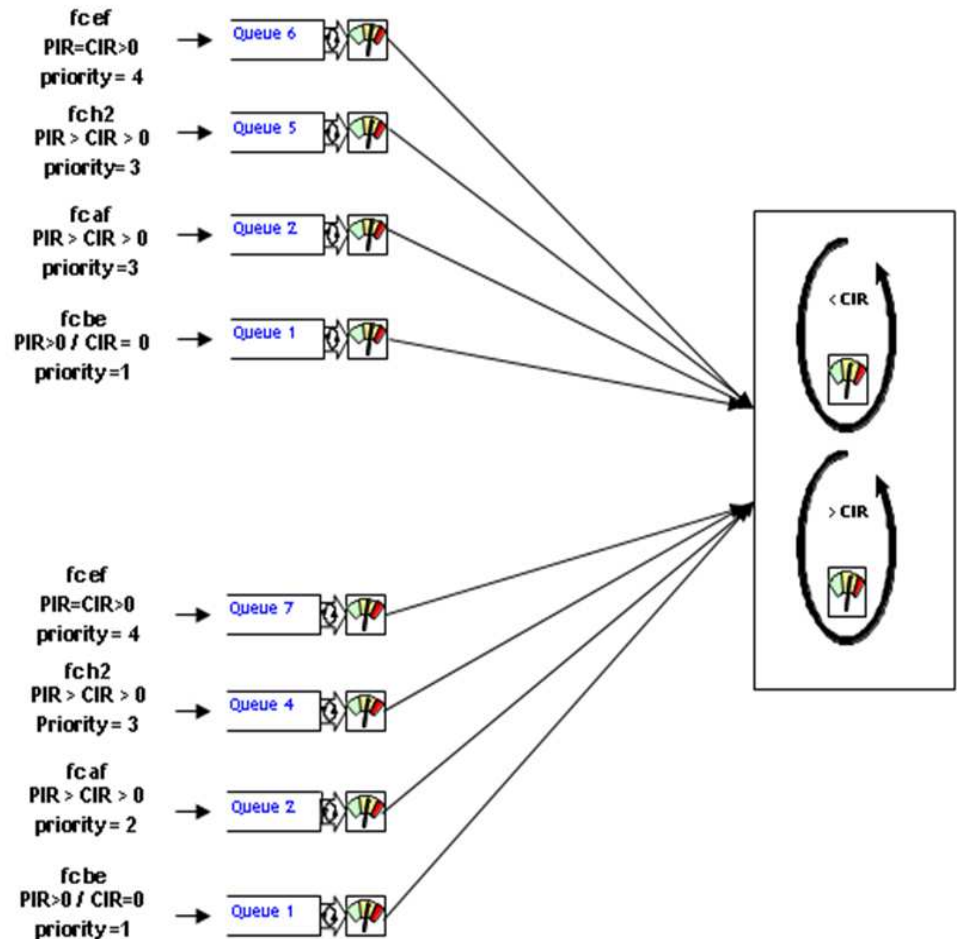
On 7210 SAS-K 2F2T1C, schedulers are used at SAP ingress, SAP egress, Access Uplink Port ingress and Access-uplink port egress.

On 7210 SAS-K 2F4T6C, schedulers are used at SAP ingress, SAP egress, Network Port ingress, and Network port egress, Access Uplink Port ingress, and Access-uplink port egress.

The scheduler uses 2 loops - the CIR loop and PIR loop, each with 4 priorities. The configured priority of the queue determines the service order of the queue in the CIR loop and the PIR loop. The scheduler first goes through the CIR loop, where it services all the queues which are operating at less than CIR rate according to the priority (that is, higher priority queues get services earlier than lower priority queues). It thereafter goes through the PIR loop, where it services all the queues which are operating above the CIR rate (but less than PIR rate) according to the priority (that is, higher priority queues get services earlier than lower priority queues). If there are multiple queues configured with the same priority, in the CIR loop the queues are scheduled using WFQ, with the configured weight of the queue used to determine the proportion of the available bandwidth that is given to the queue. In the PIR loop, the queues are scheduled using WFQ, with the configured weight of the queue used to determine the proportion of the available bandwidth that is given to the queue (using WFQ). In the PIR loop, the queues are scheduler until the PIR rate is met or until no more bandwidth is available. If the PIR rate is met, then the queues are not scheduled anymore.

The diagram below shows the scheduler implemented in 7210 SAS-K devices (both 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C)

Figure 4: Scheduler in 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C



the queues at the top belong to SAP #1 and the queues at the bottom belong to SAP #2 and that all queues have the same weight. Then the scheduling order is (assuming unlimited bandwidth is):

- Start CIR loop
- Select the highest priority queues (priority of 4), that is, SAP#1 -> Queue#6, SAP #2 -> Queue#7 and schedule them until CIR is met.
- Select the next highest priority queues (priority of 3), SAP #1-> Queue#5, SAP #1 -> Queue#2, SAP #2 -> Queue#4 and schedule them until CIR is met; Each queue is given a bandwidth in proportion to the configured weight.

Schedulers on 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C

- Select the next highest priority queues (priority of 2), SAP #2 -> Queue#2 and schedule it until CIR is met.
- No more queues left to be serviced in the CIR loop; All queues CIR is met.



Note: SAP#1-> Queue#1 and SAP#2->Queue#1 are not serviced in CIR loop as they are configured with CIR=0; They are serviced only in the PIR loop.

- Start PIR loop.
- Select the highest priority queues (priority of 4), that is, SAP#1 -> Queue#6, SAP #2 -> Queue#7. These do not have any PIR configured, therefore the user can skip them;
- Select the next highest priority queues (priority of 3), SAP #1-> Queue#5, SAP #1 -> Queue#2, SAP #2 -> Queue#4 and schedule them until their PIR is met; Each queue is given a bandwidth in proportion to the configured weight until the PIR is met.
- Select the next highest priority queues (priority of 2), SAP #2 -> Queue#2 and schedule it until PIR is met.
- Select the next highest priority queues (priority of 1), SAP#1-> Queue#1 and SAP#2->Queue#1 and schedule them until the PIR is met; Each queue is given bandwidth in proportion to the configured weight until their PIR is met.

Additionally, in this section we did not take into account the port egress rate shaper or SAP aggregate shaper. These shapers limit the available bandwidth to a port or a SAP. This does not change the scheduling mechanism but provides additional controls to the user to limit the amount of bandwidth a SAP can get or a port can transmit.

On 7210 SAS-K 2F2T1C, schedulers are available at:

- Service ingress - This scheduler distributes the available bandwidth among all the SAPs with service ingress policies.
- Service egress - This scheduler distributes the available bandwidth among all the SAPs configured on a given port with service egress policies.
- Access-uplink Port egress - This scheduler distributes the available bandwidth among up to 8 egress queues configured on the access-uplink port.
- Access-uplink Port ingress - This scheduler distributes the available bandwidth among up to 8 ingress queues configured on the access-uplink port.

On 7210 SAS-K 2F4T6C, schedulers are available at:

- Service ingress - This scheduler distributes the available bandwidth among all the SAPs with service ingress policies.
- Service egress - This scheduler distributes the available bandwidth among all the SAPs configured on a given port with service egress policies.

Overview

- Network Port & Access-uplink Port egress - This scheduler distributes the available bandwidth among up to 8 egress queues configured on the network/access-uplink port.
- Network Port and Access-uplink Port ingress - This scheduler distributes the available bandwidth among up to 8 ingress queues configured on the network/access-uplink port.

Slope QoS Policies

In This Section

This section provides information to configure slope QoS policies using the command line interface.

Topics in this section include:

- [Overview of Buffer pools and Slope policies](#)
- [Basic Configurations](#)
- [Applying Slope Policies](#)
- [Deleting QoS Policies](#)
- [Copying and Overwriting QoS Policies](#)
- [Editing QoS Policies](#)

Overview of Buffer pools and Slope policies

For an overview of buffer pools supported on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, see [Buffer Pools](#).

On 7210 SAS-K 2F2T1C, slope policies are applied to service ingress queues, service egress queues, access uplink port ingress and access uplink port egress queues.

On 7210 SAS-K 2F4T6C, Slope policies are applied to service ingress queues, service egress queues, network port ingress queues, network port egress queues, access uplink port ingress queues and access uplink port egress queues.

Each of these queuing points allocates buffers from the buffer pool and implements WRED for congestion management. During congestion WRED is used to evaluate how buffers from the pool are allocated to different FCs and to in-profile and out-of-profile traffic within a given FC. The slope policies define the WRED parameters to use for in-profile/high-priority packets and for out-of-profile/low-priority packets. The high-slope and low-slope define the parameters for in-profile/high-priority packets and for out-of-profile/low-priority packets respectively. In addition, on ring ports option is available to use separate slopes for ring traffic

(that is, traffic coming in on one ring port and going out of another ring port) and non-ring traffic (that is, traffic coming in on access port and going out of another access port or access-uplink port). For more information about ring and non-ring ports refer to the section in [Buffer Pools](#).

Basic Configurations

A basic slope QoS policy must conform to the following:

- Each slope policy must have a unique policy ID.
- High slope, low slope and non-TCP slope are shut down (default).
- Default values can be modified but parameters cannot be deleted.

Create a Slope QoS Policy

Configuring and applying slope policies is optional. If no slope policy is explicitly applied to a port, a default slope policy is applied.

To create a new slope policy for 7210 SAS-K devices, define the following:

- A slope policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- The high slope for the high priority Random Early Detection (RED) slope graph.
- The low slope for the low priority Random Early Detection (RED) slope graph.

