Frame Relay Frequently Asked Questions

Document ID: 14168

Contents

Introduction

Frame Relay is a high–performance WAN protocol that operates at the physical and data link layers of the Open System Interconnection (OSI) reference model. It is described as a streamlined version of X.25 and is commonly used over reliable WAN connections. This document discusses some of the frequently asked questions about Frame Relay.

General

Q. Why am I unable to ping my own interface address?

A. You cannot ping your own IP address on a multipoint Frame Relay interface. To make a ping successful on a serial interface, an Internet Control Message Protocol (ICMP) echo request packet must be sent, and an ICMP echo reply packet must be received. Pings to your own interface address are successful on point–to–point subinterfaces or high–level data link control (HDLC) links because the router on the other side of the link returns the ICMP echo and echo reply packets.

The same principle also applies with multipoint (sub) interfaces. To successfully ping your own interface address, another router must send back the ICMP echo request and the echo reply packets. Because multipoint interfaces can have multiple destinations, the router must have Layer 2 (L2) to Layer 3 (L3) mapping for every destination. Because mapping is not configured for our own interface address, the router does not have any L2 to L3 mapping for its own address and does not know how to encapsulate the packet. That is, the router does not know which data–link connection identifier (DLCI) to use to send echo request packets to its own IP address resulting in encapsulation failure. To be able to ping its own interface address, a static mapping must be configured pointing towards another router over the Frame Relay link which can send back the ICMP echo request and reply packets.

Q. Why am I unable to ping from one spoke to another spoke in a hub and spoke configuration using multipoint (sub) interfaces?

A. You cannot ping from one spoke to another spoke in a hub and spoke configuration using multipoint interfaces because the mapping for the other spoke's IP address is not done automatically. Only the hub's address is automatically learned by way of Inverse Address Resolution Protocol (INARP). If you configure a static map using the **frame–relay map**

command for the IP address of another spoke to use the local data−link connection identifier (DLCI), you can ping the address of the other spoke.

Q. What is the Frame Relay broadcast queue?

A. The Frame Relay broadcast queue is a major feature used in medium to large IP or Internet Package Exchange (IPX) networks in which routing and Service Advertising Protocol (SAP) broadcasts must flow across the Frame Relay network. The broadcast queue is managed independently of the normal interface queue, has its own buffers, and a configurable size and service rate. Due to timing sensitivities, Spanning Tree Protocol (STP) Bridge Protocol Data Units (BPDUs) are not transmitted using the broadcast queue.

Q. How many data−link connection identifier (DLCI)s can an interface support?

A. This question is similar to the question of how many PCs can you put on an Ethernet. In general, you can put a lot more than you should, given performance and availability constraints. When dimensioning a router in a large network, consider the following issues:

- *DLCI address space*: With a 10−bit address, about 1000 DLCIs can be configured on a single physical link. Because certain DLCIs are reserved (vendor implementation−dependent), the maximum is about 1000. The range for Cisco Local Management Interface (LMI) is 16−1007. The range for American National Standards Institute and International Telecommunication Union Telecommunication Standardization Sector (ANSI/ITU−T) is 16−992. These DLCIs carry users' data.
- *LMI status update*: The LMI protocol requires that all permanent virtual circuit (PVC) status reports fit into a single packet and generally limits the number of DLCIs to less than 800, depending on the maximum transmission unit (MTU) size. This yields the following for a configured interface MTU of 4000 bytes:

$$\text{Max DLCI's} \cong \frac{\text{MTU bytes} - 20}{5 \, \text{bytes}/\text{DLCI}}$$

$$\text{Max DLCI's} \cong \frac{4000 - 20}{5} = 796 \frac{\text{DLCI's}}{\text{interface}}$$

Note: Default MTU on serial interfaces is 1500 bytes, yielding a maximum of 296 DLCIs per interface.

- *Broadcast Replication*: When the router is sending, it must replicate the packet on each DLCI, which causes congestion on the access link. The broadcast queue reduces this problem. In general, the network should be designed to keep the routing update load to below 20 percent of the access line's speed. It is also important to consider the memory requirements for the broadcast queue. A good technique to reduce this restriction is to use the default route or extend the update timers.
- *User Data Traffic*: The number of DLCIs depends upon the traffic on each DLCI and on performance requirements. In general, Frame Relay accesses should run at lower loads than router−to−router links because the prioritization capabilities are usually not as strong. In general, the marginal cost of increasing access link speed is lower than for dedicated lines.

