RBE Client Side Encapsulation with QoS

The RBE Client Side Encapsulation with QoS feature integrates routed bridged encapsulation (RBE) with quality of service (QoS) features on the Cisco 800 and 1700 series routers.

- Finding Feature Information, page 1
- Prerequisites for RBE Client Side Encapsulation with QoS, page 1
- Information About RBE Client Side Encapsulation with QoS, page 2
- Additional References, page 4
- Feature Information for RBE Client Side Encapsulation with QoS, page 5

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see Bug Search Tool and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for RBE Client Side Encapsulation with QoS

To understand the RBE Client Side Encapsulation with QoS feature, you must be familiar with routed bridge encapsulation as described in the ATM Routed Bridge Encapsulation feature module introduced in Cisco IOS Release 12.1(2)T, and with QoS class-based weighted fair queueing (CBWFQ), low latency queueing (LLQ), and class-based marking and policing as described in the Cisco IOS Quality of Service Solutions Configuration Guide.

Information About RBE Client Side Encapsulation with QoS

RBE and QoS

The RBE Client Side Encapsulation with QoS feature provides secure connectivity to an ATM bridged network in which previously a broadband access server would not forward Address Resolution Protocol (ARP) requests or perform proxy ARP, and would respond to ARPs for its own IP address only. This feature combines RBE with QoS policy-based routing to provide security to the entire network. RBE was developed to address known issues with RFC1483 bridging such as broadcast storms and security.

From the network point of view, the ATM connection looks like a routed connection. Data traffic is received as RFC1483 packets, but are actually RFC1483 Ethernet or IEEE 802.3 frames. Instead of bridging the Ethernet or IEEE 802.3 frame, as in the case of regular RFC1483 bridging, the router routes on the Layer 3 header. With the exception of some cursory checks, the bridge header is ignored.

From an operational point of view, the router operates as if the routed-bridge interface were connected to an Ethernet LAN. RBE functions in the same way as half-bridging, except that it operates only over ATM. The operation is described in two ways: packets originating from the customer premises and packets destined for the customer premises.

For packets originating from the customer premises, the Ethernet header is skipped and the destination IP address is examined. If the destination IP address is in the route cache, the packet is fast switched to the outbound interface. If the destination IP address is not in the route cache, the packet is queued for process switching. In the process switch mode, the outbound interface through which the packet must be routed is found when software routines identifies it in the routing table. After the outbound interface is identified, the packet is routed on that interface. This routing occurs without the requirement for a bridge group or bridge group virtual interface (BVI).

For packets destined for the customer premises, the destination IP address of the packet is examined first. The destination interface is determined from the IP routing table. Next, the router checks the ARP table associated with that interface for a destination MAC address to place in the Ethernet header. If none is found, the router generates an ARP request for the destination IP address. The ARP request is forwarded to the destination interface only. This is in contrast to bridging, in which the ARP request is sent to all interfaces in the bridge group.

The RBE Client Side Encapsulation with QoS feature provides the ability, as an example, to pass packets to the network with a destination MAC address of 0.0.0.0 to populate the ARP on return traffic.

Low-Latency Queueing and Class-Based Weighted Fair Queueing

Low-latency queueing (LLQ) brings strict priority queueing to CBWFQ. Strict priority queueing allows delay-sensitive data such as voice to be dequeued and sent before packets in other queues are dequeued, thereby giving delay-sensitive data preferential treatment over other traffic.

Without LLQ, CBWFQ provides weighted fair queueing based on defined classes with no strict priority queue available for real-time traffic. CBWFQ allows you to define traffic classes and then assign characteristics to that class. For example, you can designate the minimum bandwidth delivered to the class during congestion.

For CBWFQ, the weight for a packet belonging to a specific class is derived from the bandwidth you assigned to the class when you configured it. Therefore, the bandwidth assigned to the packets of a class determines the order in which packets are sent. All packets are serviced fairly based on weight; no class of packets may

be granted strict priority. This scheme poses problems for voice traffic that is largely intolerant of delay, especially variation in delay. For voice traffic, variations in delay introduce irregularities of transmission manifesting as jitter in the heard conversation.