For 7210 SAS-K devices, use the following CLI syntax to configure a slope policy.

```
CLI Syntax:  slope-policy "default" create
                description "Default slope policy."
                high-slope
                  shutdown
                  start-avg percent
                  max-avg percent
                  max-prob percent
                exit
                low-slope
                  shutdown
                  start-avg percent
                  max-avg percent
                  max-prob percent
                exit
```

```

high-slope-ring
  shutdown
  start-avg percent
  max-avg percent
  max-prob percent
exit
low-slope-ring
  shutdown
  start-avg percent
  max-avg percent
  max-prob percent
exit

```

The following displays the slope policy configuration for 7210 SAS-K devices:

```

*A:SAH01-051>config>qos>slope-policy$ info detail
-----
no description
high-slope
  shutdown
  start-avg 70
  max-avg 90
  max-prob 80
exit
low-slope
  shutdown
  start-avg 50
  max-avg 75
  max-prob 80
exit
high-slope-ring
  shutdown
  start-avg 70
  max-avg 90
  max-prob 80
exit
low-slope-ring
  shutdown
  start-avg 50
  max-avg 75
  max-prob 80
exit
-----
*A:SAH01-051>config>qos>slope-policy$

```

Applying Slope Policies

Applying Slope Policies on 7210 SAS-K 2F2T1C and 7210

SAS-K 2F4T6C

- On 7210 SAS-K 2F2T1C, slope policies are associated with service ingress queues, service egress queues, access-uplink port egress queues, and access-uplink port ingress queues.
- On 7210 SAS-K 2F4T6C, slope policies are associated with service ingress queues, service egress queues, network port egress queues, network port ingress queues, access-uplink port egress queues and access-uplink port ingress queues.

The following CLI syntax examples may be used to apply slope policies to service ingress queues, service egress queues, access-uplink port ingress queues and access-uplink port egress queues on 7210 SAS-K 2F2T1C:

CLI Syntax:

```
config> qos> sap-ingress> queue <id> slope-policy name
config> qos> sap-egress> queue <id> slope-policy name
config> qos> network> ingress> queue <id> slope-policy
name
config> qos> network-queue> queue <id> slope-policy name
```

The following CLI syntax examples may be used to apply slope policies to service ingress queues, service egress queues, network port egress queues, network port ingress queues, access-uplink port ingress queues and access-uplink port egress queues on 7210 SAS-K 2F4T6C:

CLI Syntax:

```
config> qos> sap-ingress> queue <id> slope-policy name
config> qos> sap-egress> queue <id> slope-policy name
config> qos> network> ingress> queue <id> slope-policy
name
config> qos> network-queue> queue <id> slope-policy name
```

Default Slope Policy Values for 7210 SAS-K devices

Table 27: Default Slope Policy for 7210 SAS-K

Description	Default Slope Policy
high-slope	
Administrative state	shutdown
start-avg	70% utilization
max-avg	90% utilization

Table 27: Default Slope Policy for 7210 SAS-K

Description	Default Slope Policy
max-prob	80%
low slope	
Administrative state	shutdown
start-threshold	50% utilization
max-avg	75% utilization
max-prob	80%
high-slope-ring	
Administrative state	shutdown
start-threshold	70% utilization
max-avg	90% utilization
max-prob	80%
low-slope-ring	
Administrative state	shutdown
start-threshold	50 % utilization
max-avg	75% utilization
max-prob	80%

The following output displays the default configuration:

```
ALA7>config>qos>slope-policy# info detail
-----
description "Default slope policy."
high-slope
  shutdown
  start-threshold 75
  queue 1 drop-rate 1
  queue 2 drop-rate 1
  queue 3 drop-rate 1
  queue 4 drop-rate 1
  queue 5 drop-rate 1
  queue 6 drop-rate 1
  queue 7 drop-rate 1
  queue 8 drop-rate 1
exit
low-slope
  shutdown
  start-threshold 50
```

Deleting QoS Policies

```
queue 1 drop-rate 0
queue 2 drop-rate 0
queue 3 drop-rate 0
queue 4 drop-rate 0
queue 5 drop-rate 0
queue 6 drop-rate 0
queue 7 drop-rate 0
queue 8 drop-rate 0
exit
-----
ALA7>config>qos>slope-policy#
```

Deleting QoS Policies

A slope policy is associated by default with access and access uplink egress pools. A default policy may be replaced with a non-default policy, but a policy cannot be entirely removed from the configuration. When a non-default policy is removed, the policy association reverts to the default slope policy *policy-id* default. A QoS policy cannot be deleted until it is removed from all ports where it is applied or if the policies are using the slope-policy.

```
ALA-7>config>qos# no slope-policy slopePolicy1
MINOR: QOS #1902 Slope policy has references
ALA-7>config>qos#
```

The following CLI syntax examples can be used to remove slope policies from ports on 7210 SAS-K devices:

```
CLI Syntax: config> qos> sap-ingress> queue <id> no slope-policy
               name
config> qos> sap-egress> queue <id> no slope-policy
               name
config> qos> network> ingress> queue <id> no slope-policy
               name
config> qos> network-queue> queue <id> no slope-policy
               name
```

Remove a Policy from the QoS Configuration

To delete a slope policy, enter the following command:

```
CLI Syntax: config>qos# no slope-policy policy-id
```

```
Example: config>qos# no slope-policy slopePolicy1
```

Copying and Overwriting QoS Policies

You can copy an existing slope policy, rename it with a new policy ID value, or overwrite an existing policy ID. The overwrite option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos> copy {slope-policy} source-policy-id dest-policy id [overwrite]`

The following output displays the copied policies for (7210 SAS-K devices):

```
*A:SAH01-051>config>qos>slope-policy# info detail
-----
description "Default slope policy."
high-slope
  shutdown
  start-avg 70
  max-avg 90
  max-prob 80
exit
low-slope
  shutdown
  start-avg 50
  max-avg 75
  max-prob 80
exit
high-slope-ring
  shutdown
  start-avg 70
  max-avg 90
  max-prob 80
exit
low-slope-ring
  shutdown
  start-avg 50
  max-avg 75
  max-prob 80
exit
-----
*A:SAH01-051>config>qos>slope-policy#
```

Editing QoS Policies

You can change existing policies and entries in the CLI or NMS. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors copy the policy to a work area, make the edits, and then write over the original policy.

Slope QoS Policy Command Reference

Command Hierarchies

- [Configuration commands](#)
- [Operational Commands](#)
- [Show Commands](#)

Configuration commands

```

— config
  — qos
    — [no] slope-policy name
      — description description-string
      — no description
      — [no] high-slope
        — max-avg percent
        — no max-avg
        — max-prob percent
        — no max-prob
        — [no] shutdown
        — start-avg percent
        — no start-avg
      — [no] high-slope-ring
        — max-avg percent
        — no max-avg
        — max-prob percent
        — no max-prob
        — [no] shutdown
        — start-avg percent
        — no start-avg
      — [no] low-slope
        — max-avg percent
        — no max-avg
        — max-prob percent
        — no max-prob
        — [no] shutdown
        — start-avg percent
        — no start-avg
        — [no] shutdown
      — [no] low-slope-ring
        — max-avg percent
        — no max-avg
        — max-prob percent
        — no max-prob

```

Slope QoS Policy Command Reference

- [no] **shutdown**
- **start-avg** *percent*
- **no start-avg**

Operational Commands

- **config**
 - **qos**
 - **copy** *slope-policy src-name dst-name [overwrite]*

Show Commands

- **show**
 - **qos**
 - **slope-policy** [*slope-policy-name*] [*detail*]

Command Descriptions

Configuration Commands

Generic Commands

description

Syntax	<i>description-string</i> description no description
Context	config>qos>slope-policy
Description	<p>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The description command associates a text string with a configuration context to help identify the context in the configuration file.</p> <p>The no form of this command removes any description string from the context.</p>
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax	copy slope-policy <i>src-name</i> <i>dst-name</i> [overwrite]
Context	config>qos

Command Descriptions

- Description** This command copies existing QoS policy entries for a QoS policy-id to another QoS policy-id.
- The **copy** command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.
- Parameters** **slope-policy** — Indicates that the source policy ID and the destination policy ID are slope policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.
- overwrite** — Specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, an error will occur if the destination policy ID exists.
- ```
ALA-7>config>qos# copy slope-policy default sp1
MINOR: CLI Destination "sp1" exists - use {overwrite}.
ALA-7>config>qos#overwrite
```

## Slope Policy QoS Commands

### slope-policy

- Syntax** **[no] slope-policy** *name* [**create**]
- Context** config>qos
- Description** **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.
- This command enables the context to configure a QoS slope policy.
- Default** slope-policy “default”
- Parameters** *name* — The name of the slope policy.
- Values** Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## Slope Policy QoS Policy Commands

### queue

- Syntax** **queue** *queue-id*
- Context** config>qos>slope-policy

|                    |                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.<br><br>This command sets the context to configure the high-priority, low-priority, and non-tcp slope parameters per queue. |
| <b>Parameters</b>  | <i>queue-id</i> — Specifies the ID of the queue for which the drop-rate is to be configured.<br><br><b>Values</b> 1 — 8                                                                         |

## high-slope

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] high-slope</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>slope-policy>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.<br><br>The <b>high-slope</b> context contains the commands and parameters for defining the high priority Random Early Detection (RED) slope graph. A high priority RED slope is used managing access to the shared portion of the buffer pool for high priority or in-profile packets that ingress of non-ring ports and egress out of non-ring ports.<br><br>The <b>high-slope</b> parameters can be changed at any time and the affected buffer pool high priority RED slopes will be adjusted appropriately.<br><br>The <b>no</b> form of this command restores the high slope configuration commands to the default values. If the commands within high-slope are set to the default parameters, the high-slope node will not appear in save config and show config output unless the detail parameter is present. |

## low-slope

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] low-slope</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>slope-policy>queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.<br><br>The <b>low-slope</b> context contains the commands and parameters for defining the low priority Random Early Detection (RED) slope graph. A low priority ring RED slope is used managing access to the shared portion of the buffer pool for low priority or out-of-profile packets that are received on non-ring ports and egress out of non-ring ports.<br><br>The <b>low-slope</b> parameters can be changed at any time and the affected buffer pool high priority RED slopes will be adjusted appropriately.<br><br>The <b>no</b> form of this command restores the high slope configuration commands to the default values. If the commands within low-slope are set to the default parameters, the low-slope node will not appear in save config and show config output unless the detail parameter is present. |

## Command Descriptions

### high-slope-ring

|                    |                                                                      |
|--------------------|----------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] high-slope-ring</b>                                          |
| <b>Context</b>     | config>qos>slope-policy                                              |
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. |

The high-slope-ring context contains the commands and parameters for defining the high priority Random Early Detection (RED) slope graph. A high priority ring RED slope is used managing access to the shared portion of the buffer pool for high priority or in-profile packets that ingress of ring ports and egress out of ring ports.

