QoS Classification and Marking on Catalyst 6500/6000 Series Switches Running CatOS Software

Document ID: 23420

Contents

Introduction

This document examines what happens regarding the marking and classification of a packet at different places during its journey within the Catalyst 6000 chassis. It mentions special cases, restrictions, and provides short case studies.

This document is not intended to be an exhaustive list of all Catalyst OS (CatOS) commands regarding Quality of Service (QoS) or marking. For more information on the CatOS command−line interface (CLI), refer to the following document:

- Configuring QoS

Note: This document only considers IP traffic.

Before You Begin

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

Prerequisites

There are no specific prerequisites for this document.

Components Used

This document is valid for Catalyst 6000 family switches running CatOS Software, and using one of the following Supervisor Engines:

- SUP1A + PFC
- SUP1A + PFC + MSFC
- SUP1A + PFC + MSFC2
- SUP2 + PFC2
- SUP2 + PFC2 + MSFC2

All sample commands, however, have been tried on a Catalyst 6506 with SUP1A/PFC running software version 6.3.

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Terminology

The following is a list of terminology used in this document:

- Differentiated Services Code Point (DSCP): The first six bits of the Type of Service (ToS) byte in the IP header. DSCP is only present in the IP packet.

  Note: You also assign an internal DSCP to every packet (IP or non IP), this internal DSCP assignment will be detailed later in this document.
- IP precedence: The first three bits of the ToS byte in the IP header.
- Class of Service (CoS): The only field that can be used to mark a packet at Layer 2 (L2). It consists of any of the following three bits:

  - The three dot1p bits in the dot1q tag for the IEEE dot1q packet.
  - The three bits called "User Field" in the Inter−Switch Link (ISL) header for an ISL encapsulated packet.
  - There is not CoS present inside a non−dot1q or an ISL packet.
- Classification: The process used to select the traffic to be marked.
- Marking: The process of setting a Layer 3 (L3) DSCP value in a packet. In this document, the definition of marking is extended to include setting L2 CoS values.

Catalyst 6000 family switches are able to make classifications based on the following three parameters:

- DSCP

- IP precedence
- CoS

The Catalyst 6000 family switches are making classification and marking at different places. The following is a look at what happens at these different places:

- Input port (ingress Application−Specific Integrated Circuit (ASIC))
- Switching engine (Policy Feature Card (PFC))
- Output port (egress ASIC)

Enabling QoS

By default, QoS is disabled on Catalyst 6000 switches. QoS can be enabled by issuing the CatOS command **set qos enable**.

When QoS is disabled there is no classification or marking done by the switch, and as such, every packet leaves the switch with the DSCP/IP precedence it had when entering the switch.

Input Port Handling

The main configuration parameter for the ingress port, regarding classification, is the trust state of the port. Each port of the system can have one of the following trust states:

- trust−ip−precedence
- trust−dscp
- trust−cos
- untrusted

The remainder of this section describes how the port trust states influences the final classification of the packet. The port trust state can be set or changed using the following CatOS command:

set port qos *mod/port* **trust {untrusted | trust−cos | trust−ipprec | trust−dscp }**

Note: By default all ports are in the untrusted state when QoS is enabled.

At the input port level you can also apply a default CoS per port, as in the following example:

set port qos *mod/port* **cos** *cos−value*

If the port is set to the untrusted state, simply mark the frame with the port default CoS and pass the header to the switching engine (PFC). If the port is set to one of the trust states, apply the default port CoS (if the frame does not have a received CoS (dot1q or ISL)), or keep the CoS as it is (for dot1q and ISL frames) and pass the frame to the switching engine. The input classification is illustrated in the following flowchart:

