

Configuring Frame Relay Traffic Shaping on 7200 Routers and Lower Platforms

Document ID: 6151

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

Background Information

Configure

- Network Diagram

- Configurations

Verify

- show Commands

- Configurable Parameters

- Non-Configurable Parameters

Troubleshoot

Related Information

Introduction

This document provides a sample configuration for Frame Relay traffic shaping.

Prerequisites

Requirements

There are no specific prerequisites for this document.

Components Used

Frame-relay traffic shaping has been supported since Cisco IOS® Software Release 11.2.

It is supported on Cisco 7200 routers and lower platforms. Distributed Traffic Shaping is supported on Cisco 7500 routers, 7600 routers, and FlexWAN module.

Conventions

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

Background Information

Common implementations of Frame Relay traffic shaping are:

1. **High speed to low speed circuit mismatches:** There are two possibilities here:

- ◆ The hub site has a T1 line into the cloud, while the remote site has a lower speed (56 Kbps).

In this case, you need to rate-limit the hub site so that it does not exceed the remote side access rate.

- ◆ The hub site has a single T1 line into the cloud, while the remote sites also have a full T1 line into the cloud, connecting to the same hub site. In this case, you need to rate-limit the remote sites so as to not overrun the hub.
2. **Oversubscription:** For example, if the guaranteed rate on a permanent virtual circuit (PVC) is 64 Kbps and the access rate is 128 Kbps on both ends, it is possible to burst above the guaranteed rate when there is no congestion and fall back to the guaranteed rate when there is congestion.
 3. **Quality of Service:** For implementing FRF.12 fragmentation or low latency queuing features to achieve better quality of service, see VoIP over Frame Relay with Quality of Service.

Note: The access rate is the physical line speed of the interface connecting to the Frame Relay. The guaranteed rate is the committed information rate (CIR) the Telco has given for the PVC. Setting the CIR or minCIR at the access rate should be avoided, because it may result in output drops, causing traffic to throttle. The reason for this is that the shape rate does not take into account the overhead bytes of the flag and Cyclic Redundancy Check (CRC) fields. So, shaping at line rate is actually oversubscribing, and will cause interface congestion. Shaping at the access rate is not recommended. You should always shape the traffic at 95 percent of the access rate. More generally, the aggregate shaped rate should be no more than 95 percent of the access rate.


Configure

In this section, you are presented with the information to configure the features described in this document.

Note: To find additional information on the commands used in this document, use the IOS Command Lookup tool

Network Diagram

This document uses this network setup:



In the above example, we have the following values:

- HUB – access rate = 192 Kbps, guaranteed rate = 32Kbps
- REMOTE – access rate = 64Kbps, guaranteed rate = 32Kbps

Here, we are implementing traffic shaping at both ends so that the average transmit rate is 64Kbps. If needed, the HUB can burst above this. In case of congestion, it can drop down to 32Kbps at the minimum. Congestion notification from the cloud is via backward explicit congestion notification (BECN). Hence, the shaping is configured to adapt to BECN.

Note: Frame-relay traffic shaping is enabled on the main interface, and it applies to all data link connection identifiers (DLCIs) under that interface. We cannot enable traffic shaping only for a particular DLCI or subinterface under the main interface. If a certain DLCI has no map class attached to it, and traffic shaping is

enabled on the main interface, the DLCI is assigned a default map-class with CIR = 56000.

Configurations


This document uses these configurations:

- Hub
- Remote

Hub
<pre>interface Serial0/0 no ip address encapsulation frame-relay no fair-queue frame-relay traffic-shaping !--- Apply traffic shaping to main interface (step 3). interface Serial0/0.1 point-to-point ip address 10.1.1.1 255.255.255.0 frame-relay interface-dlci 16 frame-relay class cisco !--- Apply map class to the DLCI / subinterface (step 2). ! ! !--- Configure map class parameters (step 1). map-class frame-relay cisco frame-relay cir 64000 frame-relay mincir 32000 frame-relay adaptive-shaping becn frame-relay bc 8000 frame-relay be 16000 !</pre>