The high-slope-ring parameters can be changed at any time and the affected buffer pool high priority RED slopes will be adjusted appropriately.

The no form of this command restores the high slope configuration commands to the default values. If the commands within high-slope-ring are set to the default parameters, the high-slope node will not appear in save config and show config output unless the detail parameter is present.

### low-slope-ring

|                    |                                                                      |
|--------------------|----------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] low-slope-ring</b>                                           |
| <b>Context</b>     | config>qos>slope-policy                                              |
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. |

The low-slope-ring context contains the commands and parameters for defining the low priority Random Early Detection (RED) slope graph. A low priority ring RED slope is used managing access to the shared portion of the buffer pool for low priority or out-of-profile packets that are received on ring ports and egress out of ring ports.

The low-slope-ring parameters can be changed at any time and the affected buffer pool high priority RED slopes will be adjusted appropriately.

The no form of this command restores the high slope configuration commands to the default values. If the commands within low-slope-ring are set to the default parameters, the low-slope node will not appear in save config and show config output unless the detail parameter is present.

## Slope Policy QoS Policy Commands (for SAS-K devices)

### RED Slope Commands

#### max-avg

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>max-avg</b> <i>percent</i><br><b>no max-avg</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Context</b>     | config>qos>slope-policy>queue>high-slope<br>config>qos>slope-policy>queue>low-slope<br>config>qos>slope-policy>queue>high-slope-ring<br>config>qos>slope-policy>queue>low-slope-ring                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p><b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.</p> <p>Sets the low priority or high priority Weighted Random Early Detection (WRED) slope position for the reserved and shared buffer average utilization value where the packet discard probability rises directly to one. The percent parameter is expressed as a percentage of the shared buffer size.</p> <p>The <b>no</b> form of this command restores the max-avg value to the default setting. If the current startavg setting is larger than the default, an error will occur and the max-avg setting will not be changed to the default.</p> |
| <b>Default</b>     | <p><b>max-avg 90</b> - High slope default is 90% buffer utilization before discard probability is 1.</p> <p><b>max-avg 75</b> - Low slope default is 75% buffer utilization before discard probability is 1.</p> <p><b>max-avg 75</b> - Non-tcp slope default is 75% buffer utilization before discard probability is 1.</p>                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <p><i>percent</i> — The percentage of the reserved and shared buffer space for the buffer pool at which point the drop probability becomes 1. The value entered must be greater or equal to the current setting of startavg. If the entered value is smaller than the current value of start-avg, an error will occur and no change will take place.</p> <p><b>Values</b>     0 — 100</p>                                                                                                                                                                                                                                       |

#### max-prob

|                |                                                                                                                                                                                      |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>max-prob</b> <i>percent</i><br><b>no max-prob</b>                                                                                                                                 |
| <b>Context</b> | config>qos>slope-policy>queue>high-slope<br>config>qos>slope-policy>queue>low-slope<br>config>qos>slope-policy>queue>high-slope-ring<br>config>qos>slope-policy>queue>low-slope-ring |

## Command Descriptions

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.<br><br>Sets the low priority or high priority Random Early Detection (RED) slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one. The percent parameter is expressed as a percentage of packet discard probability where always discard is a probability of 1. A <b>max-prob</b> value of 75 represents 75% of 1, or a packet discard probability of 0.75.<br><br>The <b>no</b> form of this command restores the <b>max-prob</b> value to the default setting. |
| <b>Default</b>     | <b>max-prob 75</b> - 75% maximum drop probability corresponding to the <b>max-avg</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <i>percent</i> — The maximum drop probability percentage corresponding to the max-avg, expressed as a decimal integer.<br><br><b>Values</b> 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 25, 50, 75, 100                                                                                                                                                                                                                                                                                                                                                                                                                      |

## shutdown

|                    |                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>qos>slope-policy>high-slope<br>config>qos>slope-policy>low-slope<br>config>qos>slope-policy>queue>high-slope-ring<br>config>qos>slope-policy>queue>low-slope-ring                                                                                                                                                                                           |
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.<br><br>This command enables or disables the administrative status of the Random Early Detection slope.<br><br>By default, all slopes are shutdown and have to be explicitly enabled ( <b>no shutdown</b> ).<br><br>The <b>no</b> form of this command administratively enables the RED slope. |
| <b>Default</b>     | <b>shutdown</b> - RED slope disabled implying a zero (0) drop probability.                                                                                                                                                                                                                                                                                         |

## start-avg

|                    |                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>start-avg percent</b><br><b>no start-avg</b>                                                                                                                                      |
| <b>Context</b>     | config>qos>slope-policy>queue>high-slope<br>config>qos>slope-policy>queue>low-slope<br>config>qos>slope-policy>queue>high-slope-ring<br>config>qos>slope-policy>queue>low-slope-ring |
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.                                                                                                                 |

This command sets the low priority or high priority Random Early Detection (RED) slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero. The percent parameter is expressed as a percentage of the shared buffer size.

The **no** form of this command restores the start-avg value to the default setting. If the max-avg setting is smaller than the default, an error will occur and the start-avg setting will not be changed to the default.

- Default**    **max-avg 70** - High slope default is 70% buffer utilization.  
**max-avg 50** - Low slope default is 50% buffer utilization.  
**max-avg 50** - Non-tcp slope default is 50% buffer utilization.
- Parameters**    *percent* — The percentage of the reserved and shared buffer space for the buffer pool at which the drop starts. The value entered must be lesser or equal to the current setting of max-avg. If the entered value is greater than the current value of max-avg, an error will occur and no change will take place.
- Values**        0 — 100

## Show Commands

### slope-policy

- Syntax**        **slope-policy** [*slope-policy-name*] [**detail**]
- Context**        show>qos
- Description**    **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C
- This command displays slope policy information.
- Parameters**    *slope-policy-name* — The name of the slope policy.  
**detail** — Displays detailed information about the slope policy.
- Output**         The following table describes slope QoS policy output fields.

**Table 28: Show QoS Slope Policy Output Fields**

| Label       | Description                                                                                                        |
|-------------|--------------------------------------------------------------------------------------------------------------------|
| Policy      | The ID that uniquely identifies the policy.                                                                        |
| Description | A string that identifies the policy's context in the configuration file.                                           |
| Time Avg    | The weighting between the previous shared buffer average utilization result and the new shared buffer utilization. |

**Table 28: Show QoS Slope Policy Output Fields (Continued)**

| Label            | Description                                                                                                                                                                                                                                                                                                             |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Slope Parameters |                                                                                                                                                                                                                                                                                                                         |
| Start Avg        | Specifies the low priority or high priority RED slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero.                                                                                                                                      |
| Max Avg          | Specifies the percentage of the shared buffer space for the buffer pool at which point the drop probability becomes 1, expressed as a decimal integer                                                                                                                                                                   |
| Admin State      | Up<br>The administrative status of the RED slope is enabled.<br>Down<br>The administrative status of the RED slope is disabled.<br>Specifies the low priority or high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one. |
| Max Prob.        | Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one.                                                                                                                                                    |

**Sample for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C**

```
*A:SAH01-051>show>qos# slope-policy "default"
```

```
=====
QoS Slope Policy
=====
Policy : default
Description : Default slope policy.

High Slope Parameters

Start Avg : 70
Max Avg : 90
Admin State : Disabled
Max Prob. : 80

Low Slope Parameters

Start Avg : 50
Max Avg : 75
Admin State : Disabled
Max Prob. : 80

High Slope Ring Parameters

Start Avg : 70
Max Avg : 90
Admin State : Disabled
Max Prob. : 80
```

-----  
Low Slope Ring Parameters-----  
Start Avg : 50 Admin State : Disabled  
Max Avg : 75 Max Prob. : 80=====  
\*A:SAH01-051>show>qos#**Sample for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C**

\*A:SAH01-051&gt;show&gt;qos# slope-policy detail

-----  
QoS Slope Policy=====  
Policy : default  
Description : Default slope policy.-----  
High Slope Parameters-----  
Start Avg : 70 Admin State : Disabled  
Max Avg : 90 Max Prob. : 80-----  
Low Slope Parameters-----  
Start Avg : 50 Admin State : Disabled  
Max Avg : 75 Max Prob. : 80-----  
High Slope Ring Parameters-----  
Start Avg : 70 Admin State : Disabled  
Max Avg : 90 Max Prob. : 80-----  
Low Slope Ring Parameters-----  
Start Avg : 50 Admin State : Disabled  
Max Avg : 75 Max Prob. : 80-----  
SAP Ingress-----  
SAP Ingress Policy Id : 1  
Queue Ids : 1-----  
SAP Egress-----  
SAP Egress Policy Id : 1  
Queue Ids : 1-----  
Network Ingress-----  
Network Ingress Policy Id : 1  
Queue Ids : 1, 2, 3, 4, 5, 6, 7, 8-----  
Network Queues

## Command Descriptions

```
Network Queue Policy Name : default
Queue Ids : 1, 2, 3, 4, 5, 6, 7, 8

=====
*A:SAH01-051>show>qos#
```

# Remark Policies for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

---

## In This Section

This section provides information to configure remark policies using the command line interface. This section is applicable only to 7210 SAS-K devices.