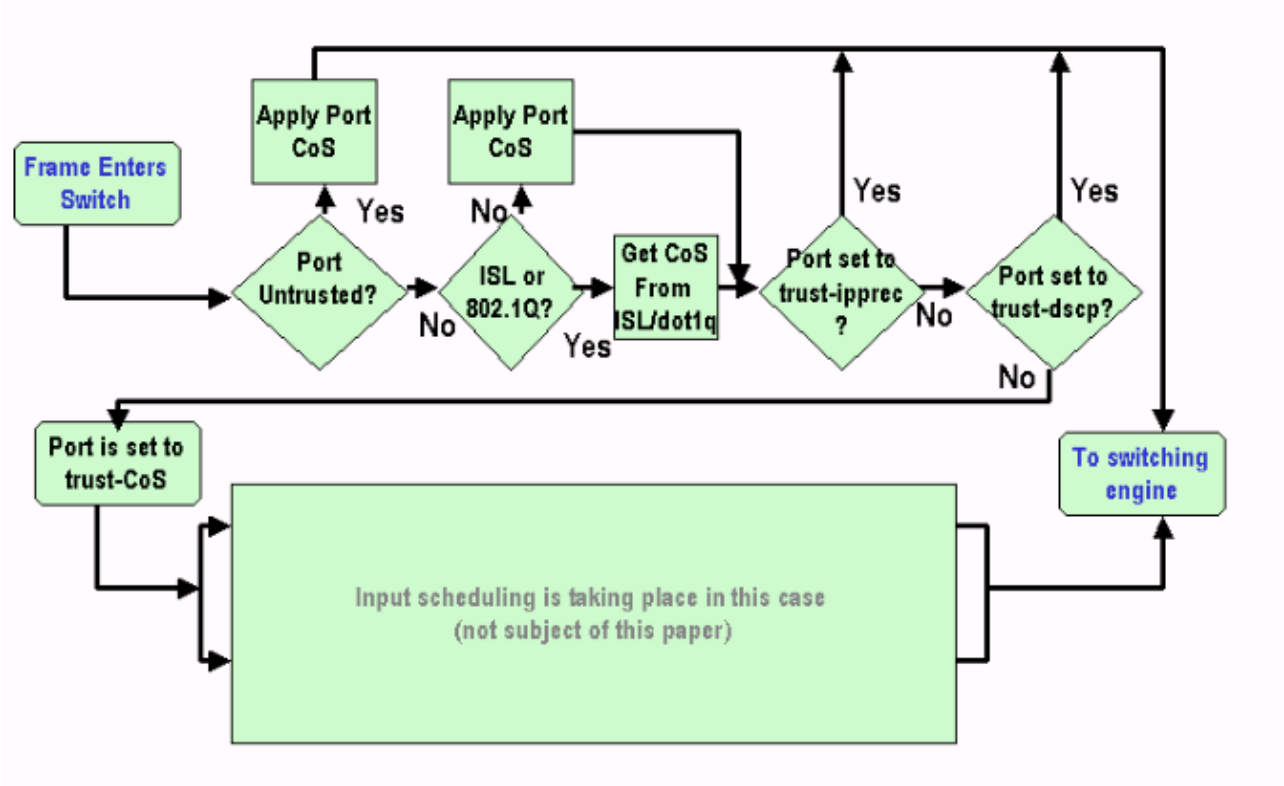

Note: As shown in the above flowchart, each frame will have an internal CoS assigned (either the received CoS, or the default port CoS), including untagged frames that do not carry any real CoS. This internal CoS and the received DSCP are written in a special packet header (called a Data Bus header) and sent over the Data Bus to the switching engine. This happens at the ingress line card and at this point it is not known yet whether this internal CoS will be carried to the egresss ASIC and inserted in the outgoing frame. This all depends on what the PFC does and is further described in the next section.

Switching Engine (PFC)

Once the header has reached the switching engine, the switching engine Encoded Address Recognition Logic (EARL) will assign each frame an internal DSCP. This internal DSCP is an internal priority assigned to the frame by the PFC as it transits the switch. This is not the DSCP in the IPv4 header. It is derived from an existing CoS or ToS setting and is used to reset the CoS or ToS as the frame exits the switch. This internal DSCP is assigned to all frames switched (or routed) by the PFC, even non−IP frames.