For estimates on the practical number of DLCIs supported on Cisco router platforms, refer to the DLCI Limitations section of Comprehensive Guide to Configuring and Troubleshooting Frame Relay.

Q. Can I use IP unnumbered with Frame Relay?

A. If you do not have the IP address space to use many subinterfaces, you can use IP unnumbered on each subinterface. You need to use static routes or dynamic routing for your traffic to get routed. And you must use point−to−point subinterfaces. For more information, refer to the Unnumbered IP over a Point−to−Point Subinterface Example section of Configuring Frame Relay.

Q. Can I configure a Cisco router to act as a Frame Relay switch?

A. Yes. You can configure Cisco routers to function as Frame Relay data communication equipment (DCE) or network−to−network interface (NNI) devices (Frame Relay switches). A router can also be configured to support hybrid data terminal equipment/data communication equipment/permanent virtual circuit (DTE/DCE/PVC) switching. . For more information, refer to the Configuring Frame Relay section of the Cisco IOS Wide−Area Networking Configuration Guide, Release 12.1.

Q. Can I bridge traffic over a Frame Relay link?

A. Yes. On multipoint interfaces, Frame Relay map statements must be configured using the **frame−relay map bridge** command to identify permanent virtual circuits (PVCs) for bridged traffic. Spanning(remove hyphen)Tree Protocol (STP) Bridge Protocol Data Units (BPDUs) are passed at regular intervals depending on the bridging protocol configured.

Q. Is a special configuration necessary to connect Cisco routers to other vendor devices over Frame Relay?

A. Cisco routers use proprietary Frame Relay encapsulation by default. The Internet Engineering Task Force (IETF) encapsulation format must be specified to interact with other vendor devices. The IETF encapsulation can be specified on an interface or per data−link connection identifier (DLCI) basis. For more information, refer to the Frame Relay Configuration Examples section of Configuring Frame Relay, in the Cisco IOS Wide−Area Networking Configuration Guide, Release 12.1.

Q. What is Frame Relay AutoInstall and how does it work? Is an additional configuration required?

A. AutoInstall allows you to configure a new router automatically and dynamically. The AutoInstall procedure involves connecting a new router to a network in which an existing router is preconfigured, turning on the new router, and enabling it with a configuration file that is downloaded from a TFTP server. For more information, refer to Using Configuration Tools.

To support AutoInstall over a link on which the existing router is configured with a point−to−point subinterface, the **frame−relay interface−dlci** command requires additions. The additional information provided with the **frame−relay interface−dlci** command is used to respond to the Bootstrap Protocol (BOOTP) request of the remote router. The addition of **protocol ip***ip−address* to the command indicates the IP address of the main interface of a new router or access server onto which a router configuration file is to be installed over a Frame Relay network. Use this option only when the device acts as the BOOTP server for automatic installation over Frame Relay.

To support AutoInstall over a link on which the existing router is configured with a multipoint (sub) interface, the **frame–relay map** command should be configured on the existing router, mapping the IP address of the new router to the local data–link connection identifier (DLCI) used for connecting to the new router.

Apart from this, the Frame Relay (sub) interface of the existing router should be configured with the **ip helper–address** command pointing to the IP address of the TFTP server.

Q. Is Frame Relay Inverse Address Resolution Protocol (IARP) on by default? The inverse–arp command does not show up in the configuration.

A. Yes.

Q. Can Frame Relay Inverse Address Resolution Protocol (IARP) work without Local Management Interface (LMI)?

A. No. It uses LMI to determine which permanent virtual circuits (PVCs) to map.

Q. Under what Local Management Interface (LMI) conditions does a Cisco router not send packets over the data–link connection identifier (DLCI)?

A. When the permanent virtual circuit (PVC) is listed as inactive or deleted.

Q. Will a Cisco router process and map an Inverse Address Resolution Protocol (IARP) if it comes across while a data–link connection identifier (DLCI) is down?