Topics in this section include:

- [Overview](#)
- [Basic Configurations on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C](#)

## Overview

### Overview of remark policies for 7210 SAS-K 2F2T1C

The remark policies are used to configure the marking behavior for the system at the egress of access SAP or access uplink port. These policies allow the user to define the forwarding class to egress marking values and allow them to use the available hardware resources efficiently. Based on the packet encapsulation used, the remark policy allows the user to define and associate service egress policies to service egress, and network QoS policies to access-uplink port. The following types of remark policies are available

- dot1p - Used for service egress policy, and network qos.
- dscp - Used for service egress and network qos .
- dot1p-dscp used for service egress and network qos.

The type of the remark policy identifies the bits marked in the packet header. Each of these remark policy types can be associated with only appropriate QoS policies and service entities as listed in [Table 29](#).

**Table 29: Summary of remark policy and attachment points for 7210 SAS-K 2F2T1C**

| <b>Remark Policy Type</b> | <b>Qos Policy</b>     | <b>Attachment Point</b> | <b>Packet Header Bits Marked</b>                                                                                                                                                                                                             |
|---------------------------|-----------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dot1p                     | Service Egress Policy | Access SAP egress       | Dot1p bits and optionally the DEI bits, in the L2 header for service packets sent out of an Access SAP.                                                                                                                                      |
|                           | Network policy        | Access-uplink Port      | Dot1p bits and optionally the DEI bits, in the L2 header for service packets sent out of an access-uplink port (all packets sent out of all access-uplink SAPs are marked).                                                                  |
| dscp                      | service egress policy | Access SAP egress       | IP DSCP bits in the IP header (if present) for service packets sent out of an Access SAP.                                                                                                                                                    |
|                           | Network policy        | Access-uplink Port      | IP DSCP bits in the IP header (if present) for service packets sent out of an Access-uplink port Marking is done for packets sent out of all the access-uplink SAPs configured on the access-uplink port.                                    |
| dot1p-dscp                | service egress policy | Access SAP egress       | Dot1p bits and optionally the DEI bits, in the L2 header for service packets sent out of access SAP. In addition to the dot1p bits, the IP DSCP bits in the IP header (if present) are marked for service packets sent out of an Access SAP. |

Table 29: Summary of remark policy and attachment points for 7210 SAS-K 2F2T1C

| Remark Policy Type | Qos Policy     | Attachment Point   | Packet Header Bits Marked                                                                                                                                                                                                                                                                                                                                             |
|--------------------|----------------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | Network policy | Access-uplink Port | Dot1p bits and optionally the DEI bits, in the L2 header for service packets sent out of access uplink port. In addition to the dot1p bits, the IP DSCP bits in the IP header (if present) are marked for service packets sent out of an access-uplink port. Marking is done for packets sent out of all the access-uplink SAPs configured on the access-uplink port. |

## Configuration Guideline for 7210 SAS-K 2F2T1C

- The 7210 SAS-K 2F2T1C marks the Dot1p configured in the marking policy, only if the node adds ones or two tags or if it replaces existing one or two tags. It does not touch the VLAN tag which represents customer payload (for example, 7210 SAS-K2F2T1C does not implement marking for packets when forwarding traffic out of dot1q range SAP or a NULL SAP).
- On 7210 SAS-K2F2T1C a limited number of unique remark policies are available to be shared among all the different types of remark policies. In other words, dot1p, dscp and dot1p-dscp remark policy types share the available resources and scaling one type of remark policy reduces the amount policies allowed for other type of remark policy.
- On 7210 SAS-K2F2T1C, user also has an option to preserve the Dot1p values in the received packet. See the CLI description for `configure> service> sap> egress> dot1p-inner` and `configure> service> sap> egress> dot1p-outer` in the 7210 SAS D,E,K Services User Guide.

## Overview of remark policies for 7210 SAS-K 2F4T6C

The remark policies are used to configure the marking behavior for the system at the egress of access SAP, network port and access uplink port. These policies allow the user to define the forwarding class to egress marking values and allow them to use the available hardware resources efficiently. Based on the packet encapsulation used, the remark policy allows the user to define and associate service egress policies to access SAP egress, and network QoS policies applied to both network port and access-uplink port. The following type of remark policies are available:

- dot1p - Used for service egress policy, and network qos.
- dscp - Used for service egress and network QoS.
- dot1p-dscp - Used for service egress and network qos.
- lsp-exp - Used for network QoS (only network ports)
- dot1p-lsp-exp - Used for network QoS (only network ports)
- dot1p-lsp-exp-dscp - Used for network QoS (only network ports)
- dscp-lsp-exp - Used for network QoS (only network ports)

The type of the remark policy identifies the bits marked in the packet header. Each of these remark policy types can be associated with only appropriate QoS policies and service entities as listed in [Table 30](#).

**Table 30: Summary of remark policy and attachment points for 7210 SAS-K 2F4T6C**

| <b>Remark Policy Type</b> | <b>Qos Policy</b>     | <b>Attachment Point</b> | <b>Packet Header Bits Marked</b>                                                                                                                                                                          |
|---------------------------|-----------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dot1p                     | Service Egress Policy | Access SAP egress       | Dot1p bits and optionally the DEI bits, in the outermost Ethernet header for service packets sent out of an Access SAP.                                                                                   |
|                           | Network policy        | Network Port            | Dot1p bits and optionally the DEI bits, in the outermost Ethernet header for all MPLS & IP packets sent out of a network port.                                                                            |
|                           | Network policy        | Access-uplink port      | Dot1p bits and optionally the DEI bits, in the Ethernet header for service packets sent out of an access-uplink port (all packets sent out of all access-uplink SAPs are marked).                         |
| dscp                      | service egress policy | Access SAP egress       | IP DSCP bits in the IP header (if present) for service packets sent out of an Access SAP.                                                                                                                 |
|                           | Network policy        | Network port            | IP DSCP bits in the IP header (if present) for all IP packets sent out of a network port.                                                                                                                 |
|                           | Network policy        | Access-uplink Port      | IP DSCP bits in the IP header (if present) for service packets sent out of an Access-uplink port Marking is done for packets sent out of all the access-uplink SAPs configured on the access-uplink port. |

**Table 30: Summary of remark policy and attachment points for 7210 SAS-K 2F4T6C**

| <b>Remark Policy Type</b> | <b>Qos Policy</b>     | <b>Attachment Point</b> | <b>Packet Header Bits Marked</b>                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|-----------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dot1p-dscp                | service egress policy | Access SAP egress       | Dot1p bits and optionally the DEI bits, in the L2 header for service packets sent out of access SAP. In addition to the dot1p bits, the IP DSCP bits in the IP header (if present) are marked for service packets sent out of an Access SAP.                                                                                                                                |
|                           | Network policy        | Network port            | Dot1p bits and optionally the DEI bits, in the Ethernet header for service packets sent out of access SAP. In addition to the dot1p bits, the IP DSCP bits in the IP header (if present) are marked for service packets sent out of an Access SAP.                                                                                                                          |
|                           | Network policy        | Access-uplink Port      | Dot1p bits and optionally the DEI bits, in the ethernet header for service packets sent out of access uplink port. In addition to the dot1p bits, the IP DSCP bits in the IP header (if present) are marked for service packets sent out of an access-uplink port. Marking is done for packets sent out of all the access-uplink SAPs configured on the access-uplink port. |
| lsp-exp                   | Network policy        | Network port            | MPLS EXP bits for the outermost MPLS labels for all MPLS packets sent out of an network port.                                                                                                                                                                                                                                                                               |

**Table 30: Summary of remark policy and attachment points for 7210 SAS-K 2F4T6C**

| Remark Policy Type | Qos Policy     | Attachment Point | Packet Header Bits Marked                                                                                                                                                                                                                                                                |
|--------------------|----------------|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dot1p-lsp-exp      | Network policy | Network port     | Dot1p bits and optionally the DEI bits, in the outermost Ethernet header for all MPLS and IP packets sent out of a network port and MPLS EXP bits for the outermost MPLS labels for all MPLS packets sent out of a network port.                                                         |
| dot1p-lsp-exp-dscp | Network policy | Network port     | Dot1p bits and optionally the DEI bits, in the outermost Ethernet header for all MPLS & IP packets sent out of a network port, IP DSCP for all IP packets sent out of the network port, and MPLS EXP bits for the outermost MPLS labels for all MPLS packets sent out of a network port. |
| dscp-lsp-exp       | Network Policy | Network port     | IP DSCP for all IP packets sent out of the network port, and MPLS EXP bits for the outermost MPLS labels for all MPLS packets sent out of a network port.                                                                                                                                |

## Configuration Guideline for 7210 SAS-K 2F4T6C

- The 7210 SAS-K 2F4T6C marks the Dot1p configured in the marking policy only if the node adds one or two tags or if it replaces existing one or two tags. It does not touch the VLAN tag which represents customer payload (for example, 7210 SAS-K 2F4T6C does not mark the packets when forwarding traffic out of dot1q range SAP or a NULL SAP).
- On 7210 SAS-K 2F4T6C a limited number of unique remark policies are available to be shared among all the different types of remark policies. In other words, dot1p, dscp, dot1p-dscp, lsp-exp, dot1p-lsp-exp, dot1p-dscp-lsp-exp remark policy types share the available resources and scaling one type of remark policy reduces the amount policies allowed for other type of remark policy.

- On 7210 SAS-K 2F4T6C, user also has an option to preserve the Dot1p values in the received packet. See the CLI description for `configure> service> sap> egress> dot1p-inner` and `configure> service> sap> egress> dot1p-outer` in the 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C Services User Guide.
- On 7210 SAS-K 2F4T6C, when users enable marking without explicitly associating the desired remark policy, then the default remark policy is used for marking packets.
- On 7210 SAS-K 2F4T6C, any remark-policy that specifies MPLS lsp-exp values cannot be associated with access-uplink port. It is blocked by software.
- On 7210 SAS-K 2F4T6C, MPLS packets sent out of network port can be marked with MPLS lsp-exp values and Dot1p, DEI values if both are configured. User has an option to configure either one of them too.
- On 7210 SAS-K 2F4T6C, IP packets sent out of network port can be marked with IP DSCP values and Dot1p, DEI values if both are configured. User has an option to configure either one of them too.
- On 7210 SAS-K 2F4T6C, the dot1p bits of the packets sent out of network ports are always marked by default even when remarking is disabled.

