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

The ingress component of the policy defines how DiffServ code points (DSCPs) and MPLS EXP bits 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 router. 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.

The egress component of the network QoS policy defines the DiffServ oriented queuing parameters associated with each forwarding class.

Each forwarding class defined within the system automatically creates a queue on each network interface. This queue gets all the parameters defined within the default network QoS policy 1 until an explicit policy is defined for the network interface.

If the egressing packet originated on an ingress SAP, or the remarking parameter is defined for the egress interface, the egress QoS policy also defines the IP DSCP or MPLS EXP bit marking based on the forwarding class and the profile state.

Network **policy-id 1** exists as the default policy that is applied to all network interfaces by default. The network **policy-id 1** cannot be modified or deleted. It defines the default DSCP-to-FC mapping and MPLS EXP-to-FC for the ingress. For the egress, it defines six forwarding classes which represent individual queues and the packet marking criteria.

New (non-default) network policy parameters can be modified. The **no** form of the command reverts the object to the default values. A new network policy must include the definition of at least one queue and specify the default-action. Incomplete network policies cannot be applied to network interfaces.

Changes made to a policy are applied immediately to all network interface 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 router devices, refer to CLI Usage chapter in the Basic System Configuration Guide.

Network Ingress Tunnel QoS Override

For Tunnel Terminated IP Routing Decisions

This section describes a mechanism that provides the ability to ignore the network ingress QoS mapping of a terminated tunnel containing an IP packet that is to be routed to a base router or VPRNdestination. This is advantageous when the mapping for the tunnel QoS marking does not accurately or completely reflect the required QoS handling for the IP routed packet. When the mechanism is enabled on an ingress network IP interface, the IP interface will ignore the tunnel's QoS mapping and derive the internal forwarding class and profile based on the precedence or DiffServe Code Point (DSCP) values within the routed IP header ToS field compared to the Network QoS policy defined on the IP interface.

Normal QoS Operation

The following types of QoS mapping decisions are applicable on a network ingress IP interface .

- Ethernet Dot1P value mapping (if defined)
- Default QoS mapping
- IP ToS precedence mapping
- IP ToS DSCP mapping
- MPLS LSP EXP mapping

The default QoS mapping always exists on an ingress IP interface and every received packet will be mapped to this default if another explicitly defined matching entry does not exist.

A tunnel that terminates on the ingress IP interface (the node is the last hop for the tunnel) is evaluated based on the type of tunnel, IP GRE or MPLS LSP. An IP tunneled packet may match a Dot1P entry, IP ToS precedence entry or IP ToS DSCP entry when defined in the applied policy. An MPLS LSP may match a Dot1P entry or MPLS EXP entry when defined.

The internal tunnel encapsulated packet is never evaluated for QoS determination when operating in normal mode.

Tunnel Termination QoS Override Operation

Tunnel termination QoS override only applies to IP routing decisions once the tunnel encapsulation is removed. Non-IP routed packets within a terminating tunnel are ignored by the override and are forwarded as described in the Normal QoS Operation section.

When tunnel termination QoS override is enabled, the ToS field within the routed IP header is evaluated against the IP ToS precedence and DSCP entries in the applied network QoS policy on the ingress IP interface. If an explicit match entry is not found, the default QoS mapping is used. Any Dot1P and MPLS LSP EXP bits within the packet are ignored. If the packet was IP GRE tunneled to the node, the tunnel IP header ToS field is ignored as well.

Any tunnel received on the ingress IP interface that traverses the node (the node is not the ultimate hop for the tunnel) is not affected by the QoS override mechanism and is forwarded as described in Normal QoS Operation section.

Enabling and Disabling Tunnel Termination QoS Override

Tunnel termination QoS override is enabled and disabled within the network QoS policy under the ingress node. The default condition within the policy is not to override tunnel QoS for IP routed packets.

DSCP Marking CPU Generated Traffic

Specific DSCP, forwarding class, and Dot1P parameters can be specified to be used by every protocol packet generated by the node. This enables prioritization or de-prioritization of every protocol (as required). The markings effect a change in behavior on ingress when queuing. For example, if OSPF is not enabled, then traffic can be de-prioritized to best effort (be) DSCP. This change de-prioritizes OSPF traffic to the CPU complex.