A. Yes, but the router will not use it until the DLCI is active.

Q. When implementing a show frame map command, data–link connection identifiers (DLCIs) are `defined and active`. This can occur when the DLCIs are not working. What does `defined and active` mean?

A. The message `defined and active` tells you that the DLCI can carry data and that the router at the far end is active.

Q. Can I change subinterfaces from point–to–point to multipoint or the reverse?

A. No, after a specific type of subinterface is created, it cannot be changed without a reload. For example, you cannot create a multipoint subinterface Serial0.2, and change it to point–to–point. To change it, delete the existing subinterface and reload the router or create another subinterface. When a subinterface is configured, an interface descriptor block (IDB) is defined by the Cisco IOS® Software. IDBs defined for subinterfaces cannot be changed without a reload. Subinterfaces that are deleted with the **no interface** command are shown as deleted by issuing the **show ip interface brief** command.

Q. What does `illegal serial line type xxx` mean?

A. This message is displayed if the encapsulation for the interface is Frame Relay (or High–Level Data Link Control [HDLC]) and the router attempts to send a packet containing an unknown packet type.

Performance

Q. What are Forward Explicit Congestion Notification (FECN) and Backward Explicit Congestion Notification (BECN) packets? How do they affect performance?

A. This congestion notification is accomplished by changing a bit in the address field of a frame as it traverses the Frame Relay network. Network DCE devices (switches) change the value of the FECN bit to one on packets traveling in the same direction as the data flow. This notifies an interface device (DTE) that congestion avoidance procedures should be initiated by the receiving device. BECN bits are set in frames that travel the opposite direction of the data flow to inform the transmitting DTE device of network congestion.

Frame Relay DTE devices may choose to ignore FECN and BECN information or may modify their traffic rates based on FECN and BECN packets received. The **frame–relay adaptive–shaping** command is used when Frame Relay traffic shaping is configured to allow the router to react to BECN packets. For information on how the router adjusts traffic rates in response to BECNs, refer to Traffic Shaping.

Q. How can I improve performance over a slow Frame Relay link?

A. Poor performance over a Frame Relay link is generally caused by congestion on the Frame Relay network and from packets that are discarded while in transit. Many service providers only provide best effort delivery on traffic that exceeds the guaranteed rate. This means that when the network becomes congested, it discards traffic over the guaranteed rate. That action can cause poor performance.

Frame Relay traffic shaping allows traffic to be shaped to the available bandwidth. Traffic shaping is frequently used to avoid performance degradation caused by congestion packet loss. For a description of Frame Relay traffic shaping and configuration examples, refer to Frame Relay Traffic Shaping or the Frame Relay Traffic Shaping section of the Comprehensive Guide to Configuring and Troubleshooting Frame Relay.

To improve performance, refer to theConfiguring Payload Compression or Configuring TCP/IP Header Compression sections of Comprehensive Guide to Configuring and Troubleshooting Frame Relay.

Q. What is Enhanced Local Management Interface (ELMI) and how is it used for dynamic traffic shaping?

A. ELMI enables automated exchange of Frame Relay Quality of Service (QoS) parameter information between the Cisco router and the Cisco switch. Routers can base congestion management and prioritization decisions on known QoS values such as committed information rate (CIR), committed burst (Bc), and excess burst (Be). The router reads QoS values from the switch and can be configured to use those values in shaping traffic. This enhancement works between Cisco routers and Cisco switches (BPX/MGX and IGX

platforms). Enable ELMI support on the router by issuing the **frame–relay qos–autosense** command. For information and configuration examples, refer to the Enabling Enhanced Local Management Interface section of the Configuring Frame Relay and Frame Relay Traffic Shaping.

Q. Can I reserve bandwidth for certain applications?

A. A recently developed Cisco feature called Class–Based Weighted Fair Queuing (CBWFQ) allows reserved bandwidth for different applications of flows depending on Access Control List (ACL) or incoming interfaces. For configuration details, refer to Configuring Weighted Fair Queueing.