## Basic Configurations on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

A basic remark policy must confirm to the following:

- Each remark policy must have a unique policy ID.
- The remark policy type must be specified.
- The forwarding class to egress marking values must be specified.

## Creating a Remark Policy

To create a new remark policy, define the following:

- A remark policy name and type is optional and by default it is 'dot1p'.
- Provide a brief description of the policy features.
- Specify the forwarding class to egress marking values.

Use the following CLI syntax to configure a remark policy:

```

*A:7210SAS>config>qos>remark# info detail (applicable to SAS-K devices)

```

## Remark Policies for 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C

```
no description
 fc af
 de-mark-outer
 dot1p-inner 4
 dot1p-outer in-profile 2 out-profile 3
 exit
 fc be
 de-mark-outer force 0
 dot1p-inner in-profile 4 out-profile 2
 dot1p-outer 4
 exit
 fc ef
 no de-mark-outer
 dot1p-inner in-profile 4 out-profile 5
 dot1p-outer in-profile 3 out-profile 1
 exit
 fc h1
 de-mark-outer
 dot1p-inner 1
 dot1p-outer in-profile 1 out-profile 2
 exit
 fc h2
 de-mark-outer
 dot1p-inner in-profile 4 out-profile 4
 dot1p-outer in-profile 6 out-profile 3
 exit
 fc l1
 de-mark-outer force 1
 dot1p-inner in-profile 2 out-profile 6
 dot1p-outer 4
 exit
 fc l2
 de-mark-outer
 dot1p-inner 7
 dot1p-outer 3
 exit
 fc nc
 no de-mark-outer
 dot1p-inner in-profile 2 out-profile 4
 dot1p-outer in-profile 5 out-profile 6
 exit
*A:7210SAS>config>qos>remark#
```

## Editing QoS Policies

Existing policies and entries can be edited through the CLI or NMS. The changes are applied immediately to all services where the policy is applicable.

To prevent configuration errors perform the following:

1. Copy the policy to a work area
2. Edit the policy

## Overview

- 
- 
3. Over write the original policy

## Remark Policy Command Reference

### Command Hierarchies

- [Configuration Commands for 7210 SAS-K 2F2T1C](#)  
[Configuration Commands for 7210 SAS-K 2F4T6C](#)
- [Show Commands](#)
- [Operational Commands](#)

### Configuration Commands for 7210 SAS-K 2F2T1C

```

config
 — qos
 — no remark policy-id
 — remark policy-id [create] [remark-type {dot1p | dscp | dot1p-dscp}]
 — [no] description description-string
 — [no] fc fc-name
 — [no] de-mark-outer [force de-value]
 — dot1p-inner dot1p-value
 — dot1p-inner [in-profile dot1p-value] [out-profile dot1p-value]
 — no dot1p-inner
 — no dot1p-outer
 — dot1p-outer dot1p-value
 — dot1p-outer [in-profile dot1p-value] [out-profile dot1p-value]
 — dscp-in-profile dscp-name
 — no dscp-in-profile
 — dscp-out-profile dscp-name
 — no dscp-out-profile

```

### Configuration Commands for 7210 SAS-K 2F4T6C

```

config
 — qos
 — no remark policy-id
 — remark policy-id [create] [remark-type {dot1p | dscp | dot1p-dscp | lsp-exp | dot1p-lsp-exp |
 dot1p-lsp-exp-dscp | dscp-lsp-exp}]
 — [no] description description-string
 — [no] fc fc-name
 — [no] de-mark-outer [force de-value]
 — dot1p-inner dot1p-value
 — dot1p-inner [in-profile dot1p-value] [out-profile dot1p-value]
 — no dot1p-inner

```

## Remark Policy Command Reference

- **no dot1p-outer**
- **dot1p-outer** *dot1p-value*
- **dot1p-outer** [**in-profile** *dot1p-value*] [**out-profile** *dot1p-value*]
- **dscp-in-profile** *dscp-name*
- **no dscp-in-profile**
- **dscp-out-profile** *dscp-name*
- **no dscp-out-profile**
- **lsp-exp-in-profile** *lsp-exp-value*
- **lsp-exp-out-profile** *lsp-exp-value*

## Show Commands

- show
  - qos
  - **remark-policy** [*policy-id*] [*association*]*detail*]

## Operational Commands

- config
  - qos
  - **copy** *remark src-pol dst-pol* [**overwrite**]

## Command Descriptions

### Configuration Commands

#### Generic Commands

##### description

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] description</b> <i>description-string</i>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>qos>remark                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p><b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C</p> <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The <b>description</b> command associates a text string with a configuration context to help identify the context in the configuration file.</p> <p>The <b>no</b> form of this command removes any description string from the context.</p> |
| <b>Default</b>     | No description is associated with the configuration context.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.                                                                                                                                       |

### Operational Commands

##### copy

|                    |                                                                     |
|--------------------|---------------------------------------------------------------------|
| <b>Syntax</b>      | <b>copy remark</b> <i>src-pol dst-pol</i> [ <b>overwrite</b> ]      |
| <b>Context</b>     | config>qos                                                          |
| <b>Description</b> | <b>Platforms Supported:</b> 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C |

## Command Descriptions

This command copies existing remark policy entries to another remark policy.

The **copy** command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

If the destination policy already exists, the key word **overwrite** must be specified.

**Parameters** *src-pol* — Specifies the source policy.

**Values** 1—65535

*dst-pol* — Specifies the destination policy.

**Values** 1—65535

**overwrite** — The information in the destination policy is overwritten by the information in the source policy.

## Remark Policy QoS Commands

### remark

**Syntax** **no remark policy-id**  
**remark policy-id [create] [remark-type remarking type]**

**Context** config>qos

**Description** **Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C**

This command creates a new remark policy of the specified type.

The following types of remark policies are available on 7210 SAS-K 2F2T1C:

- dot1p
- dscp
- dot1p-dscp

The following type of remark policies are available on 7210 SAS-K 2F4T6C:

- dot1p - Used for service egress policy, and network qos.
- dscp - Used for service egress and network QoS.
- dot1p-dscp - Used for service egress and network qos.
- lsp-exp - Used for network QoS (only network ports)
- dot1p-lsp-exp - Used for network QoS (only network ports)
- dot1p-dscp-lsp-exp - Used for network QoS (only network ports)
- dscp-lsp-exp - Used for network qos (only network ports).

## Remark Policies for 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C

The 'remark-type' of the policy also determines the values user is allowed to configure in the policy and the QoS policy with which this policy can be associated with. See the [Summary of remark policy and attachment points for 7210 SAS-K 2F2T1C](#) the different remark policies supported on the node and its use.



**Note:** On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, a limited number of unique remark policies are available to be shared among all the different types of remark policies. In other words, dot1p, dscp, dot1p-dscp, lsp-exp etc. remark policy types share the available resources and scaling one type of remark policy reduces the amount policies allowed for other type of remark policy.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no default                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b> | <i>policy-id</i> — The policy ID of the remark policy.<br><b>Values</b> [1..65535]<br><i>marking-type</i> — Specifies the type of marking values in the remark policy.<br><b>Values</b> dot1p — Specify FC to 802.1 Dot1p value to use for marking. It is the default used if user does not explicitly specify the marking-type value.<br>dscp— Specify FC to IP DSCP value to use for marking.<br>dot1p-dscp - Specify FC to both Dot1p and IP DSCP values to use for marking.<br><b>Values</b> lsp-exp- Available only on 7210 SAS-K 2F4T6C. Specify FC to MPLS EXP values to use for marking.<br><b>Values</b> dot1p-lsp-exp - Available only on 7210 SAS-K 2F4T6C. Specify FC to MPLS EXP values and Dot1p values to use for marking.<br><b>Values</b> dot1p-dscp-lsp-exp - Available only on 7210 SAS-K 2F4T6C. Specify FC to MPLS EXP values, IP DSCP values and Dot1p values to use for marking. |

fc

|                    |                                                                     |
|--------------------|---------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] fc fc-name</b>                                              |
| <b>Context</b>     | config>qos>remark                                                   |
| <b>Description</b> | <b>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C</b> |

This command specifies the forwarding class name and provides the context to configure the marking value for the FC. Based on the type of remark policy created, the FC command allows the user to specify the appropriate marking values. The fc command overrides the default parameters for the forwarding class to the values defined.

## Command Descriptions

The no form of the command removes the forwarding class to marking values map associated with the fc. The forwarding class reverts to the defined parameters in the default remark policy.

|                   |                                                                                                                        |
|-------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | none                                                                                                                   |
| <b>Parameters</b> | <i>fc-name</i> — Specifies a case-sensitive system-defined forwarding class name for which policy entries are created. |
| <b>Values</b>     | be, l2, af, l1, h2, ef, h1, nc                                                                                         |

## Remark Policy Forwarding Class Commands

### de-mark-outer

|                    |                                                                     |
|--------------------|---------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] de-mark-outer [force <i>de-value</i>]</b>                   |
| <b>Context</b>     | config>qos> remark-policy> fc                                       |
| <b>Description</b> | <b>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C</b> |

This command is used to explicitly define the marking of the DEI bit of the outermost tag for *fc-name* according to the in and out of profile status of the packet.

If no *de-value* is present, the default values are used for the marking of the DE bit: for example, 0 for in-profile packets, 1 for out-of-profile ones – see IEEE 802.1ad-2005 standard.

