



Overview

This document describes how to configure quality of service (QoS) on the Cisco Industrial Ethernet 2000U Series (IE 2000U) and Connected Grid Switches, hereafter referred to as *switch*. The QoS feature enables you to select specific network traffic, prioritize it according to its relative importance, and use traffic-management techniques to provide preferential treatment. Implementing QoS in your network makes network performance more predictable and bandwidth utilization more effective.

Features

This chapter provides a summary of the following QoS features:

- [Modular QoS CLI, page 1-1](#)
- [Input and Output Policies, page 1-2](#)
- [Classification, page 1-2](#)
- [Table Maps, page 1-2](#)
- [Policing, page 1-3](#)
- [Marking, page 1-3](#)
- [QoS Treatment for Performance-Monitoring Protocols, page 1-3](#)
- [Congestion Management and Scheduling, page 1-4](#)
- [QoS Features for Traffic Between ESM and CGR 2010, page 1-4](#)

Modular QoS CLI

Modular QoS CLI (MQC) is the framework that you use to create traffic policies and attach these policies to interfaces. A traffic policy contains a traffic class and one or more QoS features. You use a traffic class to classify traffic, and the QoS features in the traffic policy determine how to treat the classified traffic.

The MQC structure allows you to create the traffic policy (policy map) once and then apply it to as many traffic classes as needed. You can also attach the traffic policies to as many interfaces as needed.

Although you can assign the same traffic policy to multiple interfaces, each interface can have only one traffic policy attached in the input direction and only one traffic policy attached in the output direction.

MQC also supports *parent* policy levels and *child* policy levels for traffic shaping. The QoS parent-child structure is used for specific purposes where a child policy is referenced in a parent policy to provide additional control of a specific traffic type.

Related Topics

[Modular QoS CLI, page 2-3](#)

Input and Output Policies

Policy maps are either input policy maps or output policy maps, attached to packets as they enter or leave the switch by service policies applied to interfaces. Input policy maps perform policing and marking on received traffic. Policed packets can be dropped or reduced in priority (marked down) if they exceed the maximum permitted rates. Output policy maps perform scheduling and queuing on traffic as it leaves the switch.

Related Topics

[Input and Output Policies, page 2-4](#)

[Configuring Input Policy Maps, page 2-43](#)

[Configuring Output Policy Maps, page 2-61](#)

Classification

Classification distinguishes one kind of traffic from another by examining the fields in the packet header. When a packet is received, the switch examines the header and identifies all key packet fields. A packet can be classified based on an ACL, on the DSCP, the CoS, or the IP precedence value in the packet, or by the VLAN ID.

Related Topics

[Classification, page 2-6](#)

[Using ACLs to Classify Traffic, page 2-35](#)

[Using Class Maps to Define a Traffic Class, page 2-39](#)

Table Maps

You can use table maps to manage a large number of traffic flows with a single command. You can specify table maps in **set** commands and use them as mark-down mapping for the policers. You can also use table maps to map an incoming QoS marking to a replacement marking without having to configure a large number of explicit matches and sets. Table maps are used only in input policy maps.

Related Topics

[Table Maps, page 2-14](#)

[Configuring Table Maps, page 2-41](#)

Policing

After a packet is classified, you can use policing to regulate the class of traffic. The policing function limits the amount of bandwidth available to a specific traffic flow or prevents a traffic type from using excessive bandwidth and system resources. A policer identifies a packet as in or out of profile by comparing the rate of the inbound traffic to the configuration profile of the policer and traffic class. Packets that exceed the permitted average rate or burst rate are dropped or modified (marked for further processing), depending on the policer configuration.

Policing is used primarily on receiving interfaces. You can attach a policy map with a policer only in an input service policy.

Related Topics

[Policing, page 2-15](#)

[Configuring Input Policy Maps, page 2-43](#)

Marking

You can use packet marking in input policy maps to set or modify the attributes for traffic belonging to a specific class. After network traffic is organized into classes, you use marking to identify certain traffic types for unique handling. For example, you can change the CoS value in a class or set IP DSCP or IP precedence values for a specific type of traffic. These new values are then used to determine how the traffic should be treated. You can also use marking to assign traffic to a QoS group within the switch.

Traffic marking is typically performed on a specific traffic type at the ingress port. The marking action can cause the CoS, DSCP, or precedence bits to be rewritten or left unchanged, depending on the configuration. This can increase or decrease the priority of a packet in accordance with the policy used in the QoS domain so that other QoS functions can use the marking information to judge the relative and absolute importance of the packet. The marking function can use information from the policing function or directly from the classification function.

Related Topics

[Marking, page 2-21](#)

[Configuring Input Policy Maps with Marking, page 2-54](#)

QoS Treatment for Performance-Monitoring Protocols

The QoS treatment for IP-SLA and TWAMP probes must exactly reflect the effects that occur to the normal data traffic crossing the device.

The generating device should not change the probe markings. It should queue these probes based on the configured queuing policies for normal traffic.

You can use QoS marking to set or modify the attributes of traffic from the CPU. The QoS marking action can cause the CoS, DSCP, or IP precedence bits in the packet to be rewritten or left unchanged. QoS uses packet markings to identify certain traffic types and how to treat them on the local switch and the network.

Related Topics

[QoS Treatment for Performance-Monitoring Protocols, page 2-22](#)

[Configuring QoS Marking and Queuing for CPU-Generated Traffic, page 2-73](#)

Congestion Management and Scheduling

MQC provides several related mechanisms to control outgoing traffic flow. They are implemented in output policy maps to control output traffic queues. The scheduling stage holds packets until the appropriate time to send them to one of the four traffic queues. Queuing assigns a packet to a particular queue based on the packet class, and is enhanced by the Weighted Tail Drop (WTD) algorithm for congestion avoidance. You can use different scheduling mechanisms to provide a guaranteed bandwidth to a particular class of traffic while also serving other traffic in a fair way. You can limit the maximum bandwidth that can be consumed by a particular class of traffic and ensure that delay-sensitive traffic in a low-latency queue is sent before traffic in other queues.

Related Topics

[Congestion Management and Scheduling, page 2-24](#)

[Configuring Output Policy Maps, page 2-61](#)

QoS Features for Traffic Between ESM and CGR 2010

You can implement high-priority (low-latency) traffic via the internal data path between the Ethernet Switch Module (ESM) and the host CGR 2010 router by configuring these QoS functions:

- Ingress CoS to Generic Data Frame (GDF) queue mapping
- Port shaping to adjust for differences in ingress traffic bandwidth between the the internal data path on the CGR 2010 and the internal EtherChannel on the ESM
- CoS marking for CPU-generated traffic

Related Topics

[Implementing High-Priority Traffic to the Host Router, page 2-80](#)