Four Possible Sources for Internal DSCP

The internal DSCP will be derived from one of the following:

1. An existing DSCP value, set prior to the frame entering the switch.
2. The received IP precedence bits already set in the IPv4 header. Since there are 64 DSCP values and only eight IP precedence values, the administrator will configure a mapping that is used by the switch to derive the DSCP. Default mappings are in place, should the administrator not configure the maps.
3. The received CoS bits already set prior to the frame entering the switch, or from the default CoS of the incoming port if there was no CoS in the incoming frame. As with IP precedence, there are a maximum of eight CoS values, each of which must be mapped to one of the 64 DSCP values. This map can be configured, or the switch can use the default map already in place.
4. The DSCP can be set for the frame using a DSCP default value typically assigned though an Access Control List (ACL) entry.

For Nos. 2 and 3 in the above list, the static mapping used is by default, as follows:

- DSCP derived equals eight times CoS, for CoS to DSCP mapping.
- DSCP derived equals eight times IP precedence, for IP precedence to DSCP mapping.

This static mapping can be overridden by the user by issuing the following commands:

set qos ipprec−dscp−map <dscp1> <dscp2>...<dscp8>

set qos cos−dscp−map <dscp1> <dscp2>...<dscp8>

The first value of the DSCP corresponding to the mapping for the CoS (or IP precedence) is "0", the second for the CoS (or IP precedence) is "1", and continuing in that pattern.

Which of the Four Possible Sources for Internal DSCP will be Used?

This section describes the rules that determine which of the four possible sources described above will be used for each packet. That depends on the following parameters:

1. What QoS ACL will be applied to the packet? This is determined by the following rules:

   Note: Each packet goes through an ACL entry.

   - ♦ If there is no ACL attached to the incoming port or VLAN, apply the default ACL.
   - ♦ If there is an ACL attached to the incoming port or VLAN, and if the traffic matches one of the entries in the ACL, use this entry.
   - ♦ If there is an ACL attached to the incoming port or VLAN, and if the traffic *does not* match one of the entries in the ACL, use the default ACL.
2. Each entry contains a classification keyword. The following is a list of possible keywords and their descriptions:

   - ♦ trust−ipprec: The internal DSCP will be derived from the received IP precedence according to the static mapping regardless of what the port trust state may be.
   - ♦ trust−dscp: The internal DSCP will be derived from the received DSCP regardless of what the port trust state may be.
   - ♦ trust−cos: The internal DSCP will be derived from the received CoS according to the static mapping, if the port trust state is trusted (trust−cos, trust−dscp, trust−ipprec). If the port trust state is trust−xx, the DSCP will be derived from the default port CoS according to the same static mapping.
   - ♦ dscp *xx*: The internal DSCP will depend on the following incoming port trust states:

     - ◊ If the port is untrusted, the internal DSCP will be set to *xx*.
     - ◊ If the port is trust−dscp, the internal DSCP will be the DSCP received in the incoming packet.
     - ◊ If the port is trust−CoS, the internal DSCP will be derived from the CoS of the received packet.
     - ◊ If the port is trust−ipprec, the internal DSCP will be derived from the IP precedence of the received packet.
3. Each QoS ACL can be applied either to a port or to a VLAN, but there is an additional configuration parameter to take into account; the ACL port type. A port can be configured to be VLAN−based or port−based. The following is a description of the two types of configurations:

   - ♦ A port configured to be VLAN−based will only look to ACL applied to the VLAN to which the port belongs. If there is an ACL attached to the port, the ACL will be ignored for the

packet coming in on that port.

- If a port belonging to a VLAN is configured as port−based, even if there is an ACL attached to that VLAN, it will not be taken into consideration for the traffic coming in from that port.

The following is a syntax to create a QoS ACL to mark IP traffic:

set qos acl ip *acl_name* **[dscp xx | trust−cos | trust−dscp | trust−ipprec]** *acl entry rule*

The following ACL, will mark all IP traffic directed to host 1.1.1.1 with a DSCP of "40" and will trust−dscp for all other IP traffic:

set qos acl TEST_ACL dscp 40 ip any host 1.1.1.1

set qos acl TEST_ACL trust−dscp ip any any

Once the ACL has been created you need to map it to a port or a VLAN, this can be done by issuing the following command:

set qos acl map *acl_name* **[module/port | VLAN]**

By default, each port is port−based for the ACL, so if you want to attach an ACL to a VLAN, you need to configure the ports of this VLAN as vlan−based. This can be done by issuing the following command:

set port qos *module/port* **vlan−based**

It can also be reverted back to port−based mode by issuing the following command:

set port qos *module/port* **port−based**

Summary: How is the Internal DSCP Chosen?