If the *de-value* is specifically mentioned in the command line it means this value is to be used for all the packets of this forwarding class regardless of their in/out of profile status.



**Note:** If remarking is enabled, the inner tag DEI is always set to zero irrespective of *de-mark-outer*.

If remarking is enabled and *de-mark-outer* is not configured then, DEI bit of the outer tag is set to zero.

If remarking is disabled then, the DEI bits are preserved for both inner and outer tag.

|                   |                                                                             |
|-------------------|-----------------------------------------------------------------------------|
| <b>Parameters</b> | <i>de-value</i> — Specifies the DEI value to use for this forwarding class. |
| <b>Values</b>     | 0 or 1                                                                      |

### dot1p-inner

|               |                                                                |
|---------------|----------------------------------------------------------------|
| <b>Syntax</b> | <b>no dot1p-inner</b><br><b>dot1p-inner <i>dot1p-value</i></b> |
|---------------|----------------------------------------------------------------|

## Remark Policies for 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config> qos> remark> fc>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p><b>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C</b></p> <p>This command explicitly defines the egress IEEE 802.1P (dot1p) bits for the inner VLAN tag marking for fc-name. When the marking is set, all packets of fc-name that egresses out of QinQ SAP only (that is, SAPs configured with two VLAN tags explicitly defined, For example: SAP 1/1/5:10.100) use the explicitly defined dot1p-value.</p> <p>This command has no effect for egress packets sent out of all other non-QinQ SAPs for example: Dot1q SAP, NULL SAP, etc. Additionally, if the egress packets for fc-name are not IEEE 802.1Q or IEEE 802.1P encapsulated, this command has no effect. In other words, this command takes effect, only when the node adds 2 tags to the packet on the egress.</p> <p>If the no form of the command is executed, the default remarking values are used.</p> |
| <b>Default</b>     | no dot1p-inner                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <p><i>dot1p-value</i> — Specifies the 802.1p value to set for the frames in this forwarding class.</p> <p><b>Values</b>      0 — 7</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

### dot1p-inner

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] dot1p-inner [in-profile dot1p-value] [out-profile dot1p-value]</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config> qos> remark> fc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p><b>Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C</b></p> <p>The command will add the capability to mark on an egress the in and out profile status through a certain dot1p combination. It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.</p> <p>This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for fc-name. When the marking is set, all packets with in-profile status (or green color) of fc-name that egress out of QinQ SAP only (i.e. SAPs configured with two VLAN tags explicitly defined, e.g. SAP 1/1/5:10.100) use the explicitly defined in-profile and out-profile dot1p-value.</p> <p>This command has no effect for egress packets sent out of all other non-QinQ SAPs e.g. Dot1q SAP, NULL SAP, etc. Additionally, if the egress packets for fc-name are not IEEE 802.1Q or IEEE 802.1P encapsulated, this command has no effect. In other words, this command takes effect, only when the node adds 2 tags to the packet on the egress.</p> <p>This variant of the command is mutually exclusive to the use of dot1p-inner command. In other words, user has a choice to use either this command or the dot1p-inner command but not both together.</p> <p>If the no form of the command is executed, default remarking values are used for marking the inner VLAN.</p> |
| <b>Default</b>     | no dot1p-inner                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## Command Descriptions

- Parameters** *in-profile dot1p-value* — Specifies the Dot1p value to use for in-profile packets.  
**Values** 0-7
- out-profile dot1p-value* — Specifies the Dot1p bits to use for the out-profile packets.  
**Values** 0-7

### dot1p-outer

- Syntax** **no dot1p-outer**  
**dot1p-outer** *dot1p-value*
- Context** config> qos> remark> fc>
- Description** **Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C**
- This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for fc-name. When the marking is set, all packets of fc-name that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined dot1p-value. If the egress packets for fc-name are not IEEE 802.1Q or IEEE 802.1P encapsulated, the dot1p command has no effect.
- If the no form of the command is executed then software will use the default value.
- Default** no dot1p-outer
- Parameters** **dot1p-value** — Specifies the dot1p values to use.  
**Values** 0-7

### dscp-in-profile

- Syntax** **dscp-in-profile** *dscp-name*  
**no dscp-in-profile**
- Context** config> qos> remark> fc
- Description** **Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C**
- This command specifies the in-profile DSCP name for the forwarding class. When marking is set, the corresponding DSCP value is used to mark all IP packets with in-profile status, on the egress of this forwarding class queue.
- When multiple DSCP names are associated with the forwarding class in the policy, the last name entered will overwrite the previous value.
- The **no** form of this command reverts to the factory default in-profile dscp-name setting for policy-id 1.

## Remark Policies for 7210 SAS-K2F2T1C and 7210 SAS-K2F4T6C

**Parameters** *dscp-name* — System- or user-defined, case-sensitive *dscp-name*.

**Default** none

**Values** For 7210 SAS-K 2F2T1C:  
be|cp1|cp2|cp3|cp4|cp5|cp6|cp7|cs1|cp9|af11|cp11|af12|cp13|af13|  
cp15|cs2|cp17|af21|cp19|af22|cp21|af23|cp23|cs3|cp25|af31|cp27|  
af32|cp29|af33|cp31|cs4|cp33|af41|cp35|af42|cp37|af43|cp39|cs5|  
cp41|cp42|cp43|cp44|cp45|ef|cp47|nc1|cp49|cp50|cp51|cp52|cp53|  
cp54|cp55|nc2|cp57|cp58|cp59|cp60|cp61|cp62|cp63

### dscp-out-profile

**Syntax** **dscp-out-profile** *dscp-name*  
**no dscp-out-profile**

**Context** config> qos> remark> fc

**Description** **Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C**

This command specifies the out-of-profile DSCP name for the forwarding class. When marking is set, the corresponding DSCP value is used to mark all IP packets with out-of-profile status, on the egress of this forwarding class queue.

When multiple DSCP names are associated with the forwarding class in the policy, the last name entered will overwrite the previous value.

The **no** form of this command reverts to the factory default out-of-profile dscp-name setting for policy-id 1.

**Parameters** *dscp-name* — System- or user-defined, case-sensitive *dscp-name*.

**Default** none

**Values** For 7210 SAS-K 2F2T1C:  
be|cp1|cp2|cp3|cp4|cp5|cp6|cp7|cs1|cp9|af11|cp11|af12|cp13|af13|  
cp15|cs2|cp17|af21|cp19|af22|cp21|af23|cp23|cs3|cp25|af31|cp27|  
af32|cp29|af33|cp31|cs4|cp33|af41|cp35|af42|cp37|af43|cp39|cs5|  
cp41|cp42|cp43|cp44|cp45|ef|cp47|nc1|cp49|cp50|cp51|cp52|cp53|  
cp54|cp55|nc2|cp57|cp58|cp59|cp60|cp61|cp62|cp63

### dot1p-outer

**Syntax** [**no**] **dot1p-outer** [**in-profile dot1p-value**] [**out-profile dot1p-value**]

**Context** config> qos> remark> fc>

**Description** **Platforms Supported: 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C**

## Command Descriptions

The command will add the capability to mark on an egress the in and out profile status via a certain dot1p combination, similarly with the DEI options. It may be used when the internal in and out of profile status needs to be communicated to an adjacent network/customer device that does not support the DEI bit.

This command explicitly defines the egress IEEE 802.1P (dot1p) bits marking for fc-name. When the marking is set, all packets with in-profile status (or green color) of fc-name that have either an IEEE 802.1Q or IEEE 802.1P encapsulation use the explicitly defined in-profile dot1p-value. Similarly all the packet with out-of-profile status use the explicitly defined out-profile dot1p-value. If the egress packets for fc-name are not IEEE 802.1Q or IEEE 802.1P encapsulated, this command has no effect.

When this command is used the DEI Bit is left unchanged by the egress processing if a tag exists. If a new tag is added, the DEI bit is set to 0.

This variant of the command is mutually exclusive to the use of dot1p-outer command. In other words, user has a choice to use either this command or the dot1p-outer command but not both together.

If the no form of the command is executed then the default remark values are used.

**Default** no dot1p-outer

**Parameters** *in-profile dot1p-value* — Specifies the Dot1p value to use for in-profile packets.

**Values** 0-7

*out-profile dot1p-value* — Specifies the Dot1p bits to use for the out-profile packets.

**Values** 0-7

## lsp-exp-in-profile

**Syntax** **[no] lsp-exp-in-profile** *lsp-exp-value*

**Context** config> qos> remark> fc>

**Description** **Platforms Supported: 7210 SAS-K 2F4T6C**

This command specifies the in-profile MPLS EXP value for the forwarding class. When marking is set, the corresponding MPLS EXP value is used to mark all MPLS packets with in-profile status, on the egress of this forwarding class queue.

When multiple MPLS EXP names are associated with the forwarding class in the policy, the last name entered will overwrite the previous value.

The **no** form of this command reverts to the factory default in-profile MPLS EXP value setting for policy-id 1.