DSCP marking for internally generated control and management traffic by marking the DSCP value should be used for the given application. This can be configured per routing instance. For example, OSPF packets can carry a different DSCP marking for the base instance and then for a VPRN service. IS-IS and ARP traffic is not an IP-generated traffic type and is not DSCP configurable, but they are Dot1p configurable.

When an application is configured to use a specified DSCP value then the MPLS EXP, Dot1P bits will be marked in accordance with the network or access egress policy as it applies to the logical interface the packet will be egressing.

The DSCP value can be set per application. This setting will be forwarded to the egress IOM. The egress IOM does not alter the coded DSCP value and marks the LSP-EXP and IEEE 802.1p (Dot1P) bits according to the appropriate network or access QoS policy.

Sgt-qos is supported in the base router, VPRN and management contexts.

The default values for self-generated traffic are:

- Routing protocols (OSPF, BGP, etc)
  - → Forwarding class : Network Control (NC)
  - → DSCP value: NC1
  - → 802.1P value: 6 (ARP, PPPoE, IS-IS only)
- Management protocols (SSH, SNMP, etc)
  - → Forwarding class: Assured Forwarding (AF)
  - → DSCP value: AF41

Table 20: DSCP/FC Marking

Protocol	IPv4	IPv6	DSCP Marking	Dot1P Marking	Default FC
IS-IS				Yes	NC
ARP				Yes	NC
PPPoE				Yes	NC

Table 20: DSCP/FC Marking (Continued)

Protocol	IPv4	IPv6	DSCP Marking	Dot1P Marking	Default FC
LDP (T-LDP)	Yes				NC
RSVP	Yes		Yes	Yes	NC
BGP	Yes	Yes	Yes	Yes	NC
RIP	Yes	Yes	Yes	Yes	NC
MSDP	Yes				NC
PIM (SSM)	Yes	Yes	Yes	Yes	NC
OSPF	Yes	Yes	Yes	Yes	NC
SMTP	Yes				AF
IGMP/MLD	Yes	Yes	Yes	Yes	AF
Telnet	Yes	Yes	Yes	Yes	AF
TFTP	Yes		Yes	Yes	AF
FTP	Yes				AF
SSH (SCP)	Yes	Yes	Yes	Yes	AF
SNMP (get, set, etc.)	Yes	Yes	Yes	Yes	AF
SNMP trap/log	Yes	Yes	Yes	Yes	AF
syslog	Yes	Yes	Yes	Yes	AF
ICMP	Yes	Yes	Yes	Yes	AF
Traceroute	Yes	Yes	Yes	Yes	AF
TACPLUS	Yes	Yes	Yes	Yes	AF
DNS	Yes	Yes	Yes	Yes	AF
SNTP/NTP	Yes				AF
RADIUS	Yes				AF
Cflowd	Yes				AF
DHCP	Yes	Yes	Yes	Yes	AF
IPv6 Neighbor Discovery	Yes				NC

NOTE: The ICMP entry under sgt-qos is not referenced for ICMP ECHO_REQUEST (8) and ECHO_RESPONSE (0) packet types.

Default DSCP Mapping Table

```
DSCP Name   DSCP Value   DSCP Value  DSCP Value   Label
            Decimal      Hexadecimal Binary
============================================================
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     l1
af22        20           0x14        0b010100     l1
af23        22           0x16        0b010110     l1
af31        26           0x1a        0b011010     l1
af32        28           0x1c        0b011100     l1
af33        30           0x1d        0b011110     l1
af41        34           0x22        0b100010     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.

Basic Configurations

A basic network QoS policy must conform to the following:

- Each network QoS policy must have a unique policy ID.
- Include the definition of at least one queue.
- Specify the default-action.

Create a Network QoS Policy

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

To create an network QoS policy when operating, 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.
- You can modify egress criteria to customize the forwarding class queues to be instantiated. Otherwise, the default values are applied.
  - → Remarking — When enabled, this command remarks ALL packets that egress on the specified network port. The remarking is based on the forwarding class to DSCP and LSP EXP bit mapping defined under the egress node of the network QoS policy.
  - → Forwarding class criteria — The forwarding class name represents an egress queue. Specify forwarding class criteria to define the egress characteristics of the queue and the marking criteria of packets flowing through it.
  - → DSCP — The DSCP value is used for all IP packets requiring marking that egress on this forwarding class queue that are *in* or *out* of profile.
  - → LSP EXP — The EXP value is used for all MPLS labeled packets requiring marking that egress on this forwarding class queue that are *in* or *out* of profile.
- Ingress criteria — Specifies the DSCPDot1p to forwarding class mapping for all IP packets and define the MPLS EXP bits to forwarding class mapping for all labeled packets.
  - → Default action — Defines the default action to be taken for packets that have an undefined DSCP or MPLS EXP bits set. The default-action specifies the forwarding class to which such packets are assigned.
  - → DSCP — Creates a mapping between the DSCP of the network ingress traffic and the forwarding class. Ingress traffic that matches the specified DSCP will be assigned to the corresponding forwarding class.

→ LSP EXP — Creates a mapping between the LSP EXP bits of the network ingress traffic and the forwarding class. Ingress traffic that matches the specified LSP EXP bits will be assigned to the corresponding forwarding class.

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

CLI Syntax:

```
config>qos#
    network network-policy-id
        description description-string
        scope {exclusive|template}
        egress
            remarking
            fc {be|l2|af|l1|h2|ef|h1|nc}
                dot1p-in-profile dot1p-priority
                dot1p-out-profile dot1p-priority
                dscp-in-profile dscp-name
                dscp-out-profile dscp-name
                lsp-exp-in-profile mpls-exp-value
                lsp-exp-out-profile mpls-exp-value
            default-action fc {be|l2|af|l1|h2|ef|h1|nc} profile
                {in|out}
            dot1p dot1p-priority fc {fc-name} profile {in|out}
            dscp dscp-name fc {be|l2|af|l1|h2|ef|h1|nc} profile
                {in|out}
            ler-use-dscp
            lsp-exp lsp-exp-value fc fc-name profile {in|out}
```

```
A:ALA-10:A:ALA-12>config>qos# info
#------------------------------------------
echo "QoS Policy Configuration"
#------------------------------------------
...
    network 600 create
            description "Network Egress Policy"
            ingress
                default-action fc ef profile in
            exit
            egress
                remarking
            exit
        exit
...
#------------------------------------------
A:ALA-12>config>qos#
```

Applying Network Policies

Use the following CLI syntax to apply network policies to the routeraccess uplink portsIP interfaces:

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

The following output displays the configuration for router interface ALA-1-2 with network policy 600 applied to the interface.

```
A:ALA-7>config>router# info
#------------------------------------------
echo "IP Configuration"
#------------------------------------------
...
        interface "ALA-1-2"
            address 10.10.4.3/24
            qos 600
        exit
...
--------------------------------------------
A:ALA-7>config>router#
```

Default Network Policy Values

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

Table 21: Network Policy Defaults

Field	Default	
description	Default network QoS policy.	
scope	template	
ingress		
default-action	fc be profile out	
dscp:		
be	fc be	profile out
ef	fc ef	profile in
cs1	fc l2	profile in
nc1	fc h1	profile in
nc2	fc nc	profile in
af11	fc af	profile in
af12	fc af	profile out
af13	fc af	profile out
af21	fc l1	profile in
af22	fc l1	profile out
af23	fc l1	profile out
af31	fc l1	profile in
af32	fc l1	profile out
af33	fc l1	profile out
af41	fc h2	profile in
af42	fc h2	profile out

Table 21: Network Policy Defaults (Continued)

Field	Default	
af43	fc h2	profile out
lsp-exp:		
0	fc be	profile out
1	fc l2	profile in
2	fc af	profile out
3	fc af	profile in
4	fc h2	profile in
5	fc ef	profile in
6	fc h1	profile in
7	fc nc	profile in
egress		
remarking	no	
fc af:		
dscp-in-profile	af11	
dscp-out-profile	af12	
lsp-exp-in-profile	3	
lsp-exp-out-profile	2	
fc be:		
dscp-in-profile	be	
dscp-out-profile	be	
lsp-exp-in-profile	0	
lsp-exp-out-profile	0	
fc ef:		
dscp-in-profile	ef	
dscp-out-profile	ef	

Table 21: Network Policy Defaults (Continued)

Field	Default
lsp-exp-in-profile	5
lsp-exp-out-profile	5
fc h1:	
dscp-in-profile	nc1
dscp-out-profile	nc1
lsp-exp-in-profile	6
lsp-exp-out-profile	6
fc h2:	
dscp-in-profile	af41
dscp-out-profile	af42
lsp-exp-in-profile	4
lsp-exp-out-profile	4
fc l1:	
dscp-in-profile	af21
dscp-out-profile	af22
lsp-exp-in-profile	3
lsp-exp-out-profile	2
fc l2:	
dscp-in-profile	cs1
dscp-out-profile	cs1
lsp-exp-in-profile	1
lsp-exp-out-profile	1
fc nc:	
dscp-in-profile	nc2
dscp-out-profile	nc2
lsp-exp-in-profile	7

Table 21: Network Policy Defaults (Continued)

Field	Default
lsp-exp-out-profile	7

The following output displays the default configuration:

```
A:ALA-49>config>qos>network# info detail
----------------------------------------------
            description "Default network QoS policy."
            scope template
            ingress
                default-action fc be profile out
                no ler-use-dscp
                dscp be fc be profile out
                dscp ef fc ef profile in
                dscp cs1 fc l2 profile in
                dscp nc1 fc h1 profile in
                dscp nc2 fc nc profile in
                dscp af11 fc af profile in
                dscp af12 fc af profile out
                dscp af13 fc af profile out
                dscp af21 fc l1 profile in
                dscp af22 fc l1 profile out
                dscp af23 fc l1 profile out
                dscp af31 fc l1 profile in
                dscp af32 fc l1 profile out
                dscp af33 fc l1 profile out
                dscp af41 fc h2 profile in
                dscp af42 fc h2 profile out
                dscp af43 fc h2 profile out
                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
            egress
                no remarking
                fc af
                    dscp-in-profile af11
                    dscp-out-profile af12
                    lsp-exp-in-profile 3
                    lsp-exp-out-profile 2
                    dot1p-in-profile 2
                    dot1p-out-profile 2
                exit
                fc be
                    dscp-in-profile be
                    dscp-out-profile be
                    lsp-exp-in-profile 0
```

```
                    lsp-exp-out-profile 0
                    dot1p-in-profile 0
                    dot1p-out-profile 0
                exit
                fc ef
                    dscp-in-profile ef
                    dscp-out-profile ef
                    lsp-exp-in-profile 5
                    lsp-exp-out-profile 5
                    dot1p-in-profile 5
                    dot1p-out-profile 5
                exit
                fc h1
                    dscp-in-profile nc1
                    dscp-out-profile nc1
                    lsp-exp-in-profile 6
                    lsp-exp-out-profile 6
                    dot1p-in-profile 6
                    dot1p-out-profile 6
                exit
                fc h2
                    dscp-in-profile af41
                    dscp-out-profile af42
                    lsp-exp-in-profile 4
                    lsp-exp-out-profile 4
                    dot1p-in-profile 4
                    dot1p-out-profile 4
                exit
                fc l1
                    dscp-in-profile af21
                    dscp-out-profile af22
                    lsp-exp-in-profile 3
                    lsp-exp-out-profile 2
                    dot1p-in-profile 3
                    dot1p-out-profile 3
                exit
                fc l2
                    dscp-in-profile cs1
                    dscp-out-profile cs1
                    lsp-exp-in-profile 1
                    lsp-exp-out-profile 1
                    dot1p-in-profile 1
                    dot1p-out-profile 1
                exit
                fc nc
                    dscp-in-profile nc2
                    dscp-out-profile nc2
                    lsp-exp-in-profile 7
                    lsp-exp-out-profile 7
                    dot1p-in-profile 7
                    dot1p-out-profile 7
                exit
            exit
-------------------------------------------
A:ALA-49>config>qos>network#
```

Service Management Tasks

Deleting QoS Policies

A network policy is associated by default with router interfaces.

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.

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

The following output displays a sample configuration.