The internal DSCP depends on the following factors:

- port trust state
- ACL attached to the port
- default ACL
- VLAN−based or port−based in regards to the ACL

The following flow chart summarizes how the internal DSCP is chosen depending on the configuration:

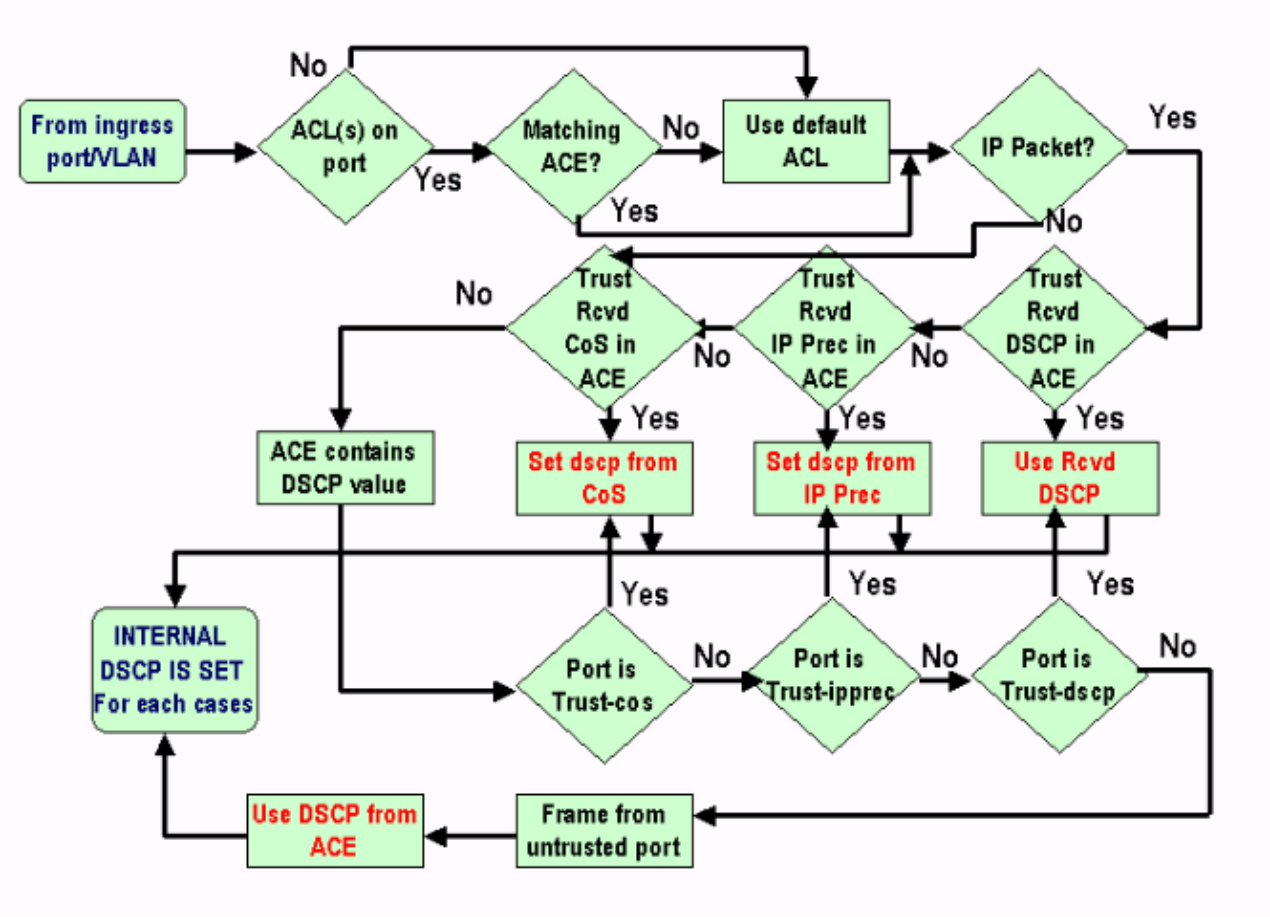

The PFC is also able to do policing. This might eventually result in a mark−down of the internal DSCP. For more details on policing, refer to the following document:

- QoS Policing on the Catalyst 6000

The following flow chart shows how the policer is applied:

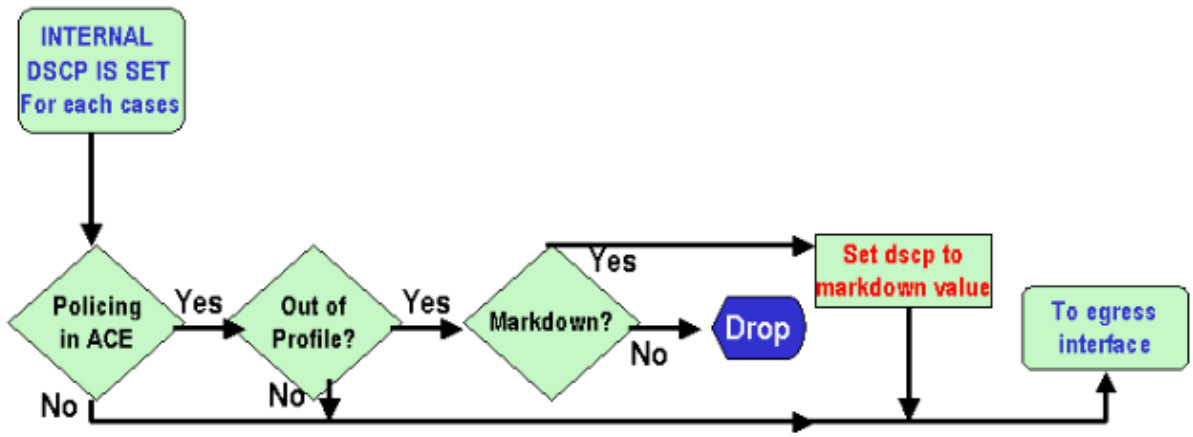

Output Port Handling

There is nothing that can be done at the egress port level to change the classification, but in this section you will mark the packet according the following rules:

- If the packet is an IPv4 packet, copy the internal DSCP assigned by the switching engine into the ToS byte of the IPv4 header.
- If the output port is configured for an ISL or dot1q encapsulation, use a CoS derived from the internal DSCP, and copy it in the ISL or dot1q frame.

Note: The CoS is derived from the internal DSCP according to a static configured by the user issuing the following command:

Note: set qos dscp−cos−map *dscp_list:cos_value*

Note: The following are the default configurations. By default the CoS will be the integer part of the DSCP divided by eight:

```
set qos dscp−cos−map 0−7:0
set qos dscp−cos−map 8−15:1
set qos dscp−cos−map 16−23:2
set qos dscp−cos−map 24−31:3
set qos dscp−cos−map 32−39:4
set qos dscp−cos−map 40−47:5
set qos dscp−cos−map 48−55:6
set qos dscp−cos−map 56−63:7
```

Once the DSCP is written into the IP header, and the CoS is derived from the DSCP, the packet will be sent to one of the output queues for output scheduling based on its CoS (even if the packet is not a dot1q or an ISL). For more information on output queue scheduling, refer to the following document:

- QoS on Catalyst 6000 Series Switches: Output Scheduling on the Catalyst 6000 with PFC or PFC 2 Using CatOS Software

The following flow chart summarizes the processing of the packet regarding marking in the output port:

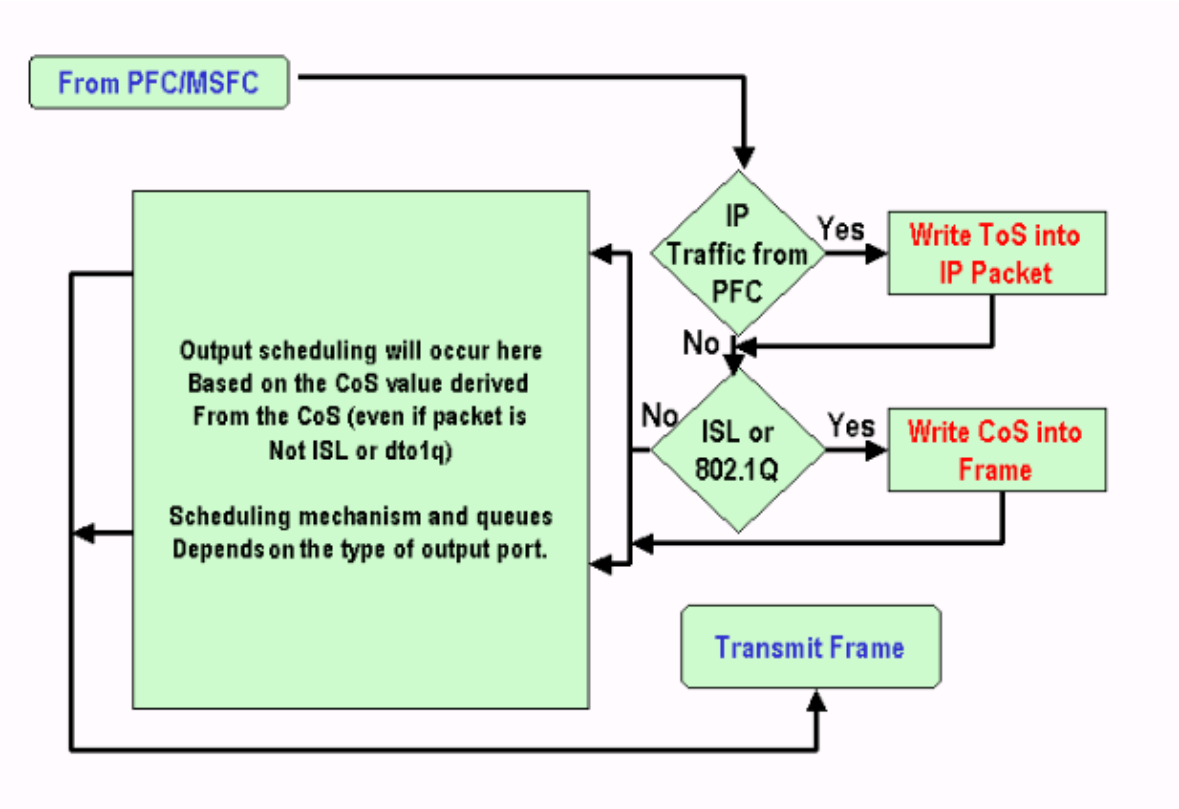

Notes and Limitations

The Default ACL

By default, the default ACL uses "dscp 0" as the classification keyword. That means that all traffic entering the switch through an untrusted port will be marked with a DSCP of "0" if QoS is enabled. You can verify the default ACL for the IP by issuing the following command:

```
Boris-1> (enable) show qos acl  info default-action  ip
set qos acl default-action
--------------------------
ip dscp 0
```

The default ACL can also be changed by issuing the following command:

set qos acl default−action ip [dscp xx | trust−CoS | trust−dscp | trust−ipprec]

trust−cos in ACL Entry Limitations

There is an additional limitation which appears when you use the trust−CoS keyword within an entry. CoS can only be trusted in an entry if the receive trust state is not untrusted. Attempting to configure an entry with trust−CoS will display the following warning:

```
Telix  (enable) set qos acl ip test_2 trust-CoS ip any any
Warning: ACL trust-CoS should only be used with ports that are also configured with port t
test_2 editbuffer modified. Use 'commit' command to apply changes.
```

This limitation is a consequence of what was seen earlier in the Input Port Handling section. As seen in the flowchart of that section, if the port is untrusted, the frame is immediately assigned the default port CoS. Therefore, the incoming CoS is not preserved and not sent to the switching engine, resulting in an inability to

trust the CoS even with a specific ACL.