**Default** no lsp-exp-in-profile

**Parameters** *lsp-exp-value* —

**Values** 0-7

## lsp-exp-out-profile

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] lsp-exp-out-profile</b> <i>lsp-exp-value</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config> qos> remark> fc>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p><b>Platforms Supported: 7210 SAS-K 2F4T6C</b></p> <p>This command specifies the out-profile MPLS EXP value for the forwarding class. When marking is set, the corresponding MPLS EXP value is used to mark all MPLS packets with out-of-profile status, on the egress of this forwarding class queue.</p> <p>When multiple MPLS EXP names are associated with the forwarding class in the policy, the last name entered will overwrite the previous value.</p> <p>The <b>no</b> form of this command reverts to the factory default out-profile MPLS EXP value setting for policy-id 1.</p> |
| <b>Default</b>     | no lsp-exp-out-profile                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>lsp-exp-value</i> —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|                    | <b>Values</b> 0-7                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Show Commands

## remark-policy

| <b>Syntax</b>      | <b>remark-policy</b> [ <i>&lt;policy-id&gt;</i> ] [ <b>association detail</b> ]                                                                                                                                                                                                                         |       |             |                  |                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|-------------|------------------|---------------------------------------------|
| <b>Context</b>     | show>qos                                                                                                                                                                                                                                                                                                |       |             |                  |                                             |
| <b>Description</b> | This command displays remark policy information.                                                                                                                                                                                                                                                        |       |             |                  |                                             |
| <b>Parameters</b>  | <p><i>policy-id</i> — The ID of the remark policy.</p> <p><b>detail</b> — Displays detailed information about the remark policy.</p>                                                                                                                                                                    |       |             |                  |                                             |
| <b>Output</b>      | <p><b>Remark Policy Output Fields</b></p> <p>The following table describes remark policy output fields.</p> <table> <thead> <tr> <th>Label</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Remark Policy ID</td> <td>The ID that uniquely identifies the policy.</td> </tr> </tbody> </table> | Label | Description | Remark Policy ID | The ID that uniquely identifies the policy. |
| Label              | Description                                                                                                                                                                                                                                                                                             |       |             |                  |                                             |
| Remark Policy ID   | The ID that uniquely identifies the policy.                                                                                                                                                                                                                                                             |       |             |                  |                                             |

## Command Descriptions

| Label   | Description                             |
|---------|-----------------------------------------|
| Type    | Displays the type of the remark policy. |
| dot1P   | dot1p value for in-profile packets.     |
| fc name | Forwarding Class name.                  |

### Sample Output for 7210 SAS-K 2F2T1C

```
*A:SAH01-051> show qos remark-policy 300 detail
=====
QoS Remarking Policies
=====

Remark Policy-id : 300 Type : dscp
Description : (Not Specified)

FC Name DSCP DSCP
 InProf OutProf

be be be
l2 cs1 cs1
af af11 af12
l1 af21 af22
h2 af41 af41
ef ef ef
h1 nc1 nc1
nc nc2 nc2

Associations

SAP Egress

No SAP Egress Associations found.

Network

No Network Policy Associations found.
=====
==
*A:SAH01-051>show>qos#
```

# 1 Standards and Protocol Support



**Note:** The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

Conventions followed:

- M(A,N) stands for 7210 SAS-M in both Access-uplink mode and Network mode. Similarly M(N) stands for 7210 SAS-M in network mode only.
- T(A,N) stands for 7210 SAS-M in both Access-uplink mode and Network mode. Similarly T(N) stands for 7210 SAS-T in network mode only.
- K5 stands for 7210 SAS-K 2F2T1C
- K12 stands for 7210 SAS-K 2F4T6C
- Sx stands for all variants of 7210 SAS-Sx-1/10GE.
- S stands for all variants of 7210 SAS-S-1/10GE platforms.
- Sx-1/10GE stands for only the variants of 7210 SAS-Sx-1/10GE
- R6 stands for 7210 SAS-R6
- R12 stands for 7210 SAS-R12
- D stands for 7210 SAS-D and 7210 SAS-D ETR, if a line item applies only to 7210 SAS-D ETR, then it is indicated as D-ETR.
- E means 7210 SAS-E.
- X means 7210 SAS-X.

## BGP

- draft-ietf-idr-add-paths-04, Advertisement of Multiple Paths in BGP (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- draft-ietf-idr-best-external-03, Advertisement of the best external route in BGP (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- draft-ietf-sidr-origin-validation-signaling-04, BGP Prefix Origin Validation State Extended Community (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1772, Application of the Border Gateway Protocol in the Internet (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1997, BGP Communities Attribute (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2385, Protection of BGP Sessions via the TCP MD5 Signature Option (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

- 
- RFC 2439, BGP Route Flap Damping (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2545, Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2858, Multiprotocol Extensions for BGP-4 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2918, Route Refresh Capability for BGP-4 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3107, Carrying Label Information in BGP-4 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3392, Capabilities Advertisement with BGP-4 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4271, A Border Gateway Protocol 4 (BGP-4) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4360, BGP Extended Communities Attribute (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4456, BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4659, BGP/MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4724, Graceful Restart Mechanism for BGP (Helper Mode) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4760, Multiprotocol Extensions for BGP-4 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4798, Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4893, BGP Support for Four-octet AS Number Space (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5004, Avoid BGP Best Path Transitions from One External to Another (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5291, Outbound Route Filtering Capability for BGP-4 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5668, 4-Octet AS Specific BGP Extended Community (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6811, Prefix Origin Validation (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

---

## Circuit Emulation

RFC 4553, Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP) (M(N))

RFC 5086, Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN) (M(N))

RFC 5287, Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks (M(N))

## Ethernet

IEEE 802.1AB, Station and Media Access Control Connectivity Discovery (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.1ad, Provider Bridges (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.1ag, Connectivity Fault Management (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.1ah, Provider Backbone Bridges (M(N), X, T(N))

IEEE 802.1ax, Link Aggregation (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.1D, MAC Bridges (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.1p, Traffic Class Expediting (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.1Q, Virtual LANs (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.1s, Multiple Spanning Trees (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.1w, Rapid Reconfiguration of Spanning Tree (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.1X, Port Based Network Access Control (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.3ab, 1000BASE-T (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.3ac, VLAN Tag (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.3ad, Link Aggregation (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.3ae, 10 Gb/s Ethernet (M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

IEEE 802.3ah, Ethernet in the First Mile (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

- IEEE 802.3ba, 40 Gb/s and 100 Gb/s Ethernet (R6, R12)
- IEEE 802.3i, Ethernet (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- IEEE 802.3u, Fast Ethernet (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- IEEE 802.3z, Gigabit Ethernet (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- ITU-T G.8032, Ethernet Ring Protection Switching (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- ITU-T Y.1731, OAM functions and mechanisms for Ethernet based networks (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

## Fast Reroute

- draft-ietf-rtgwg-lfa-manageability-08, Operational management of Loop Free Alternates (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5286, Basic Specification for IP Fast Reroute: Loop-Free Alternates (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

## IP — General

- draft-grant-tacacs-02, The TACACS+ Protocol (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 768, User Datagram Protocol (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 793, Transmission Control Protocol (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 854, TELNET Protocol Specifications (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 951, Bootstrap Protocol (BOOTP) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1034, Domain Names - Concepts and Facilities (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1035, Domain Names - Implementation and Specification (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1350, The TFTP Protocol (revision 2) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1534, Interoperation between DHCP and BOOTP (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1542, Clarifications and Extensions for the Bootstrap Protocol (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

- 
- RFC 2131, Dynamic Host Configuration Protocol (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 2347, TFTP Option Extension (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 2348, TFTP Blocksize Option (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 2349, TFTP Timeout Interval and Transfer Size Options (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 2428, FTP Extensions for IPv6 and NATs (D, E, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 2865, Remote Authentication Dial In User Service (RADIUS) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 2866, RADIUS Accounting (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 3046, DHCP Relay Agent Information Option (Option 82) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 3596, DNS Extensions to Support IP version 6 (D, E, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 3768, Virtual Router Redundancy Protocol (VRRP) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 4250, The Secure Shell (SSH) Protocol Assigned Numbers (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 4251, The Secure Shell (SSH) Protocol Architecture (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 4254, The Secure Shell (SSH) Connection Protocol (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 4632, Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 5880, Bidirectional Forwarding Detection (BFD) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 5881, Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
  - RFC 6528, Defending against Sequence Number Attacks (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

---

## IP — Multicast

- RFC 1112, Host Extensions for IP Multicasting (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2236, Internet Group Management Protocol, Version 2 (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3306, Unicast-Prefix-based IPv6 Multicast Addresses (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3376, Internet Group Management Protocol, Version 3 (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3446, Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4601, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4604, Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4607, Source-Specific Multicast for IP (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4608, Source-Specific Protocol Independent Multicast in 232/8 (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5059, Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5384, The Protocol Independent Multicast (PIM) Join Attribute Format (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6513, Multicast in MPLS/BGP IP VPNs (T(N), Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6514, BGP Encodings and Procedures for Multicast in MPLS/IP VPNs (T(N), Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6515, IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs (T(N), Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6625, Wildcards in Multicast VPN Auto-Discover Routes (T(N), Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6826, Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path (T(N), Mxp, Sx/S-1/10GE, R6, R12)
- RFC 7385, IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points (T(N), Mxp, Sx/S-1/10GE, R6, R12)

---

## IP — Version 4

- RFC 791, Internet Protocol (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 792, Internet Control Message Protocol (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 826, An Ethernet Address Resolution Protocol (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1519, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1812, Requirements for IPv4 Routers (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1981, Path MTU Discovery for IP version 6 (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2401, Security Architecture for Internet Protocol (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2460, Internet Protocol, Version 6 (IPv6) Specification (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

## IP — Version 6

- RFC 2464, Transmission of IPv6 Packets over Ethernet Networks (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3021, Using 31-Bit Prefixes on IPv4 Point-to-Point Links (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3122, Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3587, IPv6 Global Unicast Address Format (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4007, IPv6 Scoped Address Architecture (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4193, Unique Local IPv6 Unicast Addresses (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4291, Internet Protocol Version 6 (IPv6) Addressing Architecture (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4443, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4861, Neighbor Discovery for IP version 6 (IPv6) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

- 
- RFC 5095, Deprecation of Type 0 Routing Headers in IPv6 (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5952, A Recommendation for IPv6 Address Text Representation (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6106, IPv6 Router Advertisement Options for DNS Configuration (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6164, Using 127-Bit IPv6 Prefixes on Inter-Router Links (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

## **IPsec**

- RFC 2401, Security Architecture for the Internet Protocol (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2406, IP Encapsulating Security Payload (ESP) (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

## **IS-IS**

- draft-ietf-isis-mi-02, IS-IS Multi-Instance (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- draft-kaplan-isis-ext-eth-02, Extended Ethernet Frame Size Support (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- ISO/IEC 10589:2002, Second Edition, Nov. 2002, Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3359, Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3719, Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3787, Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4971, Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