Q. Can I use priority queuing with Transmission Control Protocol (TCP) header compression over Frame Relay?

A. For the TCP header compression algorithm to function, packets must arrive in order. If packets arrive out of order, the reconstruction will appear to create regular TCP/IP packets but the packets will not match the original. Because priority queuing changes the order in which packets are transmitted, enabling priority queuing on the interface is not recommended.

Q. Can Frame Relay prioritize voice traffic carried in IP packets over non–voice packets?

A. Yes. The Frame Relay IP RTP Priority feature provides a strict priority queueing scheme on a Frame Relay private virtual circuit (PVC) for delay–sensitive data, such as voice, which is identified by its Real–Time Transport Protocol (RTP) port numbers. This feature makes sure that voice traffic is given strict priority over other non–voice traffic.

Q. What is Frame Relay private virtual circuit (PVC) Interface Priority Queueing (PIPQ)?

A. The Frame Relay PVC Interface Priority Queueing (PIPQ) feature provides interface level prioritization by giving priority to one PVC over another PVC on the same interface. This feature can also be used to prioritize voice traffic over non–voice traffic when they are carried on separate PVCs on the same interface.

Routing

Q. How is IP split horizon handled on Frame Relay interfaces?

A. IP split horizon checking is disabled by default for Frame Relay encapsulation to allow routing updates to go in and out of the same interface. An exception is the Enhanced Interior Gateway Routing Protocol (EIGRP) for which split horizon must be explicitly disabled.

Certain protocols such as AppleTalk, transparent bridging, and Internetwork Packet Exchange (IPX) cannot be supported on partially meshed networks because they require split horizon to be enabled (a packet received on an interface cannot be transmitted over the same interface, even if the packet is received and transmitted on different virtual circuits).

Configuring Frame Relay subinterfaces ensures that a single physical interface is treated as multiple virtual interfaces. This capability allows you to overcome split horizon rules so

packets received on one virtual interface can be forwarded to another virtual interface, even if they are configured on the same physical interface.

Q. Does Open Shortest Path First (OSPF) require additional configuration to run over Frame Relay?

A. OSPF treats multipoint Frame Relay interfaces as NON_BROADCAST by default. This requires that neighbors be explicitly configured. There are various methods for handling OSPF over Frame Relay. The one that is implemented depends upon whether the network is fully meshed. For more information, refer to the following documents:

♦ Initial Configurations for OSPF over Non−Broadcast Links
♦ Initial Configurations for OSPF over Frame Relay Subinterfaces
♦ Problems with Running OSPF in Mode over Frame Relay

Q. How can the bandwidth consumed by routing updates over Frame Relay be calculated?

A. Reliable estimates can only be calculated for distance vector protocols that send periodic updates. This includes Routing Information Protocol (RIP) and Interior Gateway Routing Protocol (IGRP) for IP, RIP for Internetwork Packet Exchange (IPX), and Routing Table Maintenance Protocol (RTMP) for AppleTalk. A discussion of the bandwidth consumed by these protocols over Frame Relay can be found in the RIP and IGRP section of Configuring and Troubleshooting Frame Relay.

Simple Network Management Protocol (SNMP)

Q. I can issue a Simple Network Management Protocol (SNMP) ping to the router asking it to ping all data−link connection identifier (DLCI) partners, and it is successful. What does this indicate?

A. This confirms that the protocol is configured and the protocol−to−DLCI mapping is correct at both ends.

Q. Are Simple Network Management Protocol (SNMP) variables available that can provide an accurate status on the data−link connection identifiers (DLCIs)?

A. Yes. The variables are found in RFC1315 and the Frame Relay data terminal ready (DTR) management information base.

The SNMP variable for a circuit's status is **fr** *CircuitState*. Its Abstract Syntax Notation One (ASN.1) object identifier (OID) form is 1.3.6.1.2.1.10.32.2.1.3. It resides in the frCircuitTable. To get the value (the status in this case), the index and the DLCI would be the first and second instance respectively. By issuing the **SNMP Get** or **Getnext** commands, you can find out the system's internal circuit status. The following table lists valid values:

Value	State
1	invalid

2	active
3	inactive

For Cisco, you would see either 2 or 3.

Related Information

- **Frame Relay Technology Support Page**
- **Technical Support – Cisco Systems**

Updated: Nov 17, 2007   Document ID: 14168