Packet Lab

Frame Relay
Introduction and Concepts

Frame Relay

Frame Relay is a high-performance packet-switched WAN protocol. It operates at Layer 2 (Data Link) of the OSI model. It is a high-speed WAN protocol designed primarily for interconnecting geographically dispersed LANs or WANs through a shared network. Frame Relay is a Layer 2 protocol and can be used to transport a variety of routed protocols. These protocols' (we will concentrate only on IP) packets are encapsulated in frame when being sent over the Frame Relay network. So when frames are traversing the Frame Relay network, they are switched based on the Frame Relay information. When the frames eventually reach their destination, the Frame Relay headers are stripped off and the IP payload is passed on to the upper layers for processing.

Frame Relay puts data in variable-size units called "frames" and leaves any necessary error-correction (such as re-transmission of data) up to the end-points (although it does check basic packet consistency via a Frame Check Sequence (FCS)). This speeds up overall data transmission. When a Frame Relay network detects an error in a frame, it simply drops that frame. The end points have the responsibility for detecting and retransmitting dropped frames.

Frame Relay – OSI Model

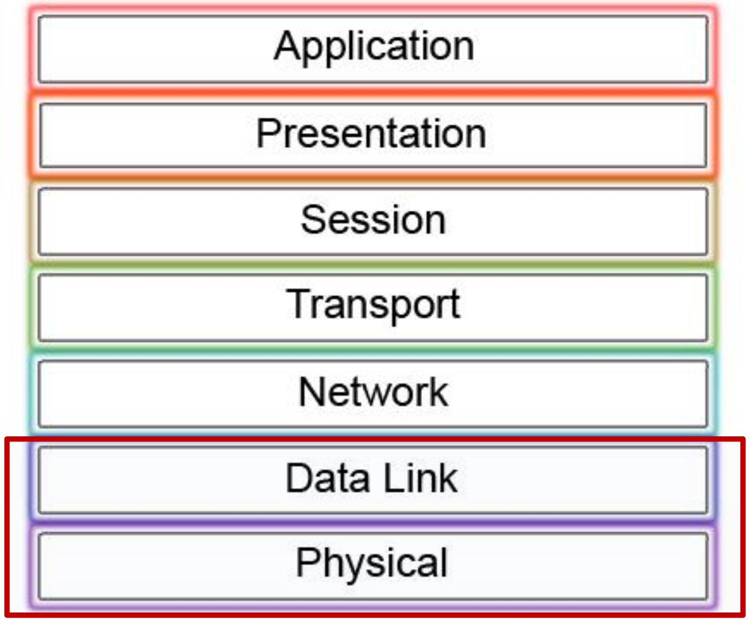

Frame Relay History

Frame Relay arose during the era of X.25 Frame Relay was first developed in 1984, but did not start to take off until the late 1980s.

In 1990 when Cisco, Digital Equipment Corporation (DEC), Northern Telecom, and StrataCom (the "Gang of Four") formed a consortium to focus on Frame Relay technology development. This consortium extended the protocol with features that provide additional capabilities for complex internetworking environments. These Frame Relay extensions are referred to collectively as the Local Management Interface (LMI). We'll look at LMI in more detail later on in this lesson.

Packet Switching

Frame Relay is an example of a packet-switched technology. Packet-switched networks enable end stations to dynamically share the network medium and the available bandwidth. The following two techniques are used in packet-switching technology:

- Variable-length packets
- Statistical multiplexing

Unlike technologies such as ATM which use a fixed packet size, Frame Relay uses variable-length packets for more efficient and flexible data transfers. These packets are switched between the various segments in the network until the destination is reached.

Frame Relay uses statistical multiplexing, which allows multiple users to access a Frame Relay virtual circuit and port from a single location. In statistical multiplexing, a communication channel is divided into an arbitrary number of variable bit-rate digital channels or data streams. Statistical multiplexing allows the bandwidth to be divided arbitrarily among a variable number of channels and ensures that slots will not be wasted. Statistical multiplexing normally implies "on-demand" service rather than one that preallocates resources for each data stream.

Frame Relay Devices

Devices attached to a Frame Relay WAN fall into the following two general categories:

- **Data terminal equipment (DTE)** - generally a router
- **Data circuit-terminating equipment (DCE)*** – generally a Frame Relay switch

DTEs generally are considered to be terminating equipment for a specific network and typically are located on the premises of a customer. In fact, they may be owned by the customer. Examples of DTE devices are terminals, personal computers, routers, and bridges.