Limitations of the WS−X6248−xx, WS−X6224−xx, and WS−X6348−xx Line Cards

This section only concerns the following line cards:

- WS−X6224−100FX−MT : CATALYST 6000 24 PORT 100 FX MULTIMODE
- WS−X6248−RJ−45 : CATALYST 6000 48−PORT 10/100 RJ−45 MODULE
- WS−X6248−TEL : CATALYST 6000 48−PORT 10/100 TELCO MODULE
- WS−X6248A−RJ−45 : CATALYST 6000 48−PORT 10/100, ENHANCED QOS
- WS−X6248A−TEL : CATALYST 6000 48−PORT 10/100, ENHANCED QOS
- WS−X6324−100FX−MM : CATALYST 6000 24−PORT 100FX, ENH QOS, MT
- WS−X6324−100FX−SM : CATALYST 6000 24−PORT 100FX, ENH QOS, MT
- WS−X6348−RJ−45 : CATALYST 6000 48−PORT 10/100, ENHANCED QO
- WS−X6348−RJ21V : CATALYST 6000 48−PORT 10/100, INLINE POWER
- WS−X6348−RJ45V : CATALYST 6000 48−PORT 10/100, ENH QOS, INLI NE POWER

These line cards, however, do have some additional limitations:

- At the port level, you cannot trust−dscp or trust−ipprec.
- At the port level, if the port trust state is trust−CoS, the following statements apply:
    - ♦ The receive threshold for input scheduling is enabled. In addition, the CoS in the receive packet is used to prioritize packets to access the bus.
    - ♦ The CoS will not be trusted and will not be used to derive the internal DSCP, unless you also configured the ACL for that traffic to trust−cos. In addition, it is not enough for the line cards to trust−cos on the port, you also need to have an ACL with trust−cos for that traffic.
- If the port trust state is untrusted, normal marking will happen (as with the standard case). This depends on the ACL applied to the traffic.

Any attempt to configure a trust state on one of these ports will display one of the following warning messages:

```
telix (enable) set port qos 3/24 trust trust−ipprec
Trust type trust−ipprec not supported on this port.

telix (enable) set port qos 8/4 trust trust−dscp
Trust type trust−dscp not supported on this port.

telix (enable) set port qos 3/24 trust trust−cos
Trust type trust−cos not supported on this port.
Receive thresholds are enabled on port 3/24.
Port  3/24 qos set to untrusted.
```

Summary of Classification

The tables below show the resulting DSCP classified by the following:

- The incoming port trust state.
- The classification keyword within the applied ACL.

Generic Table Summary for All Ports Except WS−X62*xx* and WS−X63*xx*

	dscp xx	**trust−dscp**	**trust−ipprec**	**trust−CoS**

ACL Keyword / Port Trust State				
Untrusted	xx (1)	Rx dscp	derived from Rx ipprec	0
trust−dscp	Rx−dscp	Rx dscp	derived from Rx ipprec	derived from Rx CoS or port CoS
trust−ipprec	derived from Rx ipprec	Rx dscp	derived from Rx ipprec	derived from Rx CoS or port CoS
trust−CoS	derived from Rx cos or port CoS	Rx dscp	derived from Rx ipprec	derived from Rx CoS or port CoS

(1) This is the only way to make a new marking of a frame.

Table Summary for WS−X62*xx* or WS−X63*xx*

ACL Keyword / Port Trust State	dscp xx	trust−dscp	trust−ipprec	trust−CoS
Untrusted	xx	Rx dscp	derived from Rx ipprec	0
trust−dscp	Not supported	Not supported	Not supported	Not supported
trust−ipprec	Not supported	Not supported	Not supported	Not supported
trust−CoS	xx	Rx dscp	derived from Rx ipprec	derived from Rx CoS or port CoS (2)

(2) This is the only way to preserve the incoming CoS for traffic coming from a 62*xx* or 63*xx* line card.