Remote
<pre>interface Serial0/0 no ip address encapsulation frame-relay no fair-queue frame-relay traffic-shaping ! interface Serial0/0.1 point-to-point ip address 10.1.1.2 255.255.255.0 frame-relay interface-dlci 16 frame-relay class cisco ! map-class frame-relay cisco frame-relay cir 64000 frame-relay mincir 32000 frame-relay adaptive-shaping becn frame-relay bc 8000 !</pre>

This diagram shows traffic being sent out of the HUB router:



BECNs indicating congestion received starting at 300ms mark
As a result router throttles down until BECNs stop or reaches Mincir

Assuming that the traffic is sent with a burst of 80000 bits, this is sent out of the PVC in 8 Tc intervals (125 msec each). We can achieve this because, in the first interval, the credit available is $Bc + Be = 8000 + 16000 = 24000$ bits. This means that the rate is $24000 \text{ bits} / 125 \text{ msec} = 192 \text{ Kbps}$.

In the next seven intervals it is only $Bc = 8000$ bits. Hence the rate is $8000 / 125 \text{ msec} = 64 \text{ Kbps}$.

For example, if we receive a burst of 88000 bits, we cannot send all this traffic in 8 Tc intervals. The final 8000 bits will be sent in the 9th Tc interval. Thus, this traffic is delayed by the traffic shaping mechanism.

Verify

This section provides information you can use to confirm your configuration is working properly.

show Commands

Certain **show** commands are supported by the Output Interpreter Tool (registered customers only) , which allows you to view an analysis of **show** command output.

Use the **show frame relay pvc <dldci>** command to view the configuration details:

```
Hub#show frame relay pvc 16
PVC Statistics for interface Serial0/0 (Frame Relay DTE)
DLCI = 16, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial0/0.1
  input pkts 8743          output pkts 5          in bytes 2548330
  out bytes 520           dropped pkts 0         in FECN pkts 0
  in BECN pkts 0         out FECN pkts 0       out BECN pkts 0
  in DE pkts 0           out DE pkts 0
  out bcast pkts 0       out bcast bytes 0
  Shaping adapts to BECN
  pvc create time 6d01h, last time pvc status changed 6d01h

cir 64000 bc 8000 be 16000 byte limit 3000 interval 125
mincir 56000 byte increment 1000 Adaptive Shaping BECN
pkts 5 bytes 170 pkts delayed 0 bytes delayed 0
shaping inactive
traffic shaping drops 0

Queueing strategy: fifo
```

Output queue 0/40, 0 drop, 0 dequeued

shaping inactive / active

This shows, in real time, whether the traffic shaping mechanism has been activated or not. Traffic shaping is active in the following scenarios:

1. BECNs are received, and DLCI has been configured to shape to BECNs.
2. The number of data bytes to transmit out of an interface are more than the available credit (byte limit) in a given interval (Tc).
3. FRF.12 fragmentation has been configured, and packets are waiting to be fragmented.

pkts delayed / bytes delayed

This shows the number of packets and bytes that have been delayed due to activation of the traffic shaping mechanism. This mainly applies if the number of bytes to be transmitted exceeds the available credit per interval, or if packets need to be fragmented (FRF.12). These packets and bytes are stored in the shaping queue (allocated per VC) and then transmitted in subsequent intervals when there is enough available credit.

traffic shaping drops

This shows the number of drops in the shaping queue. Bytes are first delayed by the shaping mechanism and stored in this queue. If the queue fills up, then packets are dropped. By default, the queue type is FCFS (First Come First Serve) or FIFO, but can be changed to WFQ, PQ, CQ, CBWFQ, or LLQ. See the Related Information section for documents on configuring fancy queueing mechanisms on Frame Relay PVCs.

Configurable Parameters

frame relay cir

The average rate you want to send traffic out on a given PVC in bps. This is generally higher than the guaranteed rate but less than the access rate (AR). It equals the guaranteed rate only if:

1. The service provider does not allow you to send above guaranteed rate.
2. The physical line rate on the interface is same as the guaranteed rate.
3. There are Voice (voice over IP [VOIP] or voice over Frame Relay [VOFR]) packets on this PVC, therefore you cannot afford dropped packets for quality or service.