DCEs are carrier-owned internetworking devices. The purpose of DCE equipment is to provide clocking and switching services in a network, which are the devices that actually transmit data through the WAN. In most cases, these are packet switches.

*Also referred to as **data communications equipment** and **data carrier equipment**.

Packet Lab

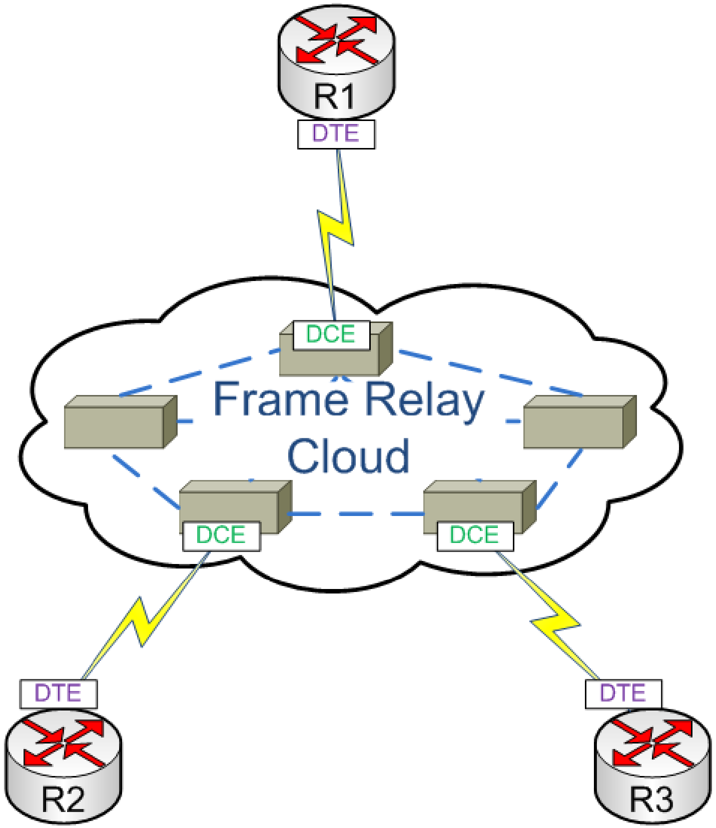

Frame Relay Cloud

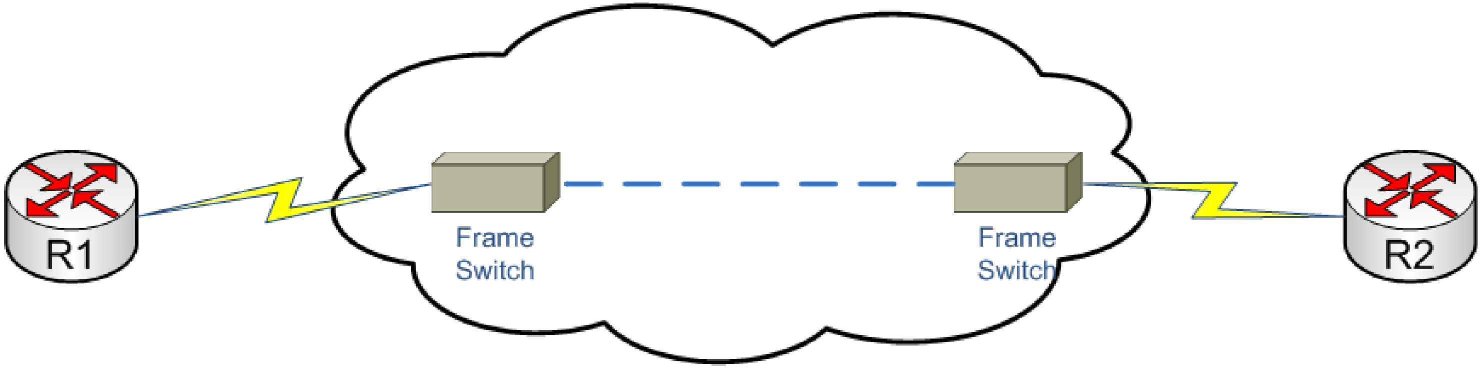

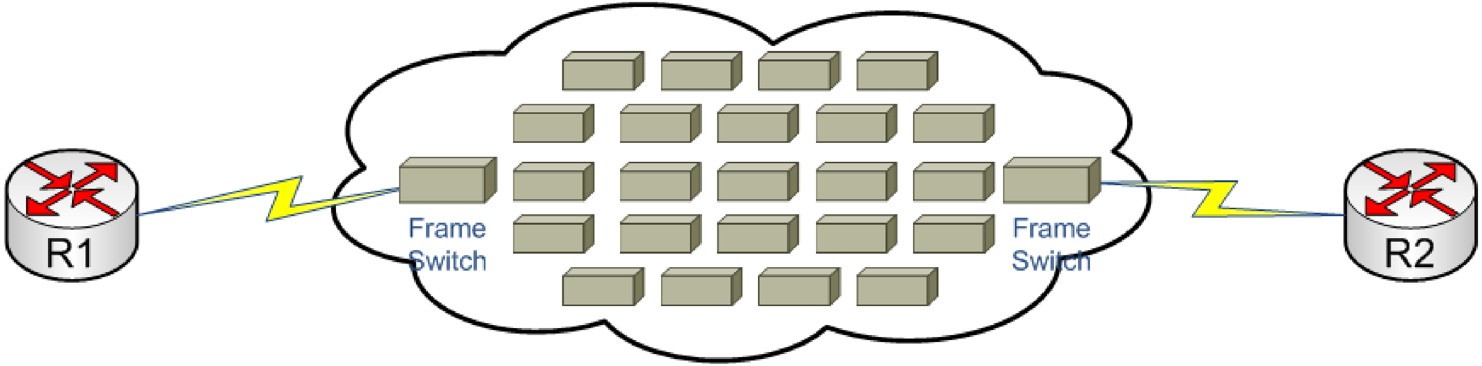

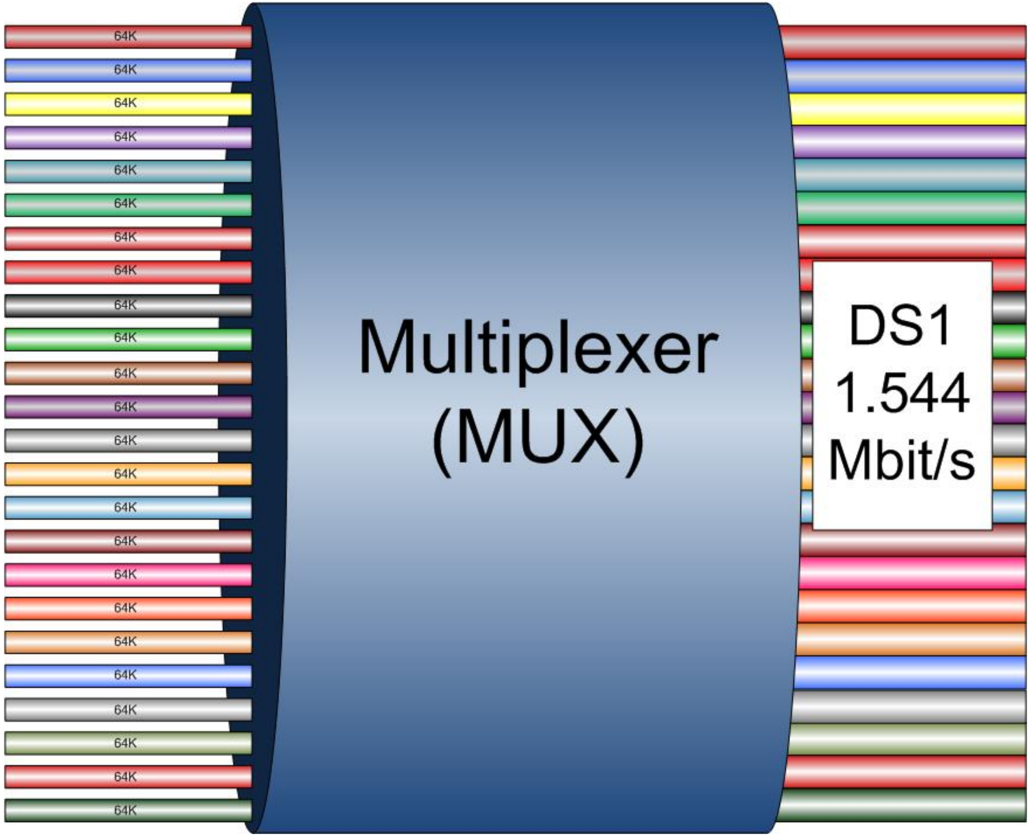
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
64K
Multiplexer
(MUX)
DS1
1.544
Mbit/s