Monitoring and Verifying a Configuration

Checking the Port Configuration

The port settings and configurations can verified by issuing the following command:

show port qos *module/port*

By issuing this command, you can verify, among other parameters, the following classification parameters:

- port−based or VLAN−based
- trust port type
- ACL attached to the port

The following is a sample of this command output with the important fields regarding classification highlighted:

```
tamer (enable) show port qos 1/1
QoS is enabled for the switch.
QoS policy source for the switch set to local.


Port  Interface Type Interface Type Policy Source Policy Source
      config         runtime        config        runtime
----- -------------- -------------- ------------- -------------
 1/1    port−based     port−based      COPS           local

Port  TxPort Type  RxPort Type  Trust Type   Trust Type    Def CoS Def CoS
                                config       runtime       config  runtime
----- ------------ ------------ ------------ ------------- ------- -------
 1/1     1p2q2t       1p1q4t     untrusted     untrusted      0       0

(*)Runtime trust type set to untrusted.

Config:
Port  ACL name                         Type
----- -------------------------------- ----
 1/1  test_2                           IP

Runtime:
Port  ACL name                         Type
----- -------------------------------- ----
 1/1  test_2                            IP
```

Note: For each field, there is the configured parameter and the runtime parameter. The one which will be applied to the packet is the runtime parameter.

Checking the ACL

You can check the ACL applied and seen in previous commands by issuing the following command:

show qos acl info runtime *acl_name*

```
tamer (enable) show qos acl info run test_2
set qos acl IP test_2
------------------------------------------------
1. dscp 32 ip any host 1.1.1.1
2. trust−dscp any
```

Sample Case Studies

The following examples are sample configurations of common cases that could appear in a network.

Case 1 : Marking at the Edge

Assume you are configuring a Catalyst 6000 used as an access switch with many users connected to slot 2, which is a WS−X6348 line card (10/100M). The users can send the following:

- Normal data traffic: This is always in VLAN 100, and needs to get a DSCP of "0."
- Voice traffic from an IP phone: This is always in the voice auxiliary VLAN 101, and needs to get a DSCP of "40."
- Mission critical application traffic: This also comes in VLAN 100, and is directed to the server 10.10.10.20. This traffic needs to get a DSCP of "32."

None of this traffic is marked by the application, therefore you will leave the port as untrusted and will configure a specific ACL to classify the traffic. One ACL will be applied to VLAN 100 and one ACL will be applied to VLAN 101. You also need to configure all ports as VLAN−based. The following is an example of the resulting configuration:

```
set qos enable
set port qos 2/1−48 vlan−based

!−−− Not needed, as it is the default.

set port qos 2/1−48 trust untrusted
set qos acl ip Data_vlan dscp 32 ip any host 10.10.10.20

!−−− Not needed, because if it is not present you would
!−−− use the default ACL which has the same effect.

Set qos acl ip Data_vlan  dscp 0 ip any any
set qos acl ip Voice_vlan dscp 40 ip any any
commit qos acl all
set qos acl map Data_vlan 100
set qos acl map Voice_vlan 101
```

Case 2: Trusting in the Core with Only a Gigabit Interface

Assume you are configuring a core Catalyst 6000 with only a Gigabit interface in slot 1 and slot 2 (no 62xx or 63xx line card in the chassis). The traffic has been correctly marked previously by the access switches, therefore you do not need to make any remarking, but you need to ensure that you do trust the incoming DSCP. This is the easiest case, as all ports will be marked as trust−dscp and that should be sufficient:

```
set qos enable
set port qos 1/1−2 trust trust−dscp
set port qos 2/1−16 trust trust−dscp
...
```

Case 3: Trusting in the Core with a 62xx or 63xx Port in the Chassis

Assume you are configuring a core/distribution device with a Gigabit link on a WS−X6416−GBIC line card (in slot 2), and a 10/100 link on a WS−X6348 line card (in slot 3). You also need to trust all incoming traffic as it has been marked earlier at the access switch level. Because you cannot trust−dscp on the 6348 line card, the easiest method in this case would be to leave all the ports as untrusted and to change the default ACL to trust−dscp, as in the following example:

```
set qos enable
set port qos 2/1−16 trust untrusted
set port qos 3/1−48 trust untrusted
set qos acl default−action ip trust−dscp
```

Related Information

- **LAN Product Support**
- **LAN Switching Technology Support**

- **Technical Support − Cisco Systems**

Updated: Nov 16, 2007 Document ID: 23420