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.

History for the RBE Client Side Encapsulation with QoS Feature

Release	Modification
12.4(2)T	This feature was introduced.

Finding Support Information for Platforms and Cisco IOS Software Images

Use Cisco Feature Navigator to find information about platform support and Cisco IOS software image support. Access Cisco Feature Navigator at http://www.cisco.com/go/fn. You must have an account on Cisco.com. If you do not have an account or have forgotten your username or password, click **Cancel** at the login dialog box and follow the instructions that appear.

Contents

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*.

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Information About RBE Client Side Encapsulation with QoS

The RBE Client Side Encapsulation with QoS feature is described in the following sections:

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

Command Reference

This section documents the modified **atm route-bridged** command.

atm route-bridged

To configure an interface to use the ATM routed bridge encapsulation (RBE), use the **atm route-bridged** command in ATM subinterface configuration mode.

atm route-bridged *protocol*

Syntax Description

protocol	Protocol to be route-bridged. IP and IPv6 are the only protocols supported for this command.

Defaults

ATM routed bridge encapsulation is not configured.

Command Modes

ATM subinterface configuration

Command History

Release	Modification
12.0(5)DC	This command was introduced.
12.1(2)T	This command was integrated in Cisco IOS Release 12.1(2)T.
12.3(4)T	The **ipv6** keyword was added to support RBE of IPv6 packets as specified in RFC 1483.
12.4(2)T	This command was updated to work with QoS policy-based routing in Cisco IOS Release 12.4(2)T.

Usage Guidelines

Use this command to configure RBE on an ATM interface. The **atm route-bridged** command can also be used to integrate RBE with quality of service (QoS) features on the Cisco 800 and 1700 series routers.

Routing of IPv6 and IP Packets

IP and IPv6 packets can be routed using RBE only over ATM point-to-point subinterfaces.

Routing of IP packets and IPv6 half-bridging, bridging, PPP over Ethernet (PPPoE), or other Ethernet 802.3-encapsulated protocols can be configured on the same subinterface.

Router Advertisements with IPv6

Router advertisements are suppressed by default. For stateless autoconfiguration, router advertisements must be allowed with the **no ipv6 nd suppress-ra** command. For static configuration, router advertisement is not required; however, the aggregator should either have the RBE interface on the same subnet as the client or have a static IPv6 route to that subnet through the RBE interface.

Examples

IP Encapsulation Example

The following example configures ATM routed bridge encapsulation on an interface:

```
interface atm 4/0.100 point-to-point
 ip address 172.16.5.9 255.255.255.0
 atm route-bridged ip
 pvc 0/32
```

IPv6 Encapsulation Example

The following example shows a typical configuration on an RBE interface to allow routing of IPv6 encapsulated Ethernet packets. IPv6 packets sent out of the subinterface are encapsulated over Ethernet over the RBE interface.

```
interface ATM1/0.1 point-to-point
 ipv6 enable
 ipv6 address 3FEE:12E1:2AC1:EA32::/64
 no ipv6 nd suppress-ra
 atm route-bridged ipv6
 pvc 1/101
```

In this example, the **ipv6 enable** command allows the routing of IPv6 packets. The **ipv6 address** command specifies an IPv6 address for the interface and an IPv6 prefix to be advertised to a peer. The **no ipv6 nd supress-ra** command enables router advertisements on the interface.

IPv6 Routing and Bridging of Other Traffic Example

The following example shows a configuration in which IPv6 packets are routed and all other packets are bridged.

```
interface ATM1/0.1 point-to-point
 ipv6 enable
 ipv6 address 3FEE:12E1:2AC1:EA32::/64
 atm route-bridged ipv6
 bridge-group 1
 pvc 1/101
```

IP and IPv6 Routing with Bridging of Other Protocols Example

IP and IPv6 routing can be configured on the same interface as shown in this example. All other packets are bridged. PPP over Ethernet (PPPoE) could also be configured on this same interface.

```
interface ATM1/0.1 point-to-point
 ipv6 enable
 ipv6 address 3FEE:12E1:2AC1:EA32::/64
 ip address 10.0.0.1 255.255.255.0
 atm route-bridged ipv6
 atm route-bridged ip
 bridge-group 1
 pvc 1/101
```

Static Configuration Example

The following example shows the IPv6 static route configured. Unlike IP, the IPv6 interface on an aggregator is always numbered and, minimally, has a link local IPv6 address.

```
Router# configure terminal
Router(config)# ipv6 route 3FEE:12E1:2AC1:EA32::/64 atm1/0.3
Router(config)# end
```

show ipv6 interface Example

Notice in this **show ipv6 interface** output display that each RBE link has its own subnet prefix. Unlike proxy Address Resolution Protocol (ARP) in IPv4 RBE configurations, the aggregator does not require proxy ND in IPv6 RBE deployments.

```
Router# show ipv6 interface atm1/0.1

ATM1/0.1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::203:FDFF:FE3B:B400
  Global unicast address(es):
    3FEE:12E1:2AC1:EA32::, subnet is 3FEE:12E1:2AC1:EA32::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:0
    FF02::1:FF3B:B400
  MTU is 4470 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses
```

Integrated Class-Based Weighted Fair Queueing and RBE on ATM Example

The following partial example configures a single PVC using AAL5SNAP encapsulation and class-based routing for traffic shaping on the interface where RBE is enabled. The following Class-Based Weighted Fair Queueing (CBWFQ) parameters are configured: access-list with different IP precedence, class map, policy map, and service policy. Different bandwidth classes are configured in the same policy.

The RBE base configuration is as follows:

```
interface FastEthernet0
 ip address 172.22.1.1 255.255.0.0
!
interface ATM0.1 point-to-point
 ip address 10.1.1.5 255.255.255.252
 atm route-bridged ip
 pvc 88/800
  encapsulation aal5snap
!
interface ATM0.1 point-to-point
 ip address 10.1.1.1 255.255.255.252
 atm route-bridged ip
 pvc 99/900
  encapsulation aal5snap
!
interface ATM0.1 point-to-point
 ip address 172.18.0.1 255.0.0.0
 pvc 100/1000
!
router eigrp 100
 network 10.1.0.0
 network 172.18.0.0
 network 172.22.0.0
.
.
.
```

The CBWFQ configuration is as follows:

```
class-map match-all voice
 match access-group 105
!
policy-map voicedatapolicy
 class voice
  bandwidth 200
 class class-default
  fair-queue
  random-detect
!
interface Ethernet0
 ip address 172.25.1.1 255.0.0.0
 hold-queue 600 in
 hold-queue 100 out
!
interface ATM0
 no ip address
 no atm ilmi-keepalive
 dsl operating-mode auto
!
interface ATM0.1 point-to-point
 ip address 10.2.3.4 255.255.255.0
 atm route-bridged ip
 pvc 1/42
  protocol ip 10.2.3.5 broadcast
  vbr-nrt 300 300
  encapsulation aal5snap
  service-policy output voicedatapolicy
.
.
.
```

Related Commands

Command	Description
no ipv6 nd suppress-ra	Suppresses IPv6 router advertisement transmissions on a LAN interface.