Frame Relay Virtual Circuits

Frame Relay provides connection-oriented data link layer communication. This means that a defined communication exists between each pair of devices and that these connections are associated with a connection identifier. This service is implemented by using a Frame Relay virtual circuit, which is a logical connection created between two data terminal equipment (DTE - generally two routers) devices across a Frame Relay packet-switched network (PSN).

Virtual circuits provide a bidirectional communication path from one DTE device to another and are uniquely identified by a data-link connection identifier (DLCI). A number of virtual circuits can be multiplexed into a single physical circuit for transmission across the network. This capability often can reduce the equipment and network complexity required to connect multiple DTE devices.

A virtual circuit can pass through any number of intermediate DCE devices (switches) located within the Frame Relay PSN.

Frame Relay virtual circuits fall into two categories: **switched virtual circuits (SVCs)** and **permanent virtual circuits (PVCs).**

Switched Virtual Circuit (SVC)

We'll begin by looking at Switched Virtual Circuits (SVCs). You are unlikely to ever run across an SVC in the wild, but it's good to learn about them, if only to compare them with PVCs.

Switched virtual circuits (SVCs) are temporary connections used in situations requiring only sporadic data transfer between DTE devices across the Frame Relay network. SVCs consist of the following four operational states:

- **Call setup** - The virtual circuit between two Frame Relay DTE devices is established.
- **Data transfer** - Data is transmitted between the DTE devices over the virtual circuit.
- **Idle** - The connection between DTE devices is still active, but no data is transferred. If an SVC remains in an idle state for a defined period of time, the call can be terminated.
- **Call termination** - The virtual circuit between DTE devices is terminated.

The advantage of an SVC is that it would only be used when you had data to send. The downside is that you have a lot of circuit setup and teardown which makes this a more complicated option than PVCs. SVCs are similar to an on-demand ISDN circuit. In practice, SVCs are pretty rare (I've never seen on in the wild).

Permanent Virtual Circuit (PVC)

As you can probably guess from the name, permanent virtual circuits (PVCs) are permanently established connections that are used for frequent and consistent data transfers between DTE devices across the Frame Relay network. Communication across a PVC does not require the call setup and termination states that are used with SVCs. PVCs always operate in one of the following two operational states:

- **Data transfer** - Data is transmitted between the DTE devices over the virtual circuit.
- **Idle** - The connection between DTE devices is active, but no data is transferred. Unlike SVCs, PVCs will not be terminated under any circumstances when in an idle state.

The overwhelming majority of Frame Relay virtual circuits will be PVCs. You can begin transferring data whenever you are ready because the circuit is permanently established.

Data Link Control Identifier (DLCI)

Frame Relay virtual circuits are identified by data-link connection identifiers (DLCIs). DLCI (often pronounced “dell-see”) values are generally assigned by the Frame Relay service provider. Frame Relay DLCIs have local significance. A DLCI is stored in a 10-bit (actually two fields, one 6-bit and the other 4-bit) in the Frame Relay header. 2^10 = 1,024 possible DLCI values, of which DLCIs 16 - 991 are available for use. The following DLCIs are reserved:

DLCIs 0 and 1023 are reserved for management (LMI).
DLCIs 1 - 15 and 1008 - 1022 have been reserved for future use.
DLCIs 992 - 1007 are reserved for various purposes.

DLCIs are Layer 2 addresses and can be though of as analogs to a MAC address in Ethernet. Unlike Ethernet, Frame Relay does not offer a dynamic mapping of Layer 2 addresses to Layer 3 addresses like ARP does in Ethernet [There is a Frame Relay Inverse-ARP that works in a similar manner though]. In many situations, it will be necessary or desirable to manually map the Layer 2 DLCI to a Layer 3 IP address.