The value of the CIR is 56000 bps is by default.

frame relay mincir

The actual guaranteed rate obtained from service provider in bps. This value should be the minimum rate you should drop to in the event of congestion (dropping below this rate implies you are not getting the bandwidth you are paying for). In certain cases (listed above) the **mincir** and **cir** values must be the same. The value of **mincir** is half of the CIR value in bps by default.

frame relay bc

The amount of data to send per each Tc interval in bits. Ideally for data PVCs $Bc = CIR/8$ so that $Tc = 125msec$. Cisco IOS recalculates the FRTS parameter when Bc is greater than 10,000 bytes. If we are doing voice on the PVC, then $Bc = CIR/100$ is preferable, so that the interval $Tc = 10msec$ (as voice packets cannot tolerate a longer delay). The value of Bc by default is shown as the CIR in bits in the output of the **show traffic-shape** command. However, internally, a different value is assigned to ensure optimal performance.

This value is shown in the "Increment Bytes" column in the **show traffic-shape** output. A value of $bc=CIR$ equates to a Tc of 1 second. Depending on how traffic arrives at the shaper, the router would have to stop transmission for close to 1 second if the burst was exhausted immediately at the start of the interval. Thus, the shaper assigns a different internal value which still allows for the configured Bc over the original Tc , only we will do it in a number of small bursts instead of one large burst.

frame relay be

The amount of excess data allowed to be sent during first Tc interval in bits once credit is built up. Configure Be only if the Frame Relay CIR value is less than the AR. For PVCs carrying voice packets, the Be must be set to zero to ensure best possible voice quality. The router only bursts (Be) when there are tokens in the token bucket. The token bucket does not accrue tokens unless the amount of traffic being sent out is less than the CIR. The router can only burst for the first Tc , after which the token bucket is empty. The value of Be by default is zero bits.

frame relay adaptive-shaping becn

Implies that the PVC adapts the rate of transmit in response to the BECNs received. The behaviour is as below:

- If the PVC receives any BECNs during the current time interval (it doesn't matter if this is one or 1000) the transmit rate is decreased by 25 percent or to **mincir** and stops if **mincir** configured value is more than 75% of the **cir** value.
- It continues to drop with each BECN (limit one drop per time interval) until the traffic rate gets to the **mincir** (guaranteed rate) where it stops.
- Once the traffic rate has decreased, it must allow 16 time intervals of receiving no BECNs before starting to increase traffic again. The amount it increases by is the byte limit that shows up in the **show frame pvc x** output divided by 16. This increase occurs only if traffic shaping is active. Thus, it takes much longer to get back to the CIR than it did to drop to **mincir**.

Non-Configurable Parameters

interval (Tc)

The interval during which you send the Bc bits in order to maintain the average rate of the CIR in seconds.

$$Tc = Bc / CIR \text{ in seconds}$$

The range for Tc is between 10 ms and 125 ms. The router internally calculates this value based on the CIR and Bc values in the map class. If Bc/CIR is more than or equal to 125 msec, it uses the internal Tc value. If Bc/CIR is less than 125 ms, it uses the Tc calculated from that equation.

byte increment

The actual number of committed bytes sent per Tc . We can calculate this using the following formula:

$$CIR * Tc / 8$$

byte limit

The actual number of bytes sent in the first Tc . We can calculate this using the following formula:

$$\text{byte increment} + Be/8 \text{ (measured in bytes)}$$

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Related Information

- [Configuring Frame Relay and Frame Relay Traffic Shaping](#)
 - [show Commands for Frame Relay Traffic Shaping](#)
 - [Configuring Generic Traffic Shaping](#)
 - [Configuring and Troubleshooting Frame Relay](#)
 - [Understanding Queueing on Frame Relay Router Interfaces](#)
 - [Configuring Packet Marking on Frame Relay PVCs](#)
 - [Configuring CBWFQ on Frame Relay PVCs](#)
 - [Priority Queueing Options on Frame Relay Virtual Circuits](#)
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: May 31, 2005

Document ID: 6151