- RFC 5120, M-ISIS: Multi Topology (MT) Routing in IS-IS (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5130, A Policy Control Mechanism in IS-IS Using Administrative Tags (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5301, Dynamic Hostname Exchange Mechanism for IS-IS (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5302, Domain-wide Prefix Distribution with Two-Level IS-IS (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5303, Three-Way Handshake for IS-IS Point-to-Point Adjacencies (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5304, IS-IS Cryptographic Authentication (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5305, IS-IS Extensions for Traffic Engineering TE (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5306, Restart Signaling for IS-IS (Helper Mode) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5308, Routing IPv6 with IS-IS (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5310, IS-IS Generic Cryptographic Authentication (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6232, Purge Originator Identification TLV for IS-IS (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6233, IS-IS Registry Extension for Purges (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

## Management

- draft-ietf-snmv3-update-mib-05, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- draft-ietf-idr-bgp4-mib-05, Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- draft-ietf-isis-wg-mib-06, Management Information Base for Intermediate System to Intermediate System (IS-IS) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- draft-ietf-mpls-ldp-mib-07, Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

- 
- draft-ietf-mpls-lsr-mib-06, Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- draft-ietf-mpls-te-mib-04, Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- draft-ietf-ospf-mib-update-08, OSPF Version 2 Management Information Base (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- ianaaddressfamilynumbers-mib, IANA-ADDRESS-FAMILY-NUMBERS-MIB (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- ianaiftype-mib, IANAifType-MIB (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- ianaiprouteprotocol-mib, IANA-RTPROTO-MIB (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- IEEE8021-CFM-MIB, IEEE P802.1ag(TM) CFM MIB (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- IEEE8021-PAE-MIB, IEEE 802.1X MIB (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- IEEE8023-LAG-MIB, IEEE 802.3ad MIB (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- LLDP-MIB, IEEE P802.1AB(TM) LLDP MIB (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1157, A Simple Network Management Protocol (SNMP) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1215, A Convention for Defining Traps for use with the SNMP (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 1724, RIP Version 2 MIB Extension (Mxp)
- RFC 2021, Remote Network Monitoring Management Information Base Version 2 using SMIv2 (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2115, Management Information Base for Frame Relay DTEs Using SMIv2 (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2138, Remote Authentication Dial In User Service (RADIUS) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2206, RSVP Management Information Base using SMIv2 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2213, Integrated Services Management Information Base using SMIv2 (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2494, Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type (M(N))
- RFC 2571, An Architecture for Describing SNMP Management Frameworks (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

- 
- RFC 2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2573, SNMP Applications (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2574, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2575, View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2578, Structure of Management Information Version 2 (SMIv2) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2579, Textual Conventions for SMIv2 (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2787, Definitions of Managed Objects for the Virtual Router Redundancy Protocol (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2819, Remote Network Monitoring Management Information Base (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2856, Textual Conventions for Additional High Capacity Data Types (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2863, The Interfaces Group MIB (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2864, The Inverted Stack Table Extension to the Interfaces Group MIB (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2933, Internet Group Management Protocol MIB (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3014, Notification Log MIB (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3164, The BSD syslog Protocol (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3165, Definitions of Managed Objects for the Delegation of Management Scripts (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3231, Definitions of Managed Objects for Scheduling Management Operations (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3273, Remote Network Monitoring Management Information Base for High Capacity Networks (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3416, Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

- 
- RFC 3417, Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3419, Textual Conventions for Transport Addresses (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3584, Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3593, Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals (K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3635, Definitions of Managed Objects for the Ethernet-like Interface Types (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3826, The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3877, Alarm Management Information Base (MIB) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3895, Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types (M(N))
- RFC 4001, Textual Conventions for Internet Network Addresses (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4022, Management Information Base for the Transmission Control Protocol (TCP) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4113, Management Information Base for the User Datagram Protocol (UDP) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4220, Traffic Engineering Link Management Information Base (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4292, IP Forwarding Table MIB (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4293, Management Information Base for the Internet Protocol (IP) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6241, Network Configuration Protocol (NETCONF) (K5, K12, R6, R12)
- RFC 6242, Using the NETCONF Protocol over Secure Shell (SSH) (K5, K12, R6, R12)

## **MPLS — General**

- RFC 3031, Multiprotocol Label Switching Architecture (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 3032, MPLS Label Stack Encoding (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 3443, Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 4182, Removing a Restriction on the use of MPLS Explicit NULL (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5332, MPLS Multicast Encapsulations (T(N), Mxp, Sx/S-1/10GE, R6, R12)

## **MPLS — GMPLS**

draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

## **MPLS — LDP**

draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

draft-pdutta-mpls-ldp-v2-00, LDP Version 2 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

draft-pdutta-mpls-tldp-hello-reduce-04, Targeted LDP Hello Reduction (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 3037, LDP Applicability (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (Helper Mode) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5036, LDP Specification (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5443, LDP IGP Synchronization (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5561, LDP Capabilities (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 6826, Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

## **MPLS — MPLS-TP**

RFC 5586, MPLS Generic Associated Channel (T(N), R6, R12)

RFC 5921, A Framework for MPLS in Transport Networks (T(N), R6, R12)

- RFC 5960, MPLS Transport Profile Data Plane Architecture (T(N), R6, R12)  
RFC 6370, MPLS Transport Profile (MPLS-TP) Identifiers (T(N), R6, R12)  
RFC 6378, MPLS Transport Profile (MPLS-TP) Linear Protection (T(N), R6, R12)  
RFC 6426, MPLS On-Demand Connectivity and Route Tracing (T(N), R6, R12)  
RFC 6428, Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile (T(N), R6, R12)  
RFC 6478, Pseudowire Status for Static Pseudowires (T(N), R6, R12)  
RFC 7213, MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing (T(N), R6, R12)

## **MPLS — OAM**

- RFC 6424, Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)  
RFC 6425, Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping (T(N), Mxp, R6, R12)

## **MPLS — RSVP-TE**

- RFC 2702, Requirements for Traffic Engineering over MPLS (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)  
RFC 2747, RSVP Cryptographic Authentication (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)  
RFC 2961, RSVP Refresh Overhead Reduction Extensions (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)  
RFC 3097, RSVP Cryptographic Authentication -- Updated Message Type Value (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)  
RFC 3209, RSVP-TE: Extensions to RSVP for LSP Tunnels (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)  
RFC 3477, Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE) (M(N), T(N), X, Mxp, R6, R12)  
RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)  
RFC 4561, Definition of a Record Route Object (RRO) Node-Id Sub-Object (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)  
RFC 4875, Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs) (T(N), Mxp, R6, R12)  
RFC 4950, ICMP Extensions for Multiprotocol Label Switching (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5712, MPLS Traffic Engineering Soft Preemption (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5817, Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

## OSPF

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 1765, OSPF Database Overflow (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 2328, OSPF Version 2 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 3101, The OSPF Not-So-Stubby Area (NSSA) Option (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 3509, Alternative Implementations of OSPF Area Border Routers (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 3623, Graceful OSPF Restart Graceful OSPF Restart (Helper Mode) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 3630, Traffic Engineering (TE) Extensions to OSPF Version 2 (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 4222, Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 4552, Authentication/Confidentiality for OSPFv3 (M(N), T(N), X, Mxp, R6, R12)

RFC 4576, Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 4970, Extensions to OSPF for Advertising Optional Router Capabilities (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5185, OSPF Multi-Area Adjacency (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5187, OSPFv3 Graceful Restart (Helper Mode) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5243, OSPF Database Exchange Summary List Optimization (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5250, The OSPF Opaque LSA Option (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5340, OSPF for IPv6 (M(N), T(N), X, Mxp, R6, R12)

RFC 5709, OSPFv2 HMAC-SHA Cryptographic Authentication (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5838, Support of Address Families in OSPFv3 (M(N), T(N), X, Mxp, R6, R12)

RFC 6987, OSPF Stub Router Advertisement (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

## Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, OAM Procedures for VPWS Interworking (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 3916, Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 3985, Pseudo Wire Emulation Edge-to-Edge (PWE3) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 4385, Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 4446, IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 4447, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 5659, An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 6073, Segmented Pseudowire (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 6310, Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 6391, Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network (Mxp, R6, R12)

RFC 6718, Pseudowire Redundancy (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 6870, Pseudowire Preferential Forwarding Status bit (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 7023, MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

RFC 7267, Dynamic Placement of Multi-Segment Pseudowires (M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

---

## Quality of Service

- RFC 2430, A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE) (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 2598, An Expedited Forwarding PHB (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3140, Per Hop Behavior Identification Codes (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 3260, New Terminology and Clarifications for Diffserv (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

## RIP

- RFC 1058, Routing Information Protocol (Mxp)
- RFC 2082, RIP-2 MD5 Authentication (Mxp)
- RFC 2453, RIP Version 2 (Mxp)

## Timing

- GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005 (D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000 (D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- IEEE 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems (D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx-1/10GE, R6, R12)
- ITU-T G.781, Synchronization layer functions, issued 09/2008 (D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003 (D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- ITU-T G.8261, Timing and synchronization aspects in packet networks, issued 04/2008 (D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- ITU-T G.8262, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007 (D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)
- ITU-T G.8264, Distribution of timing information through packet networks, issued 10/2008 (D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

- 
- ITU-T G.8265.1, Precision time protocol telecom profile for frequency synchronization, issued 10/2010 (D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx-1/10GE, R6, R12)
- ITU-T G.8275.1, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014 (R6, R12)
- RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification (D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, R12)

## **VPLS**

- RFC 4761, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 4762, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 5501, Requirements for Multicast Support in Virtual Private LAN Services (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)
- RFC 6074, Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs) (K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, R6, R12)

# Customer Document and Product Support



## Customer documentation

[Customer Documentation Welcome Page](#)



## Technical Support

[Product Support Portal](#)



## Documentation feedback

[Customer Documentation Feedback](#)