DLCI Assignment

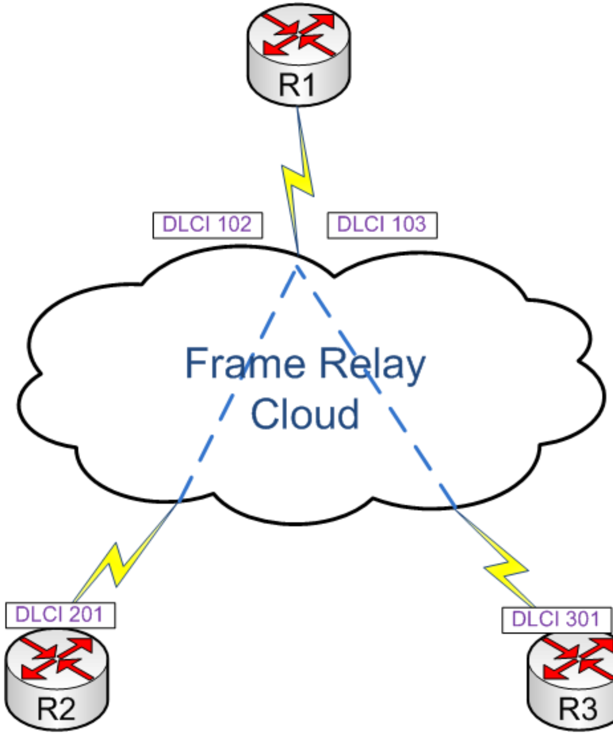

DLCI Assignment

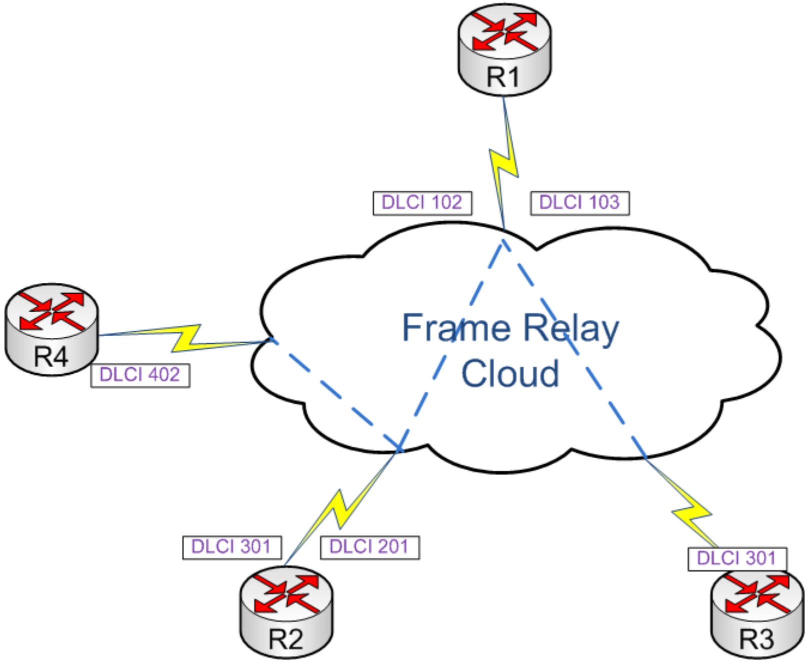

Frame Relay Header

Frame Relay Packet

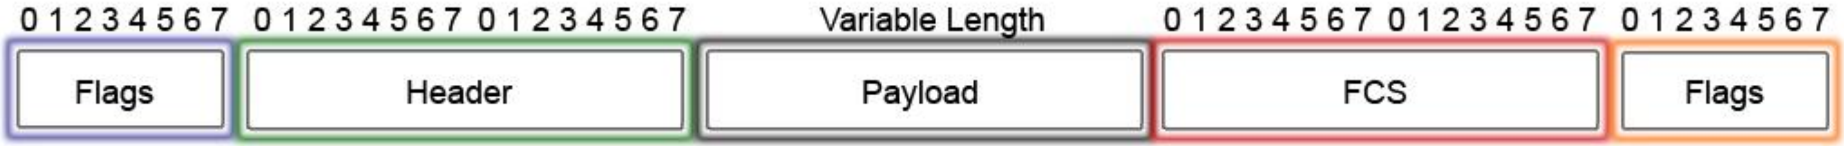

Frame Relay Header

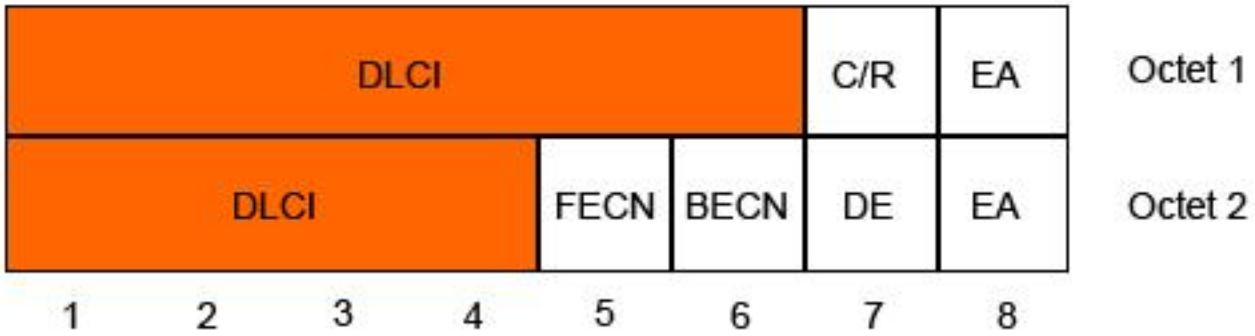

Local Management Interface (LMI)

Local Management Interface (LMI) is a signaling standard used between routers and frame relay switches. Information about keepalives, global addressing, IP Multicast and the status of virtual circuits is commonly exchanged using LMI. There are three forms of LMI:

Cisco - "Gang of Four" standard; uses **DLCI 1023**
ANSI - T1.617 Annex D standard; uses **DLCI 0**
q933a - Q.933 Annex A standard uses **DLCI 0**

The DCE device (Frame Relay switch) determines the type of LMI used. You will need to set your LMI type to match that of the Frame Relay switch. Luckily, Cisco devices use LMI Autosense which automatically determines the LMI type used by the Frame Relay switch and configures the router to match it.

There are two types of LMI messages:

- **Status-inquiry message** - Allows a user device to inquire about the status of the network.
- **Status message** - Responds to status-inquiry messages. Status messages include keepalives and PVC status messages.

No.	Time	Source	Destination	Protocol	Info
2	0.001000			Q.933	STATUS

```
⊟ Frame 2 (14 bytes on wire, 14 bytes captured)
    Arrival Time: Jul 11, 2010 13:18:43.396250000