The LLQ feature provides strict priority queueing for CBWFQ, reducing jitter in voice conversations. Configured by the **priority** command, LLQ enables use of a single, strict priority queue within CBWFQ at the class level, allowing you to direct traffic belonging to a class to the CBWFQ strict priority queue. To enqueue class traffic to the strict priority queue, you configure the **priority** command for the class after you specify the named class within a policy map. (Classes to which the **priority**command is applied are considered priority classes.) Within a policy map, you can give one or more classes priority status. When multiple classes within a single policy map are configured as priority classes, all traffic from these classes is enqueued to the same, single, strict priority queue.

One of the ways in which the strict priority queueing used within CBWFQ differs from its use outside CBWFQ is in the parameters it takes. Outside CBWFQ, by using the **ip rtp priority** command, you specify the range of User Datagram Protocol (UDP) ports whose voice traffic flows are to be given priority service. Using the **priority** command, because you can configure the priority status for a class within CBWFQ, you are no longer limited to a UDP port number to stipulate priority flows. Instead, all of the valid match criteria used to specify traffic for a class now apply to priority traffic. These methods of specifying traffic for a class include matching on access lists, protocols, and input interfaces. Moreover, within an access list you can specify that traffic matches are allowed based on the IP Differentiated Services Code Point (DSCP) value that is set using the first six bits of the Type of Service (ToS) byte in the IP header.

Class-Based Marking

In a traffic stream, a packet is classified based on the content of some portion of the packet header. The Behavior Aggregate (BA) classifier classifies packets based on the DSCP only. The Multi-field (MF) classifier selects packets based on the the value of the combination of one or more header fields, such as source address, destination address, Differentiated Services (DS) field (a replacement header field that supersedes the existing definitions of the IPv4 ToS octet and the IPv6 traffic class octet), protocol ID, source port and destination port numbers, and other information such as incoming interface and outgoing interface. The packet can be marked by a packet marker to set the DS field of a packet to a particular code point, adding the marked packet to a particular DS behavior aggregate.

Class-Based Policing

Class-based policing is applied when you attach a traffic policy containing a class-based policing configuration to an interface. A traffic policy is configured using the Modular Quality of Service Command-Line Interface (Modular QoS CLI).

Class-based policing allows you to control the maximum rate of traffic transmitted or received on an interface. Class-based policing is often configured on interfaces at the edge of a network to limit traffic into or out of the network. In most class-based policing configurations, traffic that falls within the rate parameters is transmitted, whereas traffic that exceeds the parameters is dropped or transmitted with a different priority.

Packet marking allows you to partition your network into multiple priority levels or classes of service (CoS). A packet is marked and these markings can be used to identify and classify traffic for downstream devices. In some cases, such as ATM Cell Loss Priority (CLP) marking or Frame Relay Discard Eligibility (DE) marking, the marking is used to classify traffic.

Use class-based policing to set the IP precedence or DSCP values for packets entering the network. Networking devices within your network can then use the adjusted IP precedence values to determine how the traffic

should be treated. For example, the Weighted Random Early Detection (WRED) feature uses the IP precedence values to determine the probability that a packet will be dropped.

Use class-based policing to assign packets to a QoS group. The router uses the QoS group to determine how to prioritize packets within the router.

The Single Rate Three Color Marker (srTCM) meters an IP packet stream and marks its packets either conform, exceed, or violate. Marking is based on a Committed Information Rate (CIR) and two associated burst sizes, a Committed Burst Size (CBS) and an Excess Burst Size (EBS). A packet is marked "conform" if it does not exceed the CBS, marked "exceed" if it does exceed the CBS but not the EBS, and marked "violate" otherwise.

Additional References

The following sections provide references related to the RBE Client Side Encapsulation with QoS feature.

Related Documents

Related Topic	Document Title
Routed bridge encapsulation	• Configuring Broadband Access: PPP and Routed Bridge Encapsulation Configuring PPP over ATM " chapter in the Cisco IOS Wide-Area Networking Configuration Guide • ATM Routed Bridge Encapsulation feature module
Policy-based routing with QoS	• Class-Based Weighted Fair Queueing and Low Latency Queueing sections in the Cisco IOS Quality of Service Solutions Configuration Guide

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

Feature Information for RBE Client Side Encapsulation with QoS

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to . An account on Cisco.com is not required.

Table 1: Feature Information for RBE Client Side Encapsulation with QoS

Feature Name	Releases	Feature Information
RBE Client Side Encapsulation with QoS	12.4(2)T	The RBE Client Side Encapsulation with QoS feature integrates routed bridged encapsulation (RBE) with quality of service (QoS) features on the Cisco 800 and 1700 series routers. The following commands were introduced or modified: **atm route-bridged.**