```
A:ALA-7>config>router# info
#------------------------------------------
echo "IP Configuration"
#------------------------------------------
...
        interface "ALA-1-2"
            address 10.10.4.3/24 broadcast host-ones
            no port
            no arp-timeout
            no allow-directed-broadcasts
            icmp
                mask-reply
                redirects 100 10
                unreachables 100 10
                ttl-expired 100 10
            exit
            qos 1
            ingress
                no filter
            exit
            egress
                no filter
            exit
            no mac
            no ntp-broadcast
            no cflowd
            no shutdown
        exit
        interface "ALA-1-3"
...
#------------------------------------------
A:ALA-7>config>router#
```

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
                dscp be fc be profile out
                dscp ef fc ef profile in
                dscp cs1 fc l2 profile in
                dscp nc1 fc h1 profile in
                dscp nc2 fc nc profile in
                dscp af11 fc af profile in
                dscp af12 fc af profile out
                dscp af13 fc af profile out
                dscp af21 fc l1 profile in
                dscp af22 fc l1 profile out
...
        network 600 create
            description "Default network QoS policy."
            scope template
            ingress
                default-action fc be profile out
                dscp be fc be profile out
                dscp ef fc ef profile in
                dscp cs1 fc l2 profile in
                dscp nc1 fc h1 profile in
                dscp nc2 fc nc profile in
                dscp af11 fc af profile in
                dscp af12 fc af profile out
                dscp af13 fc af profile out
                dscp af21 fc l1 profile in
```

```
                dscp af22 fc l1 profile out
...
        network 700 create
            description "Default network QoS policy."
            scope template
            ingress
                default-action fc be profile out
                dscp be fc be profile out
                dscp ef fc ef profile in
                dscp cs1 fc l2 profile in
                dscp nc1 fc h1 profile in
                dscp nc2 fc nc profile in
                dscp af11 fc af profile in
                dscp af12 fc af profile out
                dscp af13 fc af profile out
                dscp af21 fc l1 profile in
                dscp af22 fc l1 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 interfaces 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.

Resource Allocation for Network QoS policy

This section describes the allocation of QoS resources for network QoS policy (for type=ipinterface).

When an IP interface is created, a default network QoS policy is applied. For the default policy, two meters and two classification entries in hardware are allocated.

The resources are allocated to a network policy, only when a port is configured for the IP interface.

For every FC in use, the system allocates two classification entries in hardware. If multiple matchcriteria entries map to the same FC, then each of these are allocated two classification entries in hardware. For example, if there are two match-criteria entries that map to FC ‘af’, then a total of four classification entries are allocated in hardware and if there are four match-criteria entries that map to FC ‘af’, then a total of 8 classification entries are allocated in hardware.

For every meter or policer in use, the system allocates one meter in hardware. A meter or policer is considered to be in use when it is associated with an FC in use.

The number of IP interfaces allowed is limited to number of resources available in hardware, subject to system limit (a maximum of 32 IP interfaces are allowed). The system reserves a total of 512 classification entries and 256 meters in hardware for use by network policy associated with an IP interface.

For computing the number of QoS resources used by an IP interface:

- Determine number of match-criteria entries used to identify the FC.
- Determine number of FCs to use.

Only the FCs used by the match-criteria classification entries are to be considered for the 'number of FCs'. Therefore are referred to as ‘FC in use’.

Use the following rules to compute the number of classification entries per FC in use:

If a FC is in use and is created without explicit meters, use default meter#1 for unicast traffic and default meter #9 for all other traffic types (that is, broadcast, multicast and unknown-unicast). This requires two classification entries in hardware.

If a FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and use default meter #9 for all other traffic types. This requires two classification entries in hardware.

If a FC is in use and is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other kinds of traffic. This requires two classification entries in hardware.

Given the number of match criteria and the number of FCs used, use the equation given below to arrive at total number of classification entries per policy (for example TC):

$$TC = \Sigma_{i=nc,h1,ef,h2,l1,af,l2,be} \; 2 * E(i)$$

Where,

E(i) is the number of match- criteria entries that classify packets to FCi. For 7210 platforms, the maximum number of classification entries per policy can be 64 (including default).

2 is the number of classification entries that are required by FCi.

Note: In any case, only 2 classification entries are used per FC in a network policy, as only two traffic-types are supported.

Determine number of policers or meters to use (for example TP). A maximum of 12 meters per network policy is available.

Only those meters that are associated with FCs need to be considered for number of meters. Note, that only FCs in use are considered.