    [Time delta from previous captured frame: 0.001000000 seconds]
    [Time delta from previous displayed frame: 0.001000000 seconds]
    [Time since reference or first frame: 0.001000000 seconds]
    Frame Number: 2
    Frame Length: 14 bytes
    Capture Length: 14 bytes
    [Frame is marked: False]
    [Protocols in frame: fr:q933]
⊟ Frame Relay
  ⊞ First address octet: 0x00
  ⊞ Second address octet: 0x01
    DLCI: 0
  ⊟ Control field: U, func=UI (0x03)
      000. 00.. = Command: Unnumbered Information (0x00)
      .... ..11 = Frame type: Unnumbered frame (0x03)
⊟ Q.933
    Protocol discriminator: Q.933
    Call reference value length: 0
    Message type: STATUS (0x7d)
    Locking shift to codeset 5: Information elements for national use
  ⊞ Report type (ANSI)
  ⊞ Keep Alive (ANSI)
```

Additional Frame Relay Features

We will go into these features (and others) in much more detail in future lessons, but two very useful features of Frame Relay are traffic congestion notification as well as oversubscription/traffic shaping.

Frame Relay includes traffic congestion notification mechanisms (FECN and BECN – pronounced “fecken” and “becken”) which report on congestion in either direction in the Frame Relay cloud. There is also a Discard Eligible (DE) bit that can be set to prioritize certain packets to for discarding (sounds odd, but is actually really useful) if the Frame Relay VC becomes congested.

The DE bit is also used when Frame Relay Traffic Shaping (FRTS) and/or oversubscription is used. Oversubscription in Frame Relay allows a customer to lease a VC at a certain, guaranteed bandwidth (Committed Information Rate (CIR)) and (depending on certain conditions) to exceed that guaranteed rate if additional bandwidth is available. It’s kind of like borrowing bandwidth that no one else is using.

Frame Relay Traffic Shaping (FRTS)

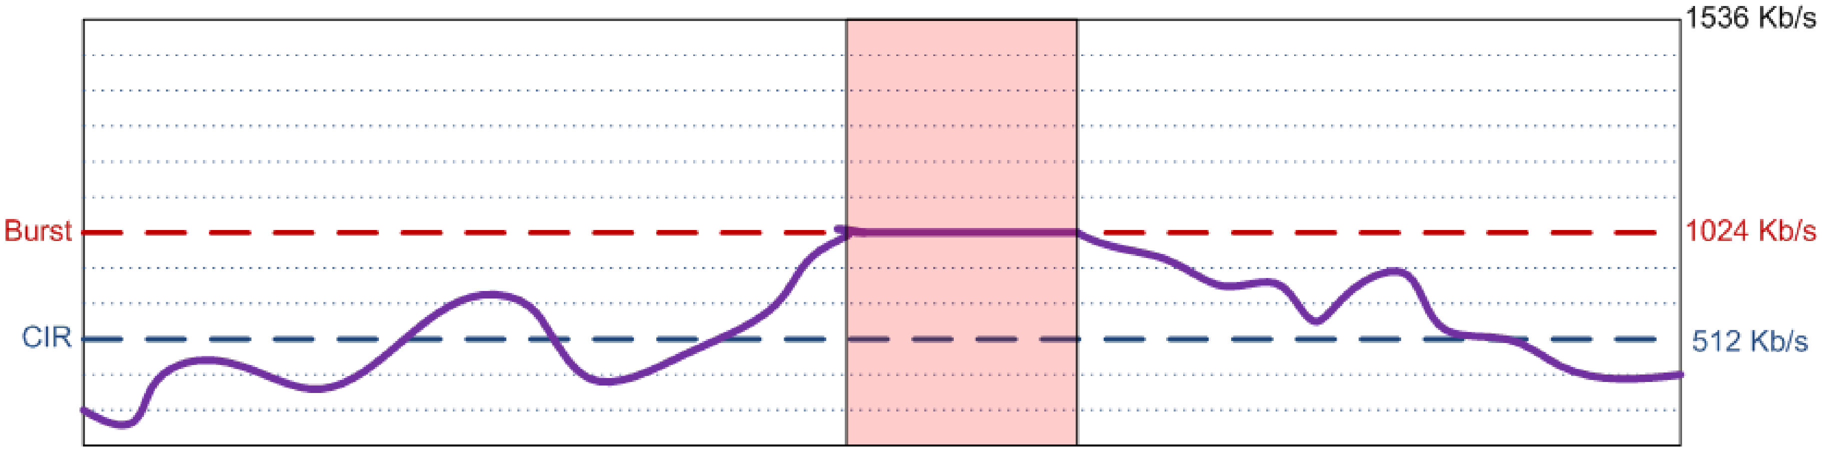

Summary

Frame Relay is a high-performance packet-switched WAN protocol which uses statistical multiplexing and variable length packets. Frame Relay operates at Layer 2 (Data Link) of the OSI model. Each customer gets a private line (or leased line) to a frame-relay node called a virtual circuit (VC) through the Frame Relay cloud. These VCs can either be a bandwidth-on-demand Switched Virtual Circuit (SVC) or (more likely) an always-on Permanent Virtual Circuit (PVC).

DTE devices such as a router communicate with a DCE Frame Relay switch. Local Management Interface (LMI) is used to get VC information and as a keepalive between the router and the Frame Relay switch. Each Frame Relay node (DTE) uses a 10-bit data-link connection identifier (DLCI) as a Layer 2 address. These DLCIs are locally significant (although there are extensions that can make them globally significant) and are analogous to MAC addresses in Ethernet.

Frame Relay has many powerful features including built-in congestion notification as well as traffic shaping (QoS) capabilities.

Although Frame Relay is getting long in the tooth (it's been around since the late 1980s) and technologies like MPLS are quickly rising in popularity, it is still a vary popular WAN protocol and once which you are very likely to work with